

FITTING WORDS: VAGUE LANGUAGE IN CONTEXT

1. INTRODUCTION

Consider the possibilities for reference using vague definite descriptions. There are two pigs in a pen: the first a runt, the second quite round but not really fat. Now the farmer says, “The fat pig won a prize”. His neighbor could protest that neither pig is really fat. But he does not protest. He understands that the rounder pig is meant, and both of them know it.

We can chalk this up to negotiation about how to use vague words. The farmer, by using the definite description, implicitly proposes to stretch the sense of ‘fat’ to cover one of the two pigs. This looser sense is needed in order to satisfy the presupposition that there is a single fat pig. Now, given the pigs’ sizes, there is really just one way to stretch. The rounder of the two will be the one that is covered, and so the definite description will come to pick out this pig. In this way, using ‘fat’ more loosely can help to distinguish between the pigs. That is why the farmer makes his proposal. When the neighbor does not protest he tacitly accepts the proposal, and lets the farmer know that he has understood.

Of course, such agreements about how to use vague words are bound to the particular contexts in which they are reached. If the rounder pig were instead among some truly fat pigs and the farmer said, “The pig that is not fat won a prize”, then the neighbor would understand that he meant the same pig. In such a context they would settle on a stricter sense of ‘fat’, not a looser one.

This contextuality and the *accommodation* that accompanies it are in our view characteristic of vague language.¹ Just as a home handyman can fit an adjustable wrench to a nut, we think, a speaker can adjust the extension of a vague expression to suit his needs, relying on the hearer to recognize his intentions and to accommodate him.²

¹ We borrow the term “accommodation” from David Lewis (1979).

² The contextuality of vague language has been pointed out by others, e.g., Manfred Pinkal (1983, 1995) and Ruth Manor (1995).



Maybe it is possible to account for examples like these without supposing a change in the senses of vague expressions. Keith Donnellan describes the case of a usurper to the throne: all refer to him somehow insincerely, but understandably enough, as “the king”, although it is common knowledge that he is not a king.³ Perhaps it is the same with the pigs. The farmer *calls* the rounder pig “fat”, though really it is borderline.

No doubt ‘fat pig’ can be used in this somehow insincere way. But the cases we have in mind differ from Donnellan’s example in two important ways. First, there are other individuals who are as entitled to be called “king”, or as lacking in entitlement, as our usurper. But whereas calling the usurper “king” does not bring with it any inclination, or obligation, also to call these other non-kings “king”, things are different with the pigs. Having called one a “fat pig”, it seems odd, even inconsistent, to withhold this description from comparable pigs.

The case of the usurper differs in another way too. In referring to the usurper as “the king”, a speaker has in mind some particular person. The description is used referentially, according to Donnellan. In contrast, in some of our examples the speaker need not have a particular referent in mind. A farmer might say, “The fat pigs are going to market”, even though he has not yet decided which are to be sent off. A hearer who understands this can nonetheless accommodate by agreeing to call pigs as round as the roundest ones “fat pigs”.

It has been suggested that the mechanism of the accommodation is in the compositional analysis of adjective-noun combinations. The import of an adjective such as ‘fat’ depends on a contextually determinate comparison class. That is how it can be that one and the same lady is a fat fashion model, but a skinny belly dancer. We think of accommodation as stretching the senses of vague expressions while keeping comparison classes fixed. But an obvious alternative is to think of it as switching comparison classes. Besides the noun phrases with which adjectives appear, comparison classes can depend on such contextual factors as salient groups of individuals. In our example, the rounder pig is not fat relative to pigs at large, but it is fat relative to the pigs in the pen. Accommodation might be switching from the one comparison class to the other.

We claim that the accommodation we have in mind can be accounted for without bringing in comparison classes. We do not have a conclusive empirical reason to leave them out, but we do have a theoretical reason. We expect that modeling accommodation as comparison-class change will require multiplying comparison classes in an *ad-hoc* manner. For most of this paper we shall hide comparison classes from view, to see how far we

³ (Donnellan, 1966).

can get without them. But we return briefly to this question in Section 6, with an example suggesting that they are needed anyway.

We shall proceed on the assumption that communication proceeds relative to a context, which includes shared presuppositions. Among these are items of general knowledge, along with whatever the situation makes obvious and whatever has been established in the course of the conversation, whether it is true or false. Often, perhaps normally, participants will *believe* these presuppositions, but they need not. Assumptions made for the sake of the argument are presuppositions in the relevant sense. So, perhaps, are propositions that children “make believe” in games. Following Robert Stalnaker, we shall call the body of these presuppositions the *common ground*.⁴ In keeping with a computational view of mind, we suppose that the common ground is represented, in the minds of the participants in a conversation, by sentences of some language. Since these presuppositions are to be shared, though, we think of the common ground metaphorically as a blackboard, visible to all participants, on which the presuppositions are written.

Utterances bring about changes in the common ground. Our purpose here is to propose a model of changes made in order to accommodate vague utterances. Our model is, in the sense of David Marr (1982), “computational”: we shall present a function for updating the common ground, but our object is not in the first place to present an algorithm that a language-understanding device might use. Rather, our object is to present a function we think such an algorithm ought to compute. With this goal in mind, we put to one side questions concerning the computational complexity of the accommodation function.

The model combines ideas from three domains: supervaluational accounts of vagueness, theories of conveyed meanings, and theories of belief revision. Supervaluational accounts of vagueness were suggested by, among others, Kit Fine (1975) and Hans Kamp (1975). From them we borrow an interpretation of vague language. By a theory of conveyed meanings we mean a specification of the presuppositions, conversational implicatures, and other meanings that are conveyed by utterances. Pieces of such a theory can be found in the work of Paul Grice (1967), and in subsequent work by Lauri Karttunen and Stanley Peters (1975, 1979) and Gerald Gazdar (1979a, 1979b), among others.⁵ We shall not present such an account, but our discussion shall take place against the background of some such theory, the existence of which we presuppose. Finally, one well-known account of belief revision is that of Carlos Alchourrón et al. (1985).

⁴ (Stalnaker, 1974).

⁵ For a survey see (Levinson, 1983) and, for more recent work, (Beaver, 1997).

The function that we shall propose for updating common ground derives directly from it.

In the next section, we shall describe a language for representing expressions of English. In Section 3, we shall briefly discuss some assumptions we make about a theory of conveyed meanings. With this background in place, in Section 4, we shall propose a model of accommodation, on which uttering a sentence with a vague description leads to changes in the common ground. We shall illustrate this accommodation in Section 5. Finally, in Section 6, we shall discuss several examples that point toward elaborations of our model.

2. VAGUE LANGUAGE

It will help to have a formal language in which to represent expressions of English. The language we have chosen is very limited in its expressive power; but it has the advantages for our purpose of its simplicity and familiarity: \mathcal{L} is a language of predicate logic augmented with the sentential operator D , with the intended interpretation “It is definitely the case that”. \mathcal{L} has individual constant symbols including *Arnold* and *Babe*, individual variable symbols x, y, \dots , and set variable symbols X, Y, \dots ; \mathcal{L} has monadic predicates including *pig*, *fat-pig*, *skinny-pig*, *winner*, and *curly-tail*; and dyadic predicates including *tail*, *at-least-as-fat-as*, and *fatter*. Some of these expressions – they include *pig*, *winner*, *tail*, *at-least-as-fat-as*, and *fatter* – are designated non-vague. The others – including *fat-pig*, *skinny-pig*, and *curly-tail* – are vague.⁶ The formulas of \mathcal{L} are defined with the usual recursion over the logical symbols.

To illustrate, the intended interpretation of $\text{fat-pig}(\text{Arnold})$ is that Arnold is a fat pig. The intended interpretations of $\text{at-least-as-fat-as}(\text{Arnold}, \text{Babe})$ and $\text{fatter}(\text{Arnold}, \text{Babe})$ are that Arnold is at least as fat as Babe, and that Arnold is fatter than Babe. The intended interpretation of $D(\text{skinny-pig}(\text{Arnold}))$ is that it is definitely the case that Arnold is a skinny pig, and the intended interpretation of $\neg D(\text{fat-pig}(\text{Arnold})) \ \&\neg D(\neg \text{fat-pig}(\text{Arnold}))$ is that there is no matter of fact about whether Arnold is a fat pig, which is to say that Arnold is a borderline fat pig.

Notice that we have not included in \mathcal{L} representations of vague adjectives like ‘fat’, which can be used to form vague composite expressions

⁶ We simplify in two ways. First, we consider only the vagueness of predicates, putting to one side questions relating to the vagueness of names and quantifiers. Second, we consider only some predicates to be vague, though just about all of them seem to be, to a greater or lesser extent.

like ‘fat pig’. We will treat such composite expressions as if they were semantically simple. A deeper treatment would interpret adjectives as functions mapping the semantic values of the nouns they combine with onto the semantic values of the resulting adjective-noun combinations. The semantic value of the noun could be thought of as picking out a “comparison class” relative to which the adjective is interpreted. An account of logical relations that hold between sentences in virtue of the adjective noun pairs that compose them would also be a part of this deeper treatment. For now, though, we shall set these matters to one side.

We turn now to the interpretation of \mathcal{L} . Imagine a pig that is on the threshold of becoming fat. Imagine also that you are well informed about the size of the pig – say, because you are looking straight at it. Now someone says that the pig is fat and someone else disagrees. In a case like this you may find yourself unable to agree with either of them. One possible explanation for this inability is that although there is a fact of the matter about whether or not the pig is fat, you are insufficiently informed about the extension of the word ‘fat’, and so you cannot tell. Either the pig in question is fat or it is not, but you just do not, and perhaps cannot, know which.⁷ Another explanation for this inability, which coheres well with the approach we take, is that there are some pigs that are neither definitely fat, nor definitely not fat, and that this pig is one of them.

This second explanation suggests a *partial* interpretation for vague predicates. A partial interpretation assigns to each predicate both a positive and a negative extension, which need not exhaust the entire domain of quantification. A *borderline* individual is one that is neither in the positive nor in the negative extension of the predicate in question.

The notion that vagueness is semantic partiality is one of the main ideas underlying supervaluational treatments of vagueness. The other is the idea that vagueness can be reduced by resolving borderline cases one way or the other, while leaving determinate cases be. The different ways of making vague expressions more precise are *precisifications*.

We shall use precisifications to explore the idea that accommodation involves the adoption of new senses of vague expressions. These senses might incorporate formerly borderline individuals either into the positive extensions of the expressions in question, or into their negative extensions. Or, they might exclude individuals from the extensions in question. The newly adopted senses correspond to new precisifications.

Precisifications are constrained by the intuitively understood meanings of expressions. Some of these constraints, following Fine (1975), we shall call *penumbral* constraints. For instance, the intuitive meanings of ‘fat pig’

⁷ Epistemic treatments are recommended by (Sorenson, 1988) and (Williamson, 1994).

and ‘skinny pig’ require that, on any *admissible* precisification, if a pig is at least as fat as a fat pig, then it too is fat, and that no skinny pig is fat. In addition, precisifications satisfy constraints placed on them by the meanings of non-vague expressions. The intuitive meaning of ‘at least as fat as’ requires that, on any precisification, given any two pigs, one is at least as fat as the other. Furthermore, on any precisification, one pig is fatter than another just in case it is at least as fat as the other, but not the other way around.

We evaluate a sentence at a precisification that is *appropriate* insofar as it retains a degree of indeterminacy that reflects the intuitively understood meanings of vague expressions. But the truth value of a sentence at this precisification depends, according to a supervaluational treatment, on the truth values it will obtain on the different ways of making it completely precise, i.e., on its *complete* precisifications. A sentence is counted true if it is true on each of its complete precisifications, and false if it is false on each of them. A sentence that is true on some but false on others has no truth value.

We shall now make these ideas precise. A *model* \mathcal{M} for \mathcal{L} is a quadruple $\langle \mathcal{U}, \mathcal{P}, \leq, \mathcal{I} \rangle$. Here \mathcal{U} is a non-empty set, the domain of quantification; \mathcal{P} is a set of evaluation points; \leq is a partial ordering on \mathcal{P} ; and \mathcal{I} is an interpretation function. We require that \leq orders \mathcal{P} as a tree, with a unique minimal element $@$, the *appropriate point*, at the base of the trunk. \mathcal{I} assigns to each n -place predicate or relation symbol R and evaluation point p two disjoint sets of n -tuples of elements of \mathcal{U} : $\mathcal{I}^+(R, p)$, the *positive extension* of R at p ; and $\mathcal{I}^-(R, p)$, the *negative extension* of R at p . \mathcal{I} must respect \leq : for each n -place relation R and all points $p, q \in \mathcal{P}$, if $p \leq q$ then $\mathcal{I}^+(R, p) \subseteq \mathcal{I}^+(R, q)$ and $\mathcal{I}^-(R, p) \subseteq \mathcal{I}^-(R, q)$. A *complete point* is a point $p \in \mathcal{P}$ such that for each n -place relation R , $\mathcal{I}^+(R, p) \cup \mathcal{I}^-(R, p)$ is the set of all n -tuples of elements of \mathcal{U} . The appropriate precisification is that which is induced by \mathcal{I} at the appropriate point, and complete precisifications are those induced by \mathcal{I} at complete points. \mathcal{I} assigns to each individual constant of \mathcal{L} some element of \mathcal{U} .

Evaluation at a complete point is classical. For any assignment a of the individual variables of \mathcal{L} into the domain of \mathcal{M} , and for any individual term t , let $[t]_{\mathcal{M}, a}$ be $\mathcal{I}(t)$ if t is a constant, and $a(t)$ if t is a variable. If ϕ is an atomic formula Rt_1, \dots, t_n and p is a complete point, then $[\phi]_{\mathcal{M}, p, a} = \text{true}$ if and only if $\langle [t_1]_{\mathcal{M}, a}, \dots, [t_n]_{\mathcal{M}, a} \rangle \in \mathcal{I}_{\mathcal{M}}^+(R, p)$. Otherwise, $[\phi]_{\mathcal{M}, p, a} = \text{false}$. Formulas composed with the classical connectives and quantifiers obtain truth values in accordance with the classical truth functions.

At any point p that is not complete, the satisfaction value of a formula is supervaluational: $[\phi]_{\mathcal{M}, p, a} = \text{true}$ [*false*] if and only if for all complete

points $q \in \mathcal{P}$ such that $p \leq q$, $[\phi]_{M,q,a} = \text{true}$ [*false*]. We say a sentence ϕ is *true in M at p*, relative to a , just in case $[\phi]_{M,p,a} = \text{true}$; ϕ is *false in M at p* just in case $[\phi]_{M,p,a} = \text{false}$; otherwise, ϕ is *undefined*. Finally, truth in a model goes by the appropriate precisification: a sentence ϕ is *true* [*false/undefined*] in \mathcal{M} just in case ϕ is *true* [*false/undefined*] in \mathcal{M} at $@$.

To complete the truth definition, we now consider expressions of the form $D(\phi)$. Informally speaking, it is definitely the case that ϕ if on all ways of making ϕ precise this sentence is true. Accordingly, we let the value of $D(\phi)$ at any point p depend on the value of ϕ at the appropriate point: $[D(\phi)]_{M,p,a} = \text{true}$ if $[\phi]_{M,@,a} = \text{true}$; otherwise, $[D(\phi)]_{M,p,a} = \text{false}$. Writing ' $I(\phi)$ ' as an abbreviation for ' $\neg D(\phi) \& \neg D(\neg\phi)$ ', it follows that $[I(\phi)]_{M,p,a} = \text{true}$ if there are complete points q and r such that $[\phi]_{M,q,a} = \text{true}$ and $[\phi]_{M,r,a} = \text{false}$; otherwise $[I(\phi)]_{M,p,a} = \text{false}$.

Note that the truth of $I(\phi)$ at a complete point p is compatible with the truth of ϕ there. This might seem odd at first, but it is as it should be. Where ϕ does not involve the operator D , the truth of ϕ at a complete point is a local matter, dependent only on the extensions of predicates at that point. But the truth of $I(\phi)$ is a global matter. It means, informally speaking, that ϕ is a sentence whose truth can go either way; on some ways of making its predicates precise it is true, on others it is false.⁸

We will be interested only in *acceptable* models. A model \mathcal{M} is acceptable if for all points p and q : (1) for each non-vague predicate or relation symbol R , $\mathcal{I}(R, p) = \mathcal{I}(R, q)$; and (2) \mathcal{I} induces an admissible precisification at p . The first constraint ensures that the interpretation of non-vague language does not vary from point to point within a model. From the second follows the logical validity, or truth in every model, of such *lexical truths* as:

$$\forall x, y (\text{at-least-as-fat-as}(x, y) \rightarrow (\text{fat-pig}(y) \rightarrow \text{fat-pig}(x))),$$

and

$$\forall x (\text{skinny-pig}(x) \rightarrow \neg \text{fat-pig}(x)).$$

⁸ For more discussion of the interpretation of ' $I(\phi)$ ' see (Fine, 1975), pp. 287–289; (Williamson, 1994), pp. 149–153; and (Morreau, 1999). This truth definition ignores the *higher order* vagueness that exists where borderlines themselves have borderlines. This simplification does not affect any of our main points.

Our models give rise to the following notion of entailment: a set of sentences Γ entails a sentence ϕ , written $\Gamma \models \phi$, if and only if ϕ is true at every point at which each member of Γ is true.⁹

We can now extend the notion of vagueness from simple expressions to sentences. A sentence ϕ is *vague* if and only if for some acceptable model \mathcal{M} and points p and q in $\mathcal{P}_{\mathcal{M}}$, it is not the case that $[\phi]_{\mathcal{M},p} = [\phi]_{\mathcal{M},q}$. Not all sentences with vague symbols are vague; no lexical truth is, for instance. But it can easily be shown that containing vague symbols is a necessary condition for vagueness. It also follows that if we exclude higher order vagueness then the result of placing a vague sentence within the scope of D or I is a non-vague sentence.¹⁰

3. CONVEYED MEANINGS

A speaker often conveys more than he has literally said. We use ‘*conveyed meanings*’ as a general term for the contribution made by literal meanings, presuppositions, and conversational implicatures to what an utterance communicates.

Consider definite descriptions. When someone says a sentence with a definite description, his utterance will in many cases presuppose that there is exactly one (salient) object that fits that description. That is why, if he is not sure whether there is a pig but is sure there is none in the pen, it is misleading for him to say “The pig is not in the pen”. Instead he ought to say “The pig is not in the pen, if there is a pig”. By saying the former sentence, but not by saying the latter, he conveys something he does not believe to be true: that there is a pig.

Our account of accommodation rests on a theory of conveyed meanings, but we shall not assume any particular theory. For some time philosophers of language and linguists have worked to provide such an account, and we shall proceed on the assumption that one can be given. Partial accounts that fill the bill can be found in (Gazdar, 1979a, 1979b), (Karttunen & Peters, 1975, 1979), (Levinson, 1983) and, more recently, (Beaver, 1997). Allowing others to fill in the details, we shall write ‘ $Cnv(s,c)$ ’ to indicate the set of conveyed meanings of an utterance of the English sentence s in a context c ; here, $Cnv(s,c)$ is understood to be a set of sentences of \mathcal{L} .

⁹ For a discussion of this notion of entailment see (Williamson, 1994), pp. 147–149. Williamson favors an alternative “global” notion: Γ entails ϕ if ϕ is true in every model in which Γ is true.

¹⁰ For a discussion of higher-order vagueness as it pertains to a supervaluational treatment of vagueness see (Fine, 1975), pp. 287–298, and (Williamson, 1994), pp. 156–161.

When we come to worked examples, in Section 5, we shall make assumptions about the function Cnv , making claims that hold for all accounts of conveyed meanings satisfying these assumptions.

The above example illustrates one of these assumptions, namely, that while an utterance with a definite description normally conveys that there is exactly one thing satisfying the description, sometimes it does not. Saying “The pig is not in the pen” is in many contexts enough to convey that there is just one pig. But things are different if you say instead “The pig is not in the pen, if there is a pig”. One conversational implicature of this utterance is that you believe neither that there is a pig nor that there is not one. In light of this implicature, the meaning that would normally be conveyed in virtue of the definite description is not conveyed.

In the next section we shall extend this assumption to sentences that contain *vague* definite descriptions. We shall assume, for instance, that the farmer’s utterance, “The fat pig won a prize”, conveys that there exists just one fat pig and that this pig won a prize. Adapting slightly the notation of Whitehead and Russell, we write ‘ $Q(\iota x Px)$ ’ to abbreviate:

$$\exists x (\forall y (P(y) \leftrightarrow y = x) \& Q(x)).^{11}$$

Our assumption is that for certain contexts c :

$$\begin{aligned} &winner(\iota x \text{ fat-pig}(x)) \\ &\in Cnv(\text{“The fat pig won a prize”}, c). \end{aligned}$$

Similarly, we shall assume that, in certain contexts, saying “The fat pigs won prizes” conveys that each pig in the most inclusive set of fat pigs won a prize. We write ‘ $Ps(X)$ ’ for:

$$\exists x, y (x \in X \& y \in X \& x \neq y \& \forall z (z \in X \rightarrow P(z))),$$

and we write: ‘ $Qs(\iota_{\max} X Ps(X))$ ’ for:

$$\exists X (Ps(X) \& \forall Y (Ps(Y) \rightarrow Y \subseteq X) \& Qs(X)).$$

Our assumption is that for certain contexts c :

$$\begin{aligned} &winners(\iota_{\max} X \text{ fat-pigs}(X)) \\ &\in Cnv(\text{“The fat pigs won prizes”}, c). \end{aligned}$$

¹¹ See (Whitehead & Russell, 1927), p. 173.

4. A MODEL OF ACCOMMODATION

What is conveyed by an utterance will often become a part of the common ground. Unless there is reason to think a speaker is mistaken (or lying or sarcastic or something of the sort), the conveyed meanings of his utterance will often simply be added to the common ground. For instance, when a speaker says, “The pig is not in the pen”, the common ground will ordinarily grow to include conveyed meanings, such as the presupposition that there is a unique pig. Let $c + s$ be the context brought about by uttering sentence s in a context c and let $g(c)$ be the common ground of c . A simple hypothesis about the effect of an utterance on the common ground is that:

$$g(c + s) = g(c) \cup Cnv(s, c).$$

This simple model accounts for many cases of accommodation. Consider again the case of the farmer who first shows his neighbor the not-quite-fat pig Arnold. Prior to the farmer’s utterance the two share the belief that Arnold is a borderline fat pig. This piece of the common ground is represented by the sentence $I(fat-pig(Arnold))$. Suppose the pig in the pen with Arnold is Babe, and that the common ground includes sentences expressing that Arnold and Babe are the only two pigs, that Arnold is the fatter of the two, and that Babe is a skinny pig. Then $g(c)$ includes the following sentences:

$$\begin{aligned} &I(fat-pig(Arnold)), \\ &\forall x(pig(x) \rightarrow x = Arnold \vee x = Babe), \\ &fatter(Arnold, Babe), \\ &skinny-pig(Babe). \end{aligned}$$

In any satisfactory account of conveyed meaning, the farmer’s utterance, “The fat pig won a prize”, will convey that there is a unique fat pig, rendered $\exists x \forall y (fat-pig(y) \leftrightarrow x = y)$ or, for short, $\exists!x fat-pig(x)$. Now, the simple model has it that $g(c + \text{“The fat pig won a prize”}) = g(c) \cup Cnv(\text{“The fat pig won a prize”}, c)$. Since the presupposition $\exists!x fat-pig(x)$ and $g(c)$ together entail:

$$\begin{aligned} &fat-pig(Arnold) \vee fat-pig(Babe), \\ &skinny-pig(Babe), \\ &I(fat-pig(Arnold)), \end{aligned}$$

the updated common ground would then entail these sentences. In light of the constraint on models that skinny pigs are not fat pigs, it would also

entail *fat-pig(Arnold)*. This is just as it should be. We expect that Arnold is the pig picked out by the definite description. Indeed, the effect of simple growth of the common ground, in this case, would be that Arnold comes to satisfy the description.

In this case it is plausible that the common ground should simply grow. Notice that this is a case where the conveyed meanings of an utterance are consistent with the prior common ground. The sentence *fat-pig(Arnold)* reflects the presupposition introduced by the utterance that Arnold and other pigs his size are for the meantime to be called “*fat pigs*”. The continued presence of *I(fat-pig(Arnold))* in the common ground reflects the participants’ mutual belief that pigs this size *could* nonetheless be denied this description.

In other cases, though, it is not plausible that the common ground should simply grow. We shall consider cases where the meanings conveyed by an utterance are *not* consistent with the prior common ground. Suppose, for example, that earlier accommodation has already brought it about that Arnold is for the meantime to be called “*fat*”, although it is a matter of common knowledge that he is a borderline case. Now the farmer moves Arnold into a pen with two very round pigs, Babe and Cornelius, so that the common ground is:

I(fat-pig(Arnold)),

fat-pig(Arnold),

$\forall x(\text{pig}(x) \rightarrow x = \text{Arnold} \vee x = \text{Babe} \vee x = \text{Cornelius})$,

fatter(Babe,Arnold),

fatter(Cornelius,Arnold).

In the company of these much fatter pigs, the farmer can no longer pick out Arnold by saying “the fat pig”. But, we think, he can pick him out by saying “the pig that isn’t fat”. When in this new context the farmer says “The pig that isn’t fat won a prize”, conveying *winner($\lambda x (\text{pig}(x) \& \neg \text{fat-pig}(x))$)*, the accommodating hearer will choose a stricter sense of “fat pig”, so as to incorporate this conveyed meaning into the common ground. Understood in its new sense, the expression will not apply to Arnold, and so the prior presupposition *fat-pig(Arnold)* ought not to be carried over to the updated common ground. The common ground in this instance should not simply grow.

Sometimes, then, presuppositions must be excluded from the prior common ground when accommodating vague utterances. Even so, we seem

to retain as much common ground as we can. To elaborate our example somewhat, accommodating the speaker by denying Arnold the description ‘fat pig’ will not lead the neighbor to call into question unrelated matters, such as the curliness of Cornelius’ tail. The change involved seems to be *minimal* change. This consideration lead us to the following recipe for accommodation:

Take the conveyed meanings of an utterance. Then add as much of the prior common ground as you can without introducing inconsistency.

Notice that the accommodation which concerns us here does not require changes in the common ground because the world represented therein has changed. Changes are not required because a participant in the discourse was mistaken, either; rather, changes are required because vague language has come to be used differently. This suggests that accommodation ought not to lead us to reject any sentence that describes the world in purely non-vague terms. When we make minimal changes to regain consistency, only vague sentences should go. This suggests the following constraint on minimal change: a minimal change to a set of sentences will be one in which no vague sentences are unnecessarily deleted from it.¹² The recipe for accommodation then comes to this:

Take the conveyed meanings of an utterance. Add the non-vague sentences of the common ground. Finally, add as many of the remaining vague sentences as you can without introducing inconsistency.

To put this recipe precisely, the following definition is useful:

For any two sets A and S of sentences, $A \perp S$ is the set $\{T : S \subseteq T \subseteq A \cup S, T \not\subseteq \perp\}$, and for all U , if $T \subseteq U \subseteq A \cup S$ and $U \not\subseteq \perp$, then $U = T$.¹³

¹² One might include in the common ground not only whatever the situation makes obvious and whatever has already been established or assumed in the conversation, but also things to which the participants are less firmly committed. These might come, say, from earlier conversations, or from more or less unreliable third parties. Accommodation should then sometimes allow such weakly entrenched non-vague sentences to be given up so as to make room for vague sentences. For now, though, we assume full commitment to all of the common ground.

¹³ This notion is familiar from work in belief revision. See (Alchourrón et al., 1985). Its application to the updating of common ground is natural given the similarity of the project: we describe how the addition of presuppositions to a set of previously accepted presuppositions leads to a modification of that set.

Let $g_{\text{vague}}(c)$ be the subset of $g(c)$ that includes just the vague sentences. Then $\text{Cnv}(s, c) \cup (g(c) \setminus g_{\text{vague}}(c))$ contains just the conveyed meanings of s and the non-vague part of the common ground. Consider any set g^* in $g_{\text{vague}}(c) \perp (\text{Cnv}(s, c) \cup (g(c) \setminus g_{\text{vague}}(c)))$. On the current proposal, g^* has the following significance: it is one possible way of making minimal changes to $g(c)$ so as to accommodate the speaker's utterance of s . Since there can be more than one such g^* , we introduce a selection function Sel which, given as its argument a non-empty set of alternatives, returns a member of this set, and given no alternatives, returns the empty set. Intuitively speaking Sel has the task of choosing, from among a number of alternatives, the best way of updating the common ground. Finally, then, the proposal we shall consider for updating the common ground following an utterance is this:

$$g(c + s) = \text{Sel}(g_{\text{vague}}(c) \perp (\text{Cnv}(s, c) \cup (g(c) \setminus g_{\text{vague}}(c))))^{14}$$

What makes one way of accommodating better than another? What constraints can be placed on Sel ? We shall put off these questions until Section 6. Meanwhile, we shall illustrate the proposed model of accommodation by returning to the examples that motivated it. Then we shall consider a further example involving plural definite descriptions. In these examples the model will yield only one option for Sel . This leaves Sel with no choice, and us able to proceed without saying more about this function.

5. EXAMPLES

Consider the farmer's utterance, "The fat pig won a prize". Here, the farmer uses a vague definite description to refer to a pig, even though none satisfies the description. This anomalous situation brings about a change in common ground to one in which something does satisfy it.

EXAMPLE 1. There are just two pigs: a borderline fat pig and a runt. The common ground $g(c)$ contains the following sentences:

- (1) $I(\text{fat-pig}(\text{Arnold}))$
- (2) $D(\text{skinny-pig}(\text{Babe}))$

¹⁴ A reviewer suggests that this definition of revision is unintuitive, being sensitive to "the superficial way in which information is presented". As this reviewer points out, we could move to some form of "theory revision", if there is a problem here. See, for example, Carlos Alchourrón et al. (1985). But we do not think there is a problem. For a discussion of the advantages of a fine-grained notion of "belief bases", see, for example, Sven Ove Hansson (1992).

$$(3) \quad \forall x(\text{pig}(x) \rightarrow x = \text{Arnold} \vee x = \text{Babe})$$

$$(4) \quad \text{fatter}(\text{Arnold}, \text{Babe})$$

Now the farmer says s : ‘The fat pig won a prize’. Intuitively speaking, ‘the fat pig’ picks out Arnold, even though Arnold is a borderline fat pig. We show that the following sentences are entailed by $g(c+s)$: $\text{fat-pig}(\text{Arnold})$, $I(\text{fat-pig}(\text{Arnold}))$, $\neg\text{fat-pig}(\text{Babe})$.

DEMONSTRATION. $g(c) \setminus g_{\text{vague}}(c) = \{1, 2, 3, 4\}$. Since $\text{winner}(\iota x \text{ fat-pig}(x)) \in \text{Cnv}(s, c)$, $\text{Cnv}(s, c) \models \exists!x \text{ fat-pig}(x)$. Since $g_{\text{vague}}(c) = \{\}$, $g_{\text{vague}}(c) \perp (\text{Cnv}(s, c) \cup (g(c) \setminus g_{\text{vague}}(c)))$ has just a single element: $\text{Cnv}(s, c) \cup \{1, 2, 3, 4\}$. This is then the set selected by Sel as the new common ground $g(c+s)$. In virtue of the constraint on models that the positive extensions of skinny-pig and fat-pig are disjoint, $\{\exists!x \text{ fat-pig}(x)\} \cup \{1, 2, 3, 4\} \models (\text{fat-pig}(\text{Arnold}) \ \& \ \neg\text{fat-pig}(\text{Babe}))$. Clearly, then, $g(c+s) \models \text{fat-pig}(\text{Arnold})$, $g(c+s) \models I(\text{fat-pig}(\text{Arnold}))$, and $g(c+s) \models \neg\text{fat-pig}(\text{Babe})$.

Ex. 1 illustrates a change in common ground that is simple growth: $g(c+s)$ is a superset of $g(c)$. The updated common ground reflects an adjustment of the positive extension of ‘ fat-pig ’ to include Arnold. Next, we consider a case where accommodation reflects an adjustment to both the positive and negative extensions of a predicate. There are two borderline fat pigs. The farmer uses the description ‘the fat pig’ to refer to the rounder one, but in doing so, brings about a common ground which contains the presupposition that the less round pig is *not* a fat pig.

EXAMPLE 2. The common ground $g(c)$ contains just the following sentences:

$$(1) \quad I(\text{fat-pig}(\text{Arnold}))$$

$$(2) \quad I(\text{fat-pig}(\text{Babe}))$$

$$(3) \quad \forall x(\text{pig}(x) \rightarrow x = \text{Arnold} \vee x = \text{Babe})$$

$$(4) \quad \text{fatter}(\text{Arnold}, \text{Babe})$$

Now the farmer says s : ‘The fat pig won a prize’. Here, too, ‘the fat pig’ ought to pick out Arnold. We show that the following sentences are entailed by $g(c+s)$: $\text{fat-pig}(\text{Arnold})$, $\neg\text{fat-pig}(\text{Babe})$.

DEMONSTRATION. It can be shown that $g(c + s) \models \exists!x \textit{fat-pig}(x) \cup \{1, 2, 3, 4\}$ and $g(c + s) \models \textit{fat-pig}(\textit{Arnold})$. (Analogous to above.) Since $\exists!x \textit{fat-pig}(x) \cup \{1, 2, 3, 4\} \models \neg \textit{fat-pig}(\textit{Babe})$, it also follows that $g(c + s) \models \neg \textit{fat-pig}(\textit{Babe})$.

So much for simple growth. We now return to the case that illustrates a shift in common ground, in Section 4. Arnold, a borderline fat pig, is in a pen with Babe and Cornelius, two very round pigs. The common ground includes the presupposition that all three pigs are fat pigs, but when the farmer uses the description ‘the pig that is not fat’ to pick out Arnold, it shifts.

EXAMPLE 3. The common ground $g(c)$ contains just the following sentences:

- (1) $I(\textit{fat-pig}(\textit{Arnold}))$
- (2) $\textit{fat-pig}(\textit{Arnold})$
- (3) $\textit{fat-pig}(\textit{Babe})$
- (4) $\textit{fat-pig}(\textit{Cornelius})$
- (5) $\exists x(\textit{tail}(\textit{Cornelius}, x) \& \textit{curly-tail}(x) \& I(\textit{curly-tail}(x)))$
- (6) $\forall x(\textit{pig}(x) \rightarrow x = \textit{Arnold} \vee x = \textit{Babe} \vee x = \textit{Cornelius})$
- (7) $\textit{fatter}(\textit{Babe}, \textit{Arnold}) \& \textit{fatter}(\textit{Cornelius}, \textit{Arnold})$

Now the farmer says s : ‘The pig that is not fat won a prize’. This time, the description that picks out Arnold is ‘the pig that is not fat’. We show that the following sentences are entailed by $g(c + s)$: $I(\textit{fat-pig}(\textit{Arnold}))$, $\neg \textit{fat-pig}(\textit{Arnold})$.

DEMONSTRATION. $g(c) \setminus g_{\text{vague}}(c) = \{1, 6, 7\}$. Since $\textit{winner}(tx(\textit{pig}(x) \& \neg \textit{fat-pig}(x))) \in \textit{Cnv}(s, c)$, $\textit{Cnv}(s, c) \models \exists!x(\textit{pig}(x) \& \neg \textit{fat-pig}(x))$, and $\textit{Cnv}(s, c) \cup (g(c) \setminus g_{\text{vague}}(c)) \models \neg \textit{fat-pig}(\textit{Arnold})$. Therefore $\textit{Cnv}(s, c) \cup (g(c) \setminus g_{\text{vague}}(c)) \cup \{2\}$ is not satisfiable. Assuming, as is reasonable, that $\textit{Cnv}(s, c) \cup (g(c) \setminus g_{\text{vague}}(c)) \cup \{3, 4, 5\}$ is satisfiable, it follows that there is a single maximal subset of $g_{\text{vague}}(c)$ that is consistent with $\textit{Cnv}(s, c) \cup (g(c) \setminus g_{\text{vague}}(c))$, namely $\{3, 4, 5\}$, and that $g_{\text{vague}}(c) \perp (\textit{Cnv}(s, c) \cup (g(c) \setminus g_{\text{vague}}(c)))$ has just a single element: $\textit{Cnv}(s, c) \cup \{1, 3, 4, 5, 6, 7\}$. This is then the set selected by \textit{Sel}

as the new common ground $g(c + s)$. Clearly, by the earlier observation, $g(c + s) \models I(\text{fat-pig}(\text{Arnold}))$ and $g(c + s) \models \neg \text{fat-pig}(\text{Arnold})$.

This example illustrates a shift in common ground. While the prior common ground includes the presupposition that Arnold is a fat pig, the updated common ground does not. It also illustrates the notion of minimal change. The updated common ground includes the non-vague sentences that are a part of the prior common ground, as well as vague sentences that are unaffected by the newly introduced presuppositions, such as 5.

In each of the examples considered above, the common ground is updated in such a way that only one thing comes to satisfy the description of the utterance. In this way, accommodation and the resolution of the referent of the description go hand in hand. Reference resolution need not accompany accommodation, however. In the next example, the farmer uses ‘the fat pigs’ to refer to a subset of the pigs that are present, but the neighbor cannot say which ones. Nonetheless, the prior common ground, which established that Arnold is a fat pig, is updated to include the presupposition that Babe is a fat pig too.

EXAMPLE 4. The common ground contains just the following sentences:

- (1) $\text{fat-pig}(\text{Arnold})$
- (2) $I(\text{fat-pig}(\text{Babe}))$
- (3) $I(\text{fat-pig}(\text{Cornelius}))$
- (4) $\text{skinny-pig}(\text{Doris})$
- (5) $\forall x(\text{pig}(x) \rightarrow x = \text{Arnold} \vee x = \text{Babe} \vee x = \text{Cornelius} \vee x = \text{Doris})$
- (6) $\text{fatter}(\text{Babe}, \text{Cornelius})$

Now the farmer says s : ‘The fat pigs won prizes’. We show that the following sentences are entailed by $g(c + s)$: $\text{fat-pig}(\text{Arnold})$, $\text{fat-pig}(\text{Babe})$, $\forall x(\text{fat-pig}(x) \rightarrow \text{winner}(x))$.

DEMONSTRATION. It can be seen that $g(c) \setminus g_{\text{vague}}(c) = \{2,3,5,6\}$. $\text{Cnv}(s, c) \models \text{winners}(\iota_{\max} X \text{ fat-pigs}(X))$. Assuming, reasonably, that $\text{Cnv}(s, c) \cup (g(c) \setminus g_{\text{vague}}(c)) \cup \{1, 4\}$ is satisfiable, it follows that $g_{\text{vague}}(c) \perp (\text{Cnv}(s, c) \cup (g(c) \setminus g_{\text{vague}}(c)))$ has just the single element: $\text{Cnv}(s, c) \cup \{1, 2, 3, 4, 5, 6\}$. This is then the set selected by Sel as the new common

ground $g(c + s)$. Since $\text{fat-pig}(\text{Arnold}) \in g(c + s)$, we have $g(c + s) \models \text{fat-pig}(\text{Arnold})$. In virtue of constraints on models, $\{\text{winners}(\iota_{\max} X \text{ fat-pigs}(X))\} \cup \{1, 2, 3, 4, 5, 6\} \models \text{fat-pig}(\text{Babe})$. Clearly, then, $g(c + s) \models \text{fat-pig}(\text{Babe})$. Finally, since $\text{winners}(\iota_{\max} X \text{ fat-pigs}(X)) \models \forall x (\text{fat-pig}(x) \rightarrow \text{winner}(x))$, it follows that $g(c + s) \models \forall x (\text{fat-pig}(x) \rightarrow \text{winner}(x))$.

It appears that the updated common ground is not committed either way on the question whether Cornelius is a fat pig. That is, $g(c + s) \not\models \text{fat-pig}(\text{Cornelius})$ and $g(c + s) \not\models \neg \text{fat-pig}(\text{Cornelius})$. Therefore the updated common ground does not settle the question of which pigs are picked out by the plural description ‘the fat pigs’.

6. ELABORATIONS

In many cases, one possible accommodation is better than the others. In Sect. 4, we introduced a selection function that tells us which one it is, but we have yet to say anything substantial about the choice. We have been able to get away with this because, in each of our examples so far, there was no real choice. There was just a single way to accommodate. In this section, we turn to cases in which there are several ways. The discussion brings out some factors in the choice between them.

EXAMPLE 5. The farmer and the neighbor are standing some distance from a feeder. Arnold is directly between them and the feeder. Babe is also between them and the feeder, but is much closer to it and a little off to one side. The farmer says, “The pig in front of the feeder won a prize”. We assume the arrangement of the pigs to be such that the neighbor understands that he means Babe.¹⁵

We assume that the positive extension of *in-front-of-the-feeder* includes everything within a cone extending from the feeder for some distance in the direction of the speakers, but that prior to the utterance it includes no pigs. The farmer’s utterance brings it about that the positive extension of *in-front-of-the-feeder* stretches to include a single pig. But it could stretch in two different ways. It could come to include Arnold, the pig at a greater distance from the feeder, or Babe, the one at a wider angle. Correspondingly, there will be two candidates for the selection function to choose between.

¹⁵ This seems reasonable if Arnold is sufficiently far away from the feeder, and Babe is sufficiently close.

We have assumed the position of the pigs to be such that, following the farmer's utterance, the common ground establishes that Babe is the unique pig in front of the feeder. The positive extension of *in-front-of-the-feeder* grows to include the pig at a wider angle rather than the one at a further distance. Somehow this change involves a lesser stretch. The selection function must be made sensitive to the fact that some "minimal" changes to the common ground reflect smaller changes in the usage of vague predicates than others. Here is another such case:

EXAMPLE 6. The farmer and his neighbor are standing some distance from a feeder. Arnold, a borderline fat pig, is directly in front of the feeder, and Babe, a definitely fat pig is slightly off to one side. The farmer says, "The fat pig in front of the feeder won a prize".

Here, again, there are two candidates to choose between, one reflecting an adjustment in the usage of 'fat pig', and the other an adjustment in the usage of 'in front of the feeder'. It seems that if Arnold is very close to being definitely fat, and if Babe is so much to the side that he is very near to being definitely *not* in front of the feeder, then the choice is clear – the updated common ground will establish that Arnold, not Babe, is the fat pig in front of the feeder. Again, it seems that among the different options for updating the common ground, the one that reflects the smallest degree of change in the usage of vague predicates is preferred.

It might happen that there are several possibilities for accommodation but no grounds for selecting among them. In the above case, suppose that Arnold were no closer to being definitely fat than Babe is to being definitely in front of the feeder. What accommodation takes place in these cases? Perhaps none at all. After all, the common ground includes only *shared* presuppositions – it does not change in ways that cannot be followed by both participants. Alternatively, a hearer might put off the choice between candidates for accommodation until more information becomes available.

We have described some of the possibilities for accommodating vague descriptions. We finish with a case in which we think accommodation is difficult or impossible, in that it stretches terms past their limits. The case we have in mind again involves three pigs, but these are no ordinary pigs. The first of them is immense, fatter than anything you have ever seen before. The second, though this seems scarcely possible, is fatter still. The third pig is truly, but truly obese – a real monster. And now the farmer says, "The fat pig is going to the fair".

One natural first response will be some sort of puzzlement. You cannot go calling just one pig fat, withholding this honorific from all except the one monster, since each pig is immense. Perhaps there has been some mistake, and the speaker does not mean for reference to be within this group of pigs after all? Perhaps something has gone wrong with the coordination of the common ground among the participants to the conversation? Perhaps one of the pigs is somehow more salient than the rest, and the hearer has not noticed? Perhaps the hearer is mistaken about which pigs are whose? Failing these and a few other possibilities, it seems to us, the farmer has just got to be joking. Whatever the case may be, the hearer ought not, without a blink, to accommodate the speaker by changing to a very exclusive sense of 'fat pig'. To do so would be to miss the difficulty with the coordination, or the joke, or whatever.

Our proposal does not allow accommodation to proceed normally in this example. To see what makes this example different from the previous ones, notice that the prior common ground will include sentences expressing that each of the pigs is *definitely* fat, which is to say, fat on all precisifications. Sentences within the scope of the operator '*D*', in keeping with the earlier exclusion of higher-order vagueness, are not vague. And since sentences expressing that all three of these pigs are definitely fat are jointly inconsistent with the presupposition that there is just one fat pig, it follows that accommodation is undefined in this case. There is no way at all to make minimal changes to the prior common ground so as to accommodate the speaker's utterance, and so our selection function finds itself with no alternatives to choose among. This is a formal reflection of the puzzlement that, as we see it, ought to prompt a hearer to wonder about coordination, jokes and the like.

Perhaps, after a few blinks, the hearer ought to understand that the farmer means the one monster pig anyway. If so, we think, the farmer's usage of 'the fat pig' might be understood in the way suggested by Donnellan's example of the usurper to the throne, discussed in the introduction. The usage is in this case not straightforwardly sincere. Alternatively, we might at this point bring in the comparison classes that we set aside earlier on, allowing accommodation to proceed by changing the comparison class that calibrates the adjective 'fat'. Whereas previously a pig had merely to be fat in relation to other pigs to qualify, after the recalibration it will have to be fat in relation to obese pigs. Similarly, the home handyman who usually gets by with his one adjustable wrench might sometimes need a larger one.

Assuming comparison classes can be brought in at this point, one might wonder why we do not recommend such a treatment for all of the ex-

amples. Since we have not seen a proposal along these lines developed in detail, a weighing of merits is likely to be unfair, both to our proposal and to its hypothetical competitor. But we can speculate. In the introduction, we gestured toward a theoretical reason to prefer a treatment such as ours: we expect that if we are to rely on comparison classes for everything, we will need to multiply them beyond the needs of the compositional semantics of noun phrases. Additionally, the case of the three definitely fat pigs may present an empirical difficulty for a treatment that involves comparison classes, without any alternative mechanism to balk at the farmer's talk about "the fat pig". Accommodation *ought* to be difficult in cases like this one. If changing comparison classes allows it to go without a hitch, then so much the worse for changing comparison classes.

REFERENCES

- Alchourrón, C., P. Gärdenfors, and P. Makinson: 1985, 'On the Theory of Logic Change: Partial Meet Functions for Contraction and Revision', *Journal of Symbolic Logic* **50**, 510–30.
- Beaver, D.: 1997, 'Presupposition', in J. Van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language*, Elsevier, Amsterdam.
- Donnellan, K.: 1966, Reference and Definite Descriptions', *Philosophical Review* **75**.
- Fine, K.: 1975, 'Vagueness, Truth, and Logic', *Synthese* **30**, 265–300.
- Gärdenfors, P.: 1988, *Knowledge in Flux*, MIT Press, Cambridge, Mass.
- Gazdar, G.: 1979a, *Pragmatics: Implicature, Presupposition and Logical Form*, Academic Press, New York.
- Gazdar, G.: 1979b, 'A Solution to the Projection Problem', in C.-K. Oh and D. A. Dinneen (eds.), *Syntax and Semantics II: Presupposition*, Academic Press, New York, 57–89.
- Grice, H. P.: 1967, 'Logic and Conversation', Unpublished MS of the William James Lectures, Harvard University, reprinted, in part, in P. Cole and J. L. Morgan (eds.) 1975, *Syntax and Semantics 3: Speech Acts*, Academic Press, New York.
- Hansson, S. O.: 1992, 'In Defense of Base Contraction', *Synthese* **91**, 239–45.
- Kamp, J. A. W.: 1975, 'Two Theories about Adjectives', in E. Keenan (ed.), *Formal Semantics of Natural Languages*, Cambridge University Press, Cambridge.
- Karttunen, L. and S. Peters: 1975, 'Conventional Implicature in Montague Grammar', *Syntax and Semantics* **11**, 1–56.
- Karttunen, L. and S. Peters: 1979, 'Conventional Implicature', *Syntax and Semantics* **11**, 1–56.
- Karttunen, L. and S. Peters: (19??) 'Requiem for Presupposition', in *Proceedings of the Third Annual Meeting of the Berkeley Linguistic Society*, 360–371.
- Levinson, S. C.: 1983, *Pragmatics*, Cambridge University Press, Cambridge.
- Lewis, D.: 1979, 'Scorekeeping in a Language Game', *The Journal of Philosophical Logic* **8**, 339–359. Reprinted in D. Lewis (1983) *Philosophical Papers*, Volume I. Oxford University Press, New York.
- Manor, R.: 1995, 'Pragmatic Considerations in Semantic Analyses', *Pragmatics and Cognition* **3**(2), 225–245.

- Marr, D.: 1982, *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*, W. H. Freeman, San Francisco.
- Morreau, M.: 1999, 'Supervaluation can Leave Truth-Value Gaps After All', *The Journal of Philosophy* **96**(3), 148–156.
- Pinkal, M.: 1983, 'On the Limits of Lexical Meaning', in R. Bauerle, C. Schwarze and A. von Stechow (eds.), *Meaning, Use and Interpretation* de Gruyter, Berlin.
- Pinkal, M.: 1995, *Logic and Lexicon*, Kluwer Academic Publishers, Dordrecht.
- Sorenson, R.: 1988, *Blindspots*, Clarendon Press, Oxford.
- Stalnaker, R.: 1974, 'Pragmatic Presuppositions', in M. K. Munitz and P. K. Unger (eds.), *Semantics and Philosophy*, NYU Press, New York, 197–213.
- Whitehead, A. N. and B. Russell: 1922, *Principia Mathematica*, 2nd edn., Cambridge University Press, Cambridge.
- Williamson, T.: 1994, *Vagueness*, Routledge, London.

Alice Kyburg
Department of Philosophy
University of Wisconsin Oshkosh
Oshkosh, WI 54901
kyburg@uwosh.edu

Michael Morreau
Department of Philosophy and UMIACS
University of Maryland College Park
College Park, MD 20742
mimo@umiacs.umd.edu