A Unified Theory of Truth and Reference

Barry Smith and Berit Brogaard[1]

Preprint version of paper published in Logique et Analyse, 43 (2000, published 2003)

The truthmaker theory rests on the thesis that the link between a true judgment and that in the world to which it corresponds is not a one-to-one but rather a one-to-many relation. An analogous thesis in relation to the link between a singular term and that in the world to which it refers is already widely accepted. This is the thesis to the effect that singular reference is marked by vagueness of a sort that is best understood in supervaluationist terms. In what follows we show that the supervaluationist approach to singular reference, when wedded to the truthmaker idea, yields a framework of surprising power, which offers a uniform set of solutions to a range of problems regarding identity, reference and knowledge, problems which have hitherto been dealt with on an ad hoc basis.

1. The Problem of the Many

You make a true judgment to the effect that Amundsen flew to the North Pole. On the theory to be advanced in what follows, it is in first approximation a certain oddly demarcated portion of reality—which might be labeled *Amundsen's flight*—which makes your judgment true. Your judgment serves to set into relief this portion of reality against a background of other portions of reality which are traced over. At the same time the designated portion of reality somehow *necessitates* the truth of your judgment, which is to say: in any possible world in which your judgment and that portion of reality exist, your judgment is true.

1. Thanks are due to Peter Baumann, Randall Dipert, Philipp Keller, Kevin Mulligan, Jonathan Sanford and Achille Varzi for helpful comments, and also to the National Science Foundation which supported our work on this paper under Research Grant BCS-9975557: "Geographic Categories: An Ontological Investigation" and to the American Philosophical Society for the award to Smith of a Sabbatical Fellowship.

The truthmaker relation thus points simultaneously in two directions: from judgment to reality, and from reality to judgment. We shall be concerned initially with the first arm of this relation, which might be called 'representation' or 'projection.' When you use the name 'Bruno' to refer to the cat in your office, this serves to make Bruno the object of your attention, to set him into relief against a certain background. Our strategy will be to seek to understand how true judgments project upon corresponding portions of reality by taking as our clue this setting into relief of one single object.

But there is a catch: the referring relation between a singular term and its object is itself more complex than might at first appear. For think of Mont Blanc, with its rabbits and foothills. Clearly, there is no single answer to the question as to what it is to which the term 'Mont Blanc' refers. Rather there are, at any given time, many answers, since there are many parcels of reality that are equally deserving of the name 'Mont Blanc.'

This does not mean that the world is vague. It is not that there are bits of physical reality that neither belong nor do not belong to other bits of physical reality. Rather, many of the terms we use to refer to objects in reality are such that, when we use these terms, we stand to the corresponding parcels of reality in a relation that is one-to-many rather than one-to-one. The same applies also when we perceive objects in reality. Here, too, our perceptual acts stand to the corresponding parcels of reality in a projection relation that is one-to-many rather than one-to-one.

Unger (1980), Lewis (1993), and others have already observed that there are many equally good boundaries for many of the entities we name and perceive in the everyday world:

> There are always outlying particles, questionable parts of things, not definitely included and not definitely not included. So there are always many aggregates, differing by a little bit here and a little bit there, with equal claim to be the thing. We have many things or we have none, but anyway not the thing we thought we had. (Lewis 1993, pp. 164f.)

We should not exaggerate this problem. A stone the size of your fist does, it is true, lose an atom from its surface about once every second—but this atom is almost always immediately recaptured by the stone. Even hard objects like stones will however manifest some version of the vagueness problem at the quantum level, and the problem arises already at quite coarse-grained levels in regard to everything that undergoes evaporation or erosion or abrasion: 'think of yourself, or any organism, with parts that gradually come loose in metabolism, or excretion or perspiration or shedding dead skin. In each case, a thing has questionable parts, and therefore is subject to the problem of the many.' (Lewis 1993, p. 165)

Unger concludes that entities such as clouds do not exist. Lewis favors the view according to which clouds exist but our uses of the word 'cloud' in descriptions such as 'that cloud over the hill' do not ordinarily pick out any single one of the many members of the relevant extension. In what follows, we defend a version of Lewis's view.

Our cat Bruno is shedding hairs. Individual strands of hair become gradually loosened and are supported for a time by the surrounding hairs. If, now, we find hairs in the kitchen and Bruno is in the living room, then we can be confident that those hairs are not parts of Bruno. But from that we cannot infer that there is, in the kitchen, some crisp aggregate that is the unique referent of the term 'Bruno.'

You point to a glass on a table, and you ask your friend if he could hand you *that glass*. Here again there is no single aggregate that could serve as designated referent. For think of all the fingerprints and soap particles on the glass or the air bubbles and neutrinos in its interior. The Unger-Lewis problem-of-the-many arises for all of those things that have questionable parts, parts that can be gained and lost without the things themselves ceasing to exist.

But as Lewis points out (1993, p. 178), the many are, in the cases treated of in the above, not *entirely* distinct, since they are not disjoint mereologically. Indeed they overlap to a considerable degree. Although no two of them are identical, any two of them are *almost identical* in that they share almost all their parts in common. They are many, but almost one.

2. Supervaluationism and Beyond

The standard framework for understanding the vagueness involved in singular reference is the method of supervaluation (van Fraassen 1966, Fine 1975). This draws on the recognition that a sentence can often be assigned a determinate truth-value independently of how the referents of its constituent singular terms are more precisely specified—which means: independently of how such referents are restricted to some one of the many available precisified aggregates of matter. A sentence is called *supertrue* if and only if it is true (and superfalse if and only if it is false) under all such precisifications. If, on the other hand, it is true under some ways of precisifying and false under others, then it is said to fall down a supervaluational truth-value gap.

It is then supertrue to say that there is just one cat in your office, or that Bruno is a nice cat. Sentences of this sort are true no matter which of the many aggregates of matter you assign as the precisified referent of the pertinent referring term.

One might, with Lowe (1995), object to this conclusion. For consider the sentence:

[A] Mont Blanc is many mountains.

If [A] were used to express a judgment in some everyday context, then this judgment would be evaluated as false under any one of the many candidate precisifications. [A] is accordingly (as we should expect) superfalse from the supervaluationist perspective. But must not the supervaluationist assume precisely that something like [A] is true? Not necessarily. For she can insist that our ordinary talk of mountains (and of 'Mont Blanc' as having just one mountain as its referent) is, *in ordinary contexts*, in perfect order as it stands. It is merely that such talk must be carefully distinguished from talk of 'precisifications' and the like in contexts of semantics and ontology. The supervaluationist must thus assume, not [A], but rather:

[A′] Mont Blanc can be any one of *these* mountain-shaped chunks of reality.

Lowe's objection is intriguing, however, for it draws attention to the fact that the very same sentence can behave differently in different contexts. This fact is of course well known to those who work on indexicals and like phenomena. What we are about to discover is that it applies in relation to the phenomenon of vageness, too. In standard semantics, of course, sentences are evaluated semantically in abstraction from any context of use. It is possible, however, to apply semantic evaluations not to sentences but to the judgments which such sentences express. This adds an extra element of realism to the supervaluationist account, since it is through judgments that terms are projected onto reality by intending subjects. It then transpires that the very same sentence may be used in different contexts to express distinct judgments even where the singular terms involved refer to what is intuitively the same parcel of reality. The supervaluations of the given judgments will then look very different, even though the sentences in question are, as syntactic objects, one and the same, and the degree and the type of vagueness by which referring terms are affected will be dependent on this context.

Focusing upon judgments in their contexts will in addition have the effect of aligning singular terms such as 'John' or 'Cleopatra's Needle' with indexical expressions such as 'him' or 'that' (so that we can speak indiscriminately in what follows of 'singular terms' or 'referring expressions,' paying no attention to the distinction between indexicals and non-indexicals). At the same time the shift in focus will mean that our linguistically mediated reference toward objects will become aligned theoretically with the non-linguistically mediated reference we find in perception.[2]

2. It may indeed be that reference-fixing is perception-dependent for all singular terms, including not only demonstratives and indexicals but also proper names; on this see Mulligan 1997.

3. Context Switch

Our goal, then, is to add contexts to the supervaluationist approach. A context, for our purposes, is a portion of reality associated with a given conversation or perceptual report and embracing also the beliefs and interests and background knowledge of the participants, their mental set, patterns of language use, ambient standards of precision, and so forth. Above all it is a matter of what is paid attention to by participant speakers and hearers on given occasions. This initially somewhat sprawling notion of context will however be tamed when we set forth our ideas in more formal terms at the end of the paper. There it will become apparent that the work done by contexts in our theory rests on one single feature, namely on the fact that contexts may be *more or less refined*, or in other words that they may determine a greater or lesser *granularity* of ways in which we relate to objects in the world.

To see how this works, consider:

[B] This glass is empty.

This sentence will come out supertrue in the context occupied by your thirsty brother after draining his glass, because it will come out true independently of how the referent of 'this glass' as he intends it in that context is further precisified. The same sentence, however, will come out superfalse when it is uttered by a public hygiene inspector examining the contents of the glass under a powerful microscope. This is because, in the former context, the molecules of water, soap and grease attaching themselves to the interior walls of the glass are not distinguished by your brother from the glass which he holds in his hand. In the latter, more refined context, however, these molecules are so salient that they must be taken account of in all precisifications consistent therewith.

The move from an everyday inspection of the glass to the more careful inspection involving powerful microscopes amounts to what we can think of as a context switch—a phenomenon analogous to the sort of Gestalt switch which occurs when you move from *duck* to *rabbit*. Consider what happens when you suddenly realize that you have wandered unknowingly onto a film set, or that you are the victim of an elaborate hoax. Such switching of context can be brought about rather easily. Even simply mentioning a certain possibility may suffice. Imagine, for example, that you are looking at Mont Blanc from a distance. There are almost no clouds in the sky. You see the mountain clearly on the horizon. It seemingly has a sharp boundary that separates it from the surrounding sky. You cannot see the people and trees on the mountain or the small rabbits crawling around under its bushes. You know perfectly well that there are such things, but in the given circumstances, in which they are in no way salient to you, you quite properly ignore them. Your perception does not separate

out the things you are seeing from the things you are ignoring. It bears upon reality as through a diffuse beam, comprehending aggregates some of which include rabbits and some of which do not.

But suppose someone now asks you whether you think that rabbits are part of Mont Blanc. This very question establishes a new context. The lazy diffuseness of your earlier perceptual projection is suddenly brought to an end. For in responding to this question your use of the term 'Mont Blanc' picks out only aggregates that do not include rabbits as parts. The beam of your referential searchlight has, in the new context, become narrowed, its focus adjusted, and this in a highly specific way. A new battery of what Fine (1975) calls 'penumbral conditions' on available precisifications has been brought into effect.

Since we pay attention in different ways and to different things in different contexts, terms refer differently in those different contexts. In the first, perceptual context—in which you do not pay attention to rabbits—your attentions in the direction of Mont Blanc relate to a family of aggregates some of which will include rabbits as parts. The rabbits are in that context not parceled out as distinct things. In the second context, however, your attentions are directed to a family of aggregates none of which includes those rabbits as parts. This is because the rabbits have been set into relief as objects distinct from the mountain: they have been carved out in your attention as objects in their own right.

Surprisingly, now, it is through all of these transformations never the case that judgments of the form:

[C] Rabbits are parts of Mont Blanc

are true. Indeed no such judgments are effected at all. For again: as soon as a judger is in a position where he might be able to formulate such a judgment, he has moved to a context within which he is, willy nilly, executing a *finer partition of reality*, in which rabbits are cognized as entities disjoint from mountains. A sentence like [C] might express a true judgment in a context in which you do not pay attention to the rabbits on the mountain; then however you would not be capable of making a judgment of that sort. Consider, for other examples of the same phenomenon, sentences such as: 'the microbes in John's ear are parts of John,' 'the dirt on your hubcaps is part of your car,' 'the umbrella in your cocktail is part of your meal.' These sentences are, we shall say, unjudgeable. (A Wittgensteinian might say that the corresponding language games *are not played*.)

Whether or not a sentence is judgeable in a given context depends on whether a speaker would feel comfortable expressing the sentence in that context (and we hazard that sentences expressing logical contradictions are unjudgeable in every normal context). Whether or not a speaker feels

comfortable depends on his own psychology and on his stock of beliefs as well as on the linguistic and epistemological standards imposed by the community to which he belongs. In our community there is a real estate market for houses which include garages as parts. Thus nobody would feel uncomfortable judging that a given garage is part of a given house. In many hunting communities, similarly, there is a real estate market for parcels of land in which the owner of a mountain is thereby also the owner of all the animals on that mountain. The sentence 'these rabbits are part of my mountain' is then perfectly judgeable in such communities.

Whether or not a sentence is judgeable in a given context depends also on what the subject perceives in that context. A subject may see differently in different contexts even in spite of the fact that his underlying perceptual experienes are phenomenologically (which means: internally) indistinguishable. This is not least because, as Mulligan points out (1999), the way your perceptions relate to external reality depends upon your trajectories of possible action. Seeing Mont Blanc *pur* is relevant to one sort of action, seeing Mont Blanc with rabbits is relevant to another. Seeing the glass without the grease is relevant to one sort of action, seeing it with the grease is relevant to another.

4. Elusive Reference

The above considerations dictate what we might call 'reference contextualism'—a view analogous to the epistemological contextualism defended by Lewis in his "Elusive Knowledge" (1996). Knowledge that *p* is elusive, for Lewis, if the very fact that one begins to discuss what possibilities there are that *not-p* brings it about that one no longer knows that *p*. He is referring specifically to the knowledge involved in presupposing, as when, in telling the time when glancing up at the church clock, you presuppose that the clock is in good working order. Such knowledge is knowledge, Lewis holds,

> but it is an *especially* elusive sort of knowledge, and consequently it is an unclaimable sort of knowledge. You do not even have to practice epistemology to make it vanish. Simply *mentioning* any particular case of this knowledge, aloud or even in silent thought, is a way to attend to the hitherto ignored possibility, and thereby render it no longer ignored, and thereby create a context in which it is no longer true to ascribe the knowledge in question to yourself or others. (1999, p. 438)

Knowledge that is not completely certain *is* knowledge, for Lewis, but it is not claimable knowledge.

And so similarly for us here: our everyday perceptions will often pick out parcels of reality which include alien objects as parts because the latter are, in everyday contexts, not projected, not set into relief, as distinct. Yet this sort of parthood is elusive: it is never claimable, since to claim it would amount to bringing about a shift in context.

Under Lewis's contextualism, the truth-value of a knowledge ascription varies with the context in which it is made. Lewis's proposal is based on the idea that, for someone to have knowledge that *p*, the evidence she has for *p* must eliminate the possibilities that not-*p*. We do not want to require that all such possibilities must in all contexts be eliminated, for that would commit us to skepticism.

We can apply an analogous distinction also to the case of reference. When a subject uses a referring term to refer in a given context, then the context helps to determine what is projected by the term. That is, if the term is vague, then the context helps to determine the family of nearly identical, overlapping parcels of reality which will fall within its scope when evaluated supervaluationistically. What is ignored or paid attention to in a given context is determined in some cases by the spatial or temporal distance between the user of a term and its referent, in other cases by the specific repertoire of concepts or interests which the user brings to bear. The greater the distance or the more general or vacuous the concepts, the more diffuse will be the corresponding projection onto reality and therefore the more items may come to be included in the corresponding parcels of reality.[3]

The considerations advanced above now dictate a quite specific version of the view according to which it is not the case that, because of some epistemological defect, we mortal creatures do not know what does and what does not belong to the mountain. There is no single answer to this question, we can now say, because the many possible answers—which is to say: the many possible specifications of relevant families of overlapping candidate parcels of reality—will vary according to the context in which the question is raised.

5. Granularity: the Source of Vagueness

Our earlier remarks about Mont Blanc and its rabbits can now be generalized. When you refer to John, who is cooking in the kitchen, you do not think of all the parts of John or of his surroundings.

---

3. Note that the determination of the kind and degree of vagueness is in every case a matter of what pertains to the context of reference currently occupied by the user. Thus we distance ourselves from those who would impute the vagueness adhering to our uses of singular terms to a lack of knowledge of the putative initial baptismal act by means of which the term is supposed to have been glued to some precisely determinate object.

The cells in John's arm and the fly next to his ear belong to the portion of the world that does not fall under the beam of your referential searchlight. They are traced over.

Every use of a referring expression accordingly brings about, in its context, a partition of reality into two domains: the foreground domain, within which the object of reference is located, and the background domain, which comprehends all entities left in the dark by the operating referential searchlight. But this partition cannot, unfortunately, be understood along simple geographical lines. That is to say, it is not as if one connected, compact (hole-free) portion of reality is set into relief in relation to its surroundings, as Beverly Hills is set into relief within the wider surrounding territory of Los Angeles. For if an object is included in the foreground domain, this does not imply that all the parts of the object are also included therein. This is because each referential partition comes, roughly speaking, with its own granularity: it does not recognize parts beneath a certain size. It is for this reason that each such partition is compatible with a range of possible views as to the ultimate constituents of the objects included in its foreground domain, and the same applies also to the partitions associated with our acts of perception.

It is the coarse-grainedness of our partitions which allows us to ignore questions as to the lower-level constituents of the objects foregrounded by our uses of singular terms. This in its turn is what allows such objects to be specified, not precisely, but rather in such a way that a range of alternative but nearly identical objects are simultaneously included within their corresponding supervaluational scopes. The unwitting author of the coarse-grained partition does not recognize this 'many' because she is focused, precisely, on those parts and moments of the matters in hand which lie above the pertinent granularity threshold.

Consider once again your thirsty brother's judgment to the effect that this glass is empty. This judgment is, as your brother intends it, true. Suppose that there are, as a matter of fact, tiny drops of water at the bottom of the glass. The partition of reality effected by your brother does not recognize these drops of water because they fall beneath the pertinent granularity threshold.

[B] This glass is empty

would in normal circumstances entail: 'There are no drops of water in this glass,' and the latter is false in the case in hand. The mentioned entailment must thus be unavailable to your brother in the given context. That it is unavailable can now be understood as follows: when the question of the presence or absence of molecules of water is raised, this will have the effect of instituting, for him, a new and more refined context, and within this new context any judgment to the effect that 'this glass is empty' would be false.

6. The Problem of Substitutability

Sentences can switch back and forth between judgeability and unjudgeability as a result of such context switch, even though as far as the objects themselves are concerned everything remains the same. Consider 'The Morning Star is not a star' or 'Karol Wojtyła is more intelligent than John Paul II.' These examples point us in the direction of a new sort of solution to the classical puzzles about substitutability. Consider:

[D] Lois believes that Superman can fly.

[E] Lois does not believe that Clark Kent can fly.

Judgments made using sentences like these may be true simultaneously, it is normally held, even though 'Superman' and 'Clark Kent' are co-referential terms. This is Frege's puzzle about beliefs, and as Kripke showed, the same problem can arise even without substitution. Consider:

[F] Peter believes of Paderewski that he had musical talent.

[G] Peter does not believe of Paderewski that he had musical talent.

[F] is true because of what Peter believes regarding a certain pianist. [G] is true because of what Peter does not believe regarding a certain statesman. The musician and the statesman are one and the same, but Peter is not aware of this. Kripke's puzzle is to explain how both [F] and [G] can be true in a way that does not impugn the rationality of the believer.

A range of divergent ways of solving such puzzles has been put forward in the literature. At the one extreme stands Frege himself, who claimed that terms occurring in the scope of intentional verbs do not have their customary references at all, but rather refer to their customary senses, so that we are forced to postulate entirely unmotivated shifts of reference in what may be the course of a single discourse.

At the other extreme stands Russell, who talks of 'propositions' as entities which relate to the objects and properties they are about as a whole relates to its parts. This yields what might be called an identity theory of truth, according to which truthmaker and truthbearer are one and the same.

Frege, we might say, postulated too large a gap between the bearer of truth and (ordinary) referents; the Russellian position, on the other hand, denies the existence of a such a gap entirely. Clearly, on the Russellian view there can be no restrictions on substitution *salva veritate.* If 'Superman' and 'Clark Kent' refer to the same individual, then the propositions expressed by sentences containing these terms are identical (they have the very same parts in the very same arrangement). The so-called 'naive' view, a modern version of Russellianism propounded by Salmon (1986) and Soames (1988), maintains accordingly (1) that referring terms preserve their customary

reference and that (2) co-referring terms are intersubstitutable even when they appear within the scope of intentional verbs. Lois then believes that Clark Kent can fly *because* she believes that Superman can fly. The difference between the two belief reports is not semantic but rather pragmatic—it is a matter of the different presuppositions brought to bear by different speakers. The problem with substituting 'Clark Kent' for 'Superman' in [D] is seen on this view as lying in the fact that the result conveys a pragmatic implication to the effect that Lois would accept a sentence—'Clark Kent can fly'—which she would not accept. But it is true, nonetheless, that Lois believes that Clark Kent can fly just in case she believes that Superman can fly.

Schiffer (1977), Richard (1990), and Crimmins (1992) share with Salmon and Soames the thesis that sentences like [D] and [E] convey information about Lois's different ways of thinking about the same designated individual; but they think that such ways of thinking ('senses,' in Fregean terminology) enter into the truth-conditions of the corresponding belief reports. For them, therefore, it is false to say that 'Lois believes that Clark Kent can fly' because the latter makes a false claim to the effect that Lois accepts the sentence 'Clark Kent can fly.' Belief-ascriptions, on the view in hand, make tacit reference to the believing subject's ways of thinking about the objects to which the beliefs are directed. This view is also known as the 'hidden indexical view,' because belief-ascription sentences are seen as being, like sentences containing indexicals, not absolutely true or false, but rather true or false with respect to the contexts in which they are used to express judgments.

7. Substitutability Contextualized

On our more thoroughgoing contextualist approach, now, we can bring order into this conflicting family of positions.[4] We first of all note that whether or not substitution of co-referring terms can occur *salva veritate* depends on the context in which the relevant judgment (about beliefs, or other matters) is made. It then transpires that in some contexts the hidden indexical view does yield exactly the right result. This is so wherever speakers and hearers care about which terms the believer would use to represent the object the belief is about.

Consider a context where John and Peter are talking about their friend Mary, and specifically about whether or not the judgment expressed by:

[H] Mary believes that Clemens is a famous author

4. It draws in this respect on the moderate contextualism discussed in Richard (1990) and Sider (1995).

is true. They already know that Mary believes that Twain is a famous author. The question is whether she would use the name 'Clemens' to represent 'Twain.'

In other contexts, however, substitution even within the scope of intentional verbs turns out to be fully permissible. Consider for example a context in which John and Peter are concerned merely to establish whether Mary believes that that man, Clemens, depicted in that picture, is a famous author, and they do not care how Mary herself would or could express this belief. Here there are no restrictions on which terms may be used to represent Clemens. In such a context [H] would express a judgment which is true. Because John and Peter are in a context where they do not focus on how Mary refers to Clemens, any referring expression will do to represent Mary's thoughts about Clemens, including mere demonstratives.

In a context in which there are no restrictions on substitution, it may be that we can truly express a person's belief by means of a contradictory sentence without thereby indicating that the person's belief is itself contradictory. Imagine that Mary knows various things about Clemens but does not believe that he is an author. John might point to a picture of Clemens and say 'Mary believes that that author is not an author.' Mary herself, of course, would in no context use 'That author is not an author' to make a judgment—and thus we see that the issue of substitutability is closely related to that of judgeability.

It would take us too far afield to address in detail how the ways in which singular terms project upon corresponding portions of reality are affected, on the contextualist view, by the fact that such terms appear within the scope of intentional verbs. We note only that our approach can be applied also to judgments which give rise to anti-substitution intuitions even though the corresponding sentences involve no intentional verbs. Imagine, for example, that you are on holiday in Metropolis and that you hear the sentences:

[I] Clark Kent went into the phone booth, but Superman came out.

[J] Lois kissed Superman before she kissed Clark Kent.

These sentences (taken from Saul 1997) are used to express judgments by people who you, a stranger to Metropolitan ways, meet on the street.

[K] *Clark Kent went into the phone booth, but Clark Kent came out

[L] *Lois kissed Superman before she kissed Superman

are on the other hand unjudgeable (a fact that is indicated by the use of the linguist's asterisk). [J] and [L] are in addition—to the degree that they can be used to express judgments at all—false. In this, as we shall see, they are similar to 'The Evening Star is brighter than the Morning Star.'

[M] Superman is Clark Kent

on the other hand is (in some, informed, Metropolitan contexts) both judgeable and true.

8. Hesperus and Phosphorus

Matutinus is a stick-in-the-mud astronomer who spends every morning looking at the heavens.[5] Matutinus does not know that the Morning Star is identical to the Evening Star, but his friend Noctis, who does his astronomy in the evenings, has told him that the Evening Star is very bright. The object Matutinus observes does not seem to him to be so bright. He thus thinks to himself, in his partial ignorance, that the Evening Star is very bright, but that the Morning Star is less so. The sentence which he uses to express this judgment is in *our* context close to a logical contradiction; it is, for us, not judgeable, although we can of course mouth the corresponding words. Matutinus, however, because he is in his context not apprised of the fact that the Morning Star and the Evening Star are one and the same, can perfectly well judge (as he sees it) that the one is brighter than the other.

Matutinus has a particular interest in the object he calls 'the Morning Star.' He is working, in our terms, with a partition of reality that traces over (is agnostic about) everything that happens to this object when it is not visible in the morning. Noctis, on the other hand, effects a different partition which picks out what is (unbeknownst to him) the same object but traces over everything that happens to this object when it is not visible in the evenings. The two partitions are compatible with a range of possible views concerning the identity or non-identity of the objects included in their respective foreground domains.

Much of what Matutinus says about the Morning Star is true. For example: 'The Morning Star is brighter than Mars,' 'The Morning Star appears in the morning,' and so on. Matutinus can even utter truths by using sentences which refer simultaneously to both the Morning Star and the Evening Star, as for example:

[N] The Morning Star does not have magical powers, and neither does the Evening Star.

Sentences of this type are however pragmatically flawed, and they will become unjudgeable when once it is discovered that the Morning Star and the Evening Star are one and the same.

Matters are even worse in regard to false judgments of the given types. For when Matutinus uses the terms 'the Morning Star' and 'the Evening Star' simultaneously in a sentence like 'The Evening

5. Our remarks here echo similar thought experiments considered and rejected by Saul (1997) and Kriegel (1998).

Star is brighter than the Morning Star,' *he is attempting to execute a partition of reality which reality cannot sustain*—a partition which seeks to recognize two distinct and mereologically non-overlapping objects in its foreground domain where no such objects can be found. Judgments effected on the basis of such partition-failure manifest a special type of falsehood which makes them quite peculiarly unstable in the sense that, when once it has been recognized that reality does not sustain partitions of the needed sort, not only the corresponding sentences but also a broad family of associated sentences are rendered unjudgeable *en masse*.

9. Theseus's Ship; the Statue and the Clay

A ship sails the seas. Over the years it has all its parts replaced, plank by plank. A museum curator collects the planks, and reassembles them in their original arrangement. There are now two ships. Both are in some way related to the original ship, but which one is identical thereto?

Our contextualist theory can be applied to this problem also, and it then transpires that the key to a solution lies once again in taking seriously the different ways in which referring terms project, diffusely, onto objects in reality. In some contexts, our terms will refer in such a way that it will be true that the ship is, even after all the repairs, still the same as the original ship. These might be contexts in which we are interested only in the ability of the ship to do its job in sailing from port to port. Our partitions in those contexts trace over the separate planks within the ship. In other contexts, however, for example inside museums of naval archaeology, our terms may refer in such a way that it is precisely these planks which are important, so that the ship may for example enjoy continued existence even when it is in a disassembled state.

Simons (1987) has proposed that these two ways of looking at identity through time involve appeals to different notions of identity: functional identity, in the eyes of the shipowner, and material identity, in the eyes of the curator. Simons comes close to providing a correct account of the problem in hand. But once again our contextualist theory is more thoroughgoing, since it grants to a much wider range of actual and possible contexts in which successor relations are tracked across time the power to determine corresponding families of true judgments of identity. Thus in particular both the shipowner and the museum curator can make true judgments of identity relating to the original ship, though there is of course no context in which these two sets of judgments can come out true together. Gallery directors, sculptors, mineralogists and chemists can likewise make true judgments of identity relating to statues and to the lumps of clay and aggregates of clay molecules by which such statues are at any given time constituted. And then again: there is no context in which their

respective sets of judgments can be made to come out true together and thus also no context in which we can use the transitivity of identity to yield a contradiction.

10. Skepticism and the Barn

The problem with skepticism is this: if maximal justification of beliefs is required in order that something be known, then it is impossible to know anything. The reason is that knowledge is closed under known implication. This means that if you know that *p*, and if you know that *p* entails *q*, then you must also know that *q*. The skeptic now uses this closure principle to prove that you have no ordinary knowledge. He first of all explains to you why most everyday beliefs (such as that you have two hands) entail the denial of the skeptical hypothesis (e.g. the hypothesis that you are a brain in a vat). He then, by contraposing on the closure principle, argues that if you fail to know that you are not the victim of a skeptical hypothesis, then it must follow that you also fail to know any of those everyday propositions.

On our theory, however, as we saw already in dealing with the empty glass and your thirsty brother, there are contexts in which otherwise obvious implications are blocked. When your thirsty brother utters 'this glass is empty,' then the inference to 'this glass contains no drops of water' is, to him, unavailable. This is because carrying out the given inference would entail a switch of context which would immediately put your brother in a position where any judgment he might make to the effect that 'this glass is empty' would be false.

Something similar holds, now, in every context in which you enjoy ordinary knowledge. If you say, in an ordinary context, 'I know that it is 4 o'clock,' then a sentence like 'I do not know whether or not I am a brain in a vat' is unjudgeable. There is no way in which you can remain in the ordinary context and still sincerely utter the latter sentence, since uttering it would initiate a switch of context. Once the switch is made, then it is indeed false to say 'I know that it is 4 o'clock' because the stricter context imposes stricter criteria of justification. But this does not detract from the existence of that ordinary knowledge which you possess in those original, ordinary contexts in which you are, like the rest of us, for most of the time happy to live.

Consider now the barn. You are looking at the only genuine barn in an area spawned with barn facades. You say: 'I know that this is a barn.' Intuitively, you fail to know that this is a barn because it is merely accidental that you are looking at the only real barn in the area and not at one of the many barn facades. On our contextualist view, in contrast, you could very well be in a context in which this case does count as genuine knowledge. Your context may, for example, be a very local

one in which it is just this barn which falls within your focus of concern; or it may even be a very far-reaching context in which the many real barns in the surrounding areas suffice to make the barn you see once more typical. But there are of course also intermediate contexts in which the fact that there are so many fake barns in your immediate surrounding area suffices to cast doubt on your claim to knowledge. These are contexts in which you ought to have had the additional knowledge about the surrounding area, so that it is merely an accident that you did not do so and thus also merely an accident that you acquired the belief that the thing in front of you is a barn. If you had known about the barn facades you would have been much more skeptical with respect to the barn in front of you. You do not now, as it turns out, have this extra knowledge. But in a context of the given sort it is very likely that you will soon acquire it. Perhaps you will continue along the road and you will notice that the next barn is merely a facade; or you will meet an actor who will tell you that a film set has been constructed in the area. Any knowledge you might possess to the effect that the thing in front of you is a barn is, therefore, in a context of the given sort, at best knowledge of an unstable sort.

In the first, very local context, the only barn partition which comes to be associated with your judgment is the one which projects upon the barn before you. In the second, very wide-ranging context, your judgment is associated in addition with a much-wider ranging barn partition whose domain might be the whole of Wisconsin. The latter projects in indeterminate fashion across all the barns in this domain and the cluster of barn facades in your immediate surrounding area is thereby traced over, in much the way in which chemical impurities are customarily traced over when we use a term like 'water.' In the intermediate, problematic context, in contrast, your judgment can too easily become associated with the attempt to impose a barn partition upon your immediate surrounding area. This attempt would, however, fail—with fatal consequences for your original claim to knowledge—because reality cannot sustain a partition of the needed sort.

You are thus in a position very like that which is involved in another Gettier scenario: the stopped clock. You look at the clock and you see that it is noon. But you fail to appreciate that the clock is broken. It is just a coincidence that you happen to look at it exactly at noon, the one time during daylight when it shows the time correctly. Your knowledge that it is noon (if you have such knowledge at all), is here of a maximally unstable sort. It is unstable because your judgment, like all judgments of the given type and aetiology, can all too easily become associated with a partition of the successive (distinct) positions of the clock's hands over time—a partition which reality, here, cannot sustain. Note that the fact that your knowledge is maximally unstable does not *ipso*

facto deprive it of the claim to being knowledge. For consider the scenario in which the clock has been in perfect working order until one second *after* you use it to tell the time.

11. New Light on Supervaluational Truth-Value Gaps

Taking contexts and partitions into account in this fashion reveals to us that many commonly advanced examples of sentences which fall down supervaluational truth-value gaps do not in fact do so, either because there are no normal contexts in which they can serve as vehicles for judgment or because their corresponding normal contexts have features which render them immune from indeterminacy. The sentence:

[O] Rabbits are part of Mont Blanc,

on a supervaluationist account, might be true on some ways of making the unmade decision of putting a specific aggregate of molecules into the extension of 'Mont Blanc' and false on others. On the contextual account of reference, however, no such conclusion follows. For as we have seen, [O] is in our normal contexts not judgeable. And so also in many other types of case. We still, to be sure, need to recognize three different alternatives as far as the corresponding *sentences* are concerned. Now, however, these will have the labels: *judgeable and true*, *judgeable and false*, and *not judgeable*, and one and the same sentence can in principle (in different contexts) instantiate all three.

Yet when it comes to our judgments themselves, no such variation is possible: it is in every case an absolute affair whether or not a judgment is true. For it is as if each judgment brings its context along with it, so that to evaluate a judgment as to supertruth and -falsehood is in every case to evaluate that judgment in its context. This implies the following principle:

A judgment of the form '$P(a_1, \ldots, a_n)$' is supertrue if and only if:

(T1) the judgment successfully imposes in its context C a partition of reality assigning to its constituent singular terms 'a_1,' …, 'a_n' corresponding families of aggregates $F_1, \ldots, F_n$, and

(T2) the corresponding families of aggregates are such that, however we select individual f_i from the many F_i, '$P(f_1, \ldots, f_n)$' is true.

A judgment is, in other words, supertrue if and only if it is true under all ways of putting members of the pertinent 'many' into the extensions of the corresponding terms. 'Bruno is in the living room' is supertrue when your use of 'Bruno' singles out a certain portion of reality and all of the aggregates consistent with this singling out are in fact in the living room.

Importantly, now, the contextualist approach will imply that supertruth and superfalsehood are not symmetrical. For while there is only one way in which a judgment of the given form can be supertrue, it can be superfalse in two distinct ways, which we can formulate as follows:

A judgment of the form '$P(a_1, \ldots, a_n)$' is superfalse if and only if

either:

(F0) the judgment fails to impose in its context C a partition of reality in which families of aggregates corresponding to its constituent singular referring terms 'a_1,' …, 'a_n' are recognized,

or both:

(F1) the judgment successfully imposes in its context C a partition of reality assigning to its constituent singular terms 'a_1,' …, 'a_n' corresponding families of aggregates $F_1, \ldots, F_n$, and

(F2) the corresponding families of aggregates are such that, however we select individual f_i from the many F_i, '$P(f_1, \ldots, f_n)$' is false.

Suppose Bruno is in the kitchen but your bleary-eyed husband, looking at a cat-shaped piece of furniture in your living room, utters: 'Your cat is in the living room.' This judgment is then superfalse in virtue of (F0). There are no qualified feline-in-the-living-room portions of reality which are here able to sustain a partition of the needed sort. Suppose that you look at Bruno in the kitchen and utter 'Bruno is a unicorn.' This judgment is superfalse because (F1) and (F2) are satisfied. Your judgment does project successfully a family of aggregates of the appropriate sort, but it is false of every single one of those aggregates that it is a unicorn. The job of (F0) is to set to one side, before evaluation proper begins, spurious attempts at partition (for example those which result from a Strawsonian failure of presupposition, and especially from what Evans, in *The Varieties of Reference* (p. 123), calls 'the conniving use of empty singular terms').[6]

The precise workings of the above can be clarified only when the concept of partition and the associated concept of recognition (setting into relief, foregrounding) of an object by a partition have been made more precise (see below). For the moment, however, our present formulations will still suffice to show why judgments commonly held to fall down supervaluational truth-value gaps do not in fact do so.

6. This suggests also an alternative treatment of sentences falling within the orbit of (F0): that they do not express judgments at all, and thus not have a truth-value.

12. Fred's House

In support of the view that everyday judgments may manifest supervaluational truth-value gaps, Lewis has offered the following example, which we can think of as his best case. A stranger is visiting Fred's house. She raises the question of whether or not the garage attached to the house is or is not a *part* of the house. You can tell lots of supertrue stories about Fred's house without ever considering whether or not the house includes the garage as part. As Lewis puts it,

> you say that a famous architect designed Fred's house; it never crossed your mind to think whether by 'house' you meant something that did or that didn't include the attached garage; neither does some established convention or secret fact decide the issue; no matter, you knew that what you said was true either way (1993, p. 172).

Consider now:

> [P] The garage is not a part of Fred's house.

This sentence would ordinarily be taken to fall into a supertruth-value gap, because a corresponding judgment is true under some ways of making the unmade decision as to whether or not the garage is a part of Fred's house, and false under others.

But Lewis's best case is not yet made. For he still needs to describe in coherent fashion a context C in which [P] would serve as vehicle for judgment. And he needs to do this in such a way that both (i) the strict conditions he himself imposes are indeed satisfied in C, and (ii) the resultant judgment is assigned the value *true* on some evaluations consistent with C and *false* on others.

It is however difficult to imagine contexts in which (i) and (ii) are satisfied together. This is because most naturally occurring contexts impose strong constraints on the sorts of families of aggregates available for supervaluation. We can certainly conceive of contexts in which [P] is judgeable and in which (i) is satisfied but (ii) fails. That such contexts are conceivable is of course of no import in establishing that gap-sustaining contexts cannot exist for [P]. To describe one such context will however help us on our way to establishing the latter claim.

The hurdle we need to overcome is that of sincerity: why, if the facts are as Lewis describes, should anyone use [P] to express a genuine judgment? Well consider this: you are approached by the stranger and you assert [P] because you dimly remember the plans of Fred's house as including a boundary dividing what is labeled 'garage' from what is labeled 'house' along the plane where they meet. You thus presuppose, as concerns the reality within the vicinity of Fred's house, that it contains, prior to your judgment, a boundary of the given sort. If the conditions—'no established convention or secret fact'—as imposed by Lewis are satisfied, however, then there is no such

boundary in reality. Hence your attempt to impose a partition of the given sort fails, and your judgment is superfalse, by (F0).

13. Lake Constance

Matters are complicated in the above by the fact that the boundary you believe to exist is not of the bona fide, physical sort that is illustrated by the walls or roof of Fred's house. Rather, it is a fiat boundary, the sort of boundary by which census tracts or postal districts are typically demarcated, and the latter are the result of a more or less arbitrary imposition (Smith 1995, 2001; Smith and Varzi 2000). Such fiat boundaries are nonetheless, when once they have been established in a duly authorized fashion, real parts of reality (and as the history of international warfare shows, they can be parts of reality of great consequence for the life of humanity). Because fiat boundaries are the results of arbitrary imposition, however, the temptation arises to suppose (1) that there is no fact of the matter as to where they lie or do not lie, and (2) that it is of no consequence if they are infiltrated into reality in idiosyncratic ways.

The issue in hand can be clarified by adverting to a real case of the (fiat) boundarylessness of the sort Lewis's fictional example involves. Switzerland, Germany, and Austria meet in the heart of Europe somewhere in the vicinity of Lake Constance. As it happens, there is as of this writing no international treaty which establishes where, in or around the Lake, their respective borders lie (a state of affairs which still occasionally gives rise to disputes, for example as concerns fishing rights in different portions of the Lake).[7] Suppose, now, that you point to a certain kilometer-wide volume of water in the center of the Lake, and you assert:

[Q] That water is in Switzerland.

Here, too, there is no established convention or secret fact which decides the issue. What this means, however, is not that [Q] asserts a truth on some precisifications and a falsehood on others. Rather, again by criterion (F0) above, [Q] is simply (super)false. Whoever uses [Q] to make a judgment in the context of currently operative international law is making the same sort of radical mistake as is Matutinus when he judges that the Evening Star is brighter than the Morning Star. For in both cases *reality is not such as to sustain a partition of the needed sort.* It is thus as if the relevant judgment

7. In fact Switzerland, Germany and Austria officially embrace distinct positions on the matter. Switzerland takes the view that the border runs through the middle of the Lake. Austria takes the view that all three countries have shared sovereignty over the whole Lake. And Germany, insofar as its pertinent pronouncements can be given a single coherent interpretation, takes the view that it takes no view on the matter.

does not even reach the starting gate as concerns our ability to evaluate its truth and falsehood via assignments of specific portions of reality to its constituent singular terms.

And [P], too, is simply (super)false, for an exactly analogous reason.

This, now, tells us more clearly what the friend of supervaluational truth-value gaps needs to find in order to support her position. Such gaps can arise only if (T1) (and thus also (F1)) is satisfied. The proponent of gaps thus needs to find an example of a sentence, for example of the form 'P(a_1, a_2),' and of a non-fangled context C in which that sentence is used to make a judgment, which is of such a sort that C determines families, F_1 and F_2, comprehending pairs of aggregates: f_1 and f_1', and f_2 and f_2', respectively, which are such that P(f_1, f_2), and not-P(f_1',f_2'). The corresponding precisifications are, we might say, *authorized* within the context C.

Well consider:

[R] This slurry is part of Mont Blanc

uttered by a mountaineer pointing to an aggregate of moistened rock which is on the point of beginning to move slowly down the mountain face. Take $f_1 = f_2$ = the slurry, f_1' =Mont Blanc plus the slurry, and f_2' = Mont Blanc minus the slurry. Set P = *is part of*, and we are home.

Again, however, there is one further task which needs to be performed by the friend of gaps. For the latter still needs to describe in coherent fashion some non-fangled context in which [R] would indeed be used to make a genuine judgment. It seems, however, that as soon as he succeeds in describing such a context—for example our mountaineer is at the same time a government surveyor authorized to determine where Mont Blanc ends for purposes of regulating the extraction of minerals—then a gap-eliminating argument along the lines of the one presented above for the cases of Fred's house and Lake Constance can be rolled out once again.

There are also temporal analogues of the same phenomenon. Suppose Geraldine is dying, and that her final death pangs lie so close to midnight that a reasonable case can be made for the thesis that 'Geraldine died before midnight' is neither supertrue nor superfalse (compare Heller (2000)). Such a case, too, presents a problem for the gap-eliminator only if he is presented also with a naturally occurring context in which the corresponding sentence will be used to make a judgment, and this in such a way as to preserve its putatively gappy status. But such a context will be difficult to find. Suppose, for example, that the issue of the precise time of Geraldine's death plays a crucial role in the execution of her will. When it comes to the evaluation of the given judgment the relevant authorities will then necessarily have recourse to some gap-eliminating procedure, for example

involving inspection of the entry under 'time of death' on Geraldine's death certificate. Truth-value indeterminacy is once again eliminated via the mechanism of fiat boundaries.

Our goal here is rather modest. We aim merely to indicate some of the power of the contextualised theory, recognizing that the general project of gap-elimination will face problems—not least of the sort which arise out of the higher-order vagueness of terms like 'genuine' and 'succeeds' in the paragraphs above. We also concede that we will face problems when it comes to judgments expressed by means of predicates expressing first-order vagueness, for example in:

[S] John is bald

(for John a borderline case). Here, too, however at least some progress is made when the contextualization of judgment is taken into account. Thus consider the way in which sentences predicating flatness come to be assigned the value *true* in some contexts and *false* in others. Persons from Switzerland, for example, may truly assert, while persons from Holland with stricter standards truly deny, that a particular road is flat. The judgments made by the two sets of subjects will, in their respective contexts, be true absolutely: true and supertrue.

Similarly for a judgment to the effect that John is bald. John is a borderline case—someone who, in four years' time, will be entirely hairless but who, four year's ago, was still renowned for the simian abundance of hair on his head. There are contexts in which 'John is bald' might be used to make a judgment now, already, today. You have heard news of John's recent loss of hair and you jump too rapidly to conclusions. You see John in a bad light which exaggerates his recent loss of hair. You have an idiosyncratic use of 'bald' (the baldness cell in your partition of the space of hair-density states is larger than is standardly the case). In contexts in which things do not, in these various ways, go wrong, however, [S] is simply unjudgeable. *This is part of what we mean when we say that John is, as far as baldness is concerned, a borderline case.* In most contexts, therefore, even [S] will not be capable of expressing judgments susceptible to supervaluational truth-value gaps.

14. From Singular Reference to Judgment

The vagueness involved in our ordinary uses of terms like 'the sun', 'this cloud', or 'Bruno' is, be it noted, not in any way diminished by the fact that it is hard to produce judgments whose truth-value would be correspondingly indeterminate. That Bruno is losing or gaining molecules from one moment to the next is of no consequence for our everyday purposes: it falls below our normal threshold of concern. And even where this threshold is breeched, the natural contexts in which we live—including institutional contexts such as courts of law and university examining boards—have

understandably thrown up constraints against the easy formulability of judgments marked by truth-value indeterminacy. The vagueness underlying our singular reference is thereby tamed for purposes of judgment.

The threshold is breeched, most certainly, in the case of Fred's house with and Fred's house without its garage, and in the case of Switzerland with and Switzerland without its claimed portion of Lake Constance. These are far from being 'almost identical.' 'Fred's house' and 'Switzerland' project upon reality in ways which are radically diffuse, and the same applies also to terms like 'Christmas' or 'Mary's wedding.' If I say, 'We'll have a party over at Fred's house next Christmas,' then you will know what I mean even where the underlying projective diffuseness is still at its most radical. It will in due course—to the degree that is necessary to ensure that the guests will know when and where to arrive for the party—be moderated, via a process which we can think of as one of zooming from a less to a more refined partition.

15. Truthmaking and Vagueness

Imagine that John kisses Mary on her cheek at noon on a particular Wednesday in your favorite café. John often kisses Mary on the cheek, but kissing Mary on the cheek at noon on this particular Wednesday is something he can do only once. Suppose you are sitting in the corner of the café, observing John and Mary. Almost exactly at noon you judge that John is kissing Mary. Your judgment is a particular mental episode. In virtue of what is it true? Once again, our answer is, in first approximation: a certain oddly demarcated portion of reality, which might be labeled *John's kissing of Mary*. Once again, however, we face the problem that the relation between a judgment and the pertinent truthmaking portion of reality is more complex than might at first appear to be the case. For just as a simple referring term projects diffusely upon whatever is the pertinent foregrounded portion of reality, so also do the sorts of (true, empirical, logically non-compound) judgments that concern us here project diffusely upon whatever is the pertinent, truthmaking, portion of reality.

The projective relation is carried first of all by the *verb* of the relevant sentence—and just as it is through judgments that terms are projected onto reality by intending subjects, so also it is through judgments that verbs are projected onto reality by those same intending subjects. It was this insight that led Mulligan *et al.*, to elaborate a view of truthmaking according to which

> what makes it true that Socrates died is Socrates' death, what makes it true that Amundsen flew to the pole is his flight, what makes it true that Mary is smiling is her (present) smile, and so on. Or, in other words, ... for many simple sentences about spatio-temporal objects the truth-makers

> for these sentences are the moments [tropes, events, individual accidents] picked out by gerundials and other nominalised expressions closely related to the main verbs of the sentences in question … If all atomic sentences contain a main verb, and all nominalisations denote moments, then it would follow, in fact, that all truth-makers are moments, that what makes it true that *a* is F is *a*'s being F, what makes it true that *a* R's *b* is *a*'s R-ing *b*, and so on. (Mulligan *et al.*, 1984)

But as in the case of our singular reference to Fred's house, so also in the case of our judgments and that which makes them true. Here, too, eligible truthmaking portions of reality may overlap mereologically only to a minimal degree. This is because the truthmaker role can be played not just by parcels of reality that are 'almost identical' to the more obvious candidate real-world correlates of the corresponding verbs, but also by a variety of larger and smaller mereological fusions—as for example, in the case in hand, by the mereological sum of John and the kissing, of Mary and the kissing, of John and Mary and the kissing, of all of these summed together with John's prior handing over of the flowers, and so forth.

Every truthmaking portion of reality must in the given case include a certain kissing event as part. It is this latter, more than anything else, which is responsible for the truth of your judgment that John is kissing Mary. But which event (which portion of reality) should this be? When, for example, does the pertinent kissing begin? Does the kissing event include the Mary-related thoughts that are simultaneously passing through John's mind or the quickening of his heartbeat? Surely, there is not just one parcel of reality that is more deserving than every other of playing the role of truthmaker for the given judgment. Rather, there are many parcels of reality which play this role, some of which differ just a little from each other. Moreover, we encounter a no less opulent abundance of mereological fusions qualified to serve as truthmakers for the given judgment if we move in the opposite direction and attempt to isolate smaller truthmaking portions of reality in the region where John makes contact with Mary's cheek. This is because John's kissing takes time, and so there are many more or less instantaneous slices of the kissing, all of which are equally qualified to serve as truthmaker for the judgment in hand.

16. The Granularity Problem Posed

How, now, do the two distinct occurrences—John's kissing Mary, and your simultaneously executed judgment—relate to each other? Your judgment is a perceptual judgment: it is a case of *seeing that*, and it can be compared, once again, to a beam of light that picks out a certain portion of reality. It

thereby partitions the world into two parts: the complex token event which is *John's kissing Mary*, on the one hand, and its complement, comprehending all of those portions of reality that do not fall within the scope of your projection, on the other.

But again: this partition is effected, not in any simple geographical way, but rather in such a way as to be marked by a certain granularity. The molecules in John's knees are not part of the object of your perceptual act. This, however, creates a problem for all of those, like the present authors, who believe that mereology is the indispensable instrument of any coherent ontology; for the relation of a part to its whole is transitive. Consider a truthmaker *x* for your judgment *p* to the effect that John is kissing Mary. Suppose that *x* consists of all those parts of John, Mary and the relational event of kissing (however this latter is determined) which are involved in making it true that John is kissing Mary. We can now construct the following inconsistent triad:

A. The molecules inside John are parts of John.

B. John is part of *x*.

C. The molecules inside John are not a part of *x*.

Some readers may prefer to reconstruct this trilemma in relation to facts, or states of affairs, or situations, or in relation to whatever is their favored candidate truthmaking entity. And counterpart trilemmas can be constructed also in regard to a host of other familiar animals in the philosophical zoo: sense data, surfaces, aspects, pluralities, visual fields, persons, Husserl's noemata, Kant's 'phenomenal world,' Fine's 'qua objects', and so on.

The trilemma makes itself felt also in extraphilosophical contexts. Indeed the transitivity of the part-whole relation threatens to give mereology something like the same corrosive force in relation to our common-sense ontology (of persons, hands, clocks, etc.) as is possessed by skepticism in relation to our common-sense theory of knowledge. Mereology thereby forces many to conclude that the only good ontology (and science) is one of ultimate (partless) atoms.

If, on the other hand, some way can be found to resolve the trilemma, and so block the transitivity of parthood,[8] then we may be in a position to do justice to entities of the mentioned sorts mereologically after all. In this way we can make

8. This will imply also the rejection of Varzi's "No Double Counting Principle" according to which an inventory of the world is to include an entity *x* if and only if *x* does not overlap any other entity *y* that is itself included in that inventory (Varzi 2000).

safe for ontology the middle world of ordinary objects (and so also make room, within the orbit of the truthmaker theory, for the truths of the middle sciences of biology, linguistics, geography, military history, land surveying, and so forth).

Standardly, of course, our problem is avoided by refusing to treat such entities seriously *as entities*, for example by appeal to the notion of a 'category mistake.' One is then simply not allowed to ask, for example, whether molecules of paint are or are not a part of the sense data which John sees when he focuses on a painted wall. The sentences which might serve as answers to such questions were indeed for a long time rendered unjudgeable in certain contexts.

Braver souls will however persist in raising such questions nonetheless. One standard answer utilizes the phrase 'under a description' or some comparable locution. For a solution along these lines to work, however, it must be true that this molecule is part of John under one description (for example: *physical body*), but that it is not a part of John under some other description (for example: *object visible with the naked eye*). If, however, John under these different descriptions is one and the same entity, then he thereby also has, under each description, all the same parts. If, on the other hand, John under this description is a different entity from John under that description, then we are still in need of an account of how this difference is to be understood, and this brings us back to the trilemma with which we began.

A more promising starting point for the resolution of our trilemma rests on set theory: the set–membership relationship is after all not transitive. But to use set theory as a means of blocking transitivity brings for our purposes too great a cost. For if set theory is taken realistically, then this forces us to identify elements (*Urelemente*) from out of which the larger structures which concern us would be constructed by set-theoretical means. But what would such elements be in the case of a complex event such as John's kissing Mary? And even where appropriate elements—for example atoms or molecules—do seem to present themselves for purposes of set-theoretic construction, problems arise because we then find that our ontology is either constrained, once again, to work exclusively with putative ultimate atoms—entities as yet unknown—or it is cluttered up with multiple copies of reconstructed objects existing at different levels of granularity (for example, John and Mary as sets of atoms, as sets of molecules, as sets of cells, and so on). This problem of supernumerary copies does not arise for mereology, of course, since the mereological sums of the atoms, molecules, cells, etc. constituting Mary are all one and the same object. It is precisely this, however, which makes the mereological approach susceptible to our trilemma.

We note in passing that standard set-theoretical semantics avoids the problem of granular copies in a different way: by spurning ontological realism. For such semantics deals not with John and Mary (and the kissing event by which they are briefly joined together) at all, but rather with denatured substitutes. For the domain of each pertinent model has as its members not John and Mary or any other denizen of the ever-changing world of flesh and blood reality, but rather abstract counterparts thereof, tailored to play a certain algebraic role. The trilemma does not arise, because what is not made of flesh and blood can have no flesh and blood parts. On the other hand, however, because standard semantics thus operates with surrogates, it leaves unanswered those questions which concern the ways in which our judgments relate to objects in reality.

17. The Granularity Problem Solved

Consider what happens when you observe a chessboard. You are working with a partition of the world into that, in the region of the chessboard, which you are focusing on, and that which is traced over. Your focus brings with it a certain granularity: you are interested, not in the atoms or molecules within the board and its pieces, but rather only in the board and pieces themselves. Moreover, you are interested in the latter not as constituting a mere list, or set, but rather as they exist within a certain arrangement. The board is divided into cells (squares). In some of these cells pieces of specific kinds are located. To understand what is going on here, we need to focus in more detail on the notion of partition and on the associated notion of cell. The first thing that we recognize is that partitions have their granularity built in, as it were from the very start. A map of France depicting its 91 *départements* or its 311 *arrondissements* is a good illustration of what is meant by a partition in our intended sense. Such maps are the result of applying a certain coarse- or fine-grained grid of cells—the minimal units of the partition—to a certain portion of reality.

Partitions are at work, we now want to claim, whenever judgments are effected in relation to the empirical world of what happens and is the case. For a partition to do its work, it needs to have cells large enough to contain the objects that are of interest in the portion of reality which concerns the judging subject, but at the same time these cells must somehow serve to factor out the details which are of no concern. A partition, as here conceived, is accordingly a device for focusing upon what is salient and also for masking what is not salient. We can think of it as being laid like a net over whatever is the relevant object-domain, and, like a net (or a latticed window of the type employed in Alberti's reticular painting grid), it is to a large degree transparent. Thus, importantly, it does not in any way change the reality to which it is applied.

This reality, and each of the objects within it, is what and where it is, and it has all its parts and moments, independently of any acts of human fiat and independently of our efforts to understand it theoretically. A partition, by contrast, is precisely an artefact of our judging, classifying, theorizing, or mapping activity.

We can now begin to see how the notion of partition can help us solve the problem of granularity. Granularity as it has been treated in the above is properly at home only in the fiat realm: it pertains not to the objects themselves on the side of reality but rather to the ways we partition these objects in different contexts. All three clauses of the trilemma will be retained; now, however, they will take the following forms:

A*. The molecules inside John are parts of John.

B*. John is recognized by a partition associated with the judgment *p*.

C*. The molecules inside John are not recognized by a partition associated with the judgment *p* (the available partitions lack appropriately fine-grained cells).

The resolution works because we are in effect exploiting an analogue of the transitivity-blocking feature of set theory. And if the notion of a partition is in some respects a generalization of the notion of set, so the notion of a cell is correspondingly a generalization of the notion of singleton. Where, however, the elements of a set exist within the set without order or location—they can be permuted at will and the set remains identical—a grid or partition standardly comes with a specific order and location of its constituent cells. The latter fit together in a determinate arrangement, like pieces in a jigsaw or like molecules in a strand of DNA.

This arrangement may be purely spatial, as in a map, where the relative positions of neighboring cells are determined by the corresponding positions of those portions of geographic reality to which the cells relate. Or it may be determined by a linear ordering, as for example where partitions are determined via quantitative scales reflecting age cohorts or tax brackets or frequency bands. The arrangement may also be determined in more complex (for example hierarchical) ways, as in the case of a partition determined by *kinds* or *concepts* (for example a partition of the animals in your local zoo into *lions*, *tigers*, *giraffes*, *small marsupials*, etc.). The partitions which come closest to mere lists are those associated with our uses of proper names. These we shall call *nominal partitions* in what follows, and we note in passing that even nominal partitions may deviate structurally from the corresponding sets in being non-extensional. Thus Peter, author of the puzzle sentences [F] and [G] above, employs a nominal partition containing *inter alia* two distinct cells, both of which are labeled ‘Paderewski’ and both of which pick out the same object in reality. A nominal partition

relating to continuants is marked in addition by the fact that its cells are able to keep track of the corresponding objects in the world as objects that are identical from one moment to the next in spite of the fact that molecules are gained and lost.

Complex multidimensional partitions may rest upon combinations of these different types of cell arrangements. A map of the zoo, for example, might indicate not only the places where animals are located but also the sorts and sizes and proper names of the animals which are located in those places.

18. Better than Sets

Set theory rests on one central relation: the relation between element and singleton. This relation is, as Lewis notes, enveloped in mystery:

> since all classes are fusions of singletons, and nothing over and above the singletons they're made of, our utter ignorance about the nature of the singletons amounts to utter ignorance about the nature of classes generally. … What do we know about singletons when we know only that they are atoms, and wholly distinct from the familiar individuals? What do we know about other classes, when we know only that they are composed of these atoms about which we know next to nothing? (1991, p. 31)

Our machinery of partitions, in contrast, rests not just on one central and mysterious relation between element and singleton, but rather on a whole series of location relations between objects and cells, each one of which is independently familiar and well understood. The relation between an object and its proper name is one such. Others include the relation between an object and its spatial location, or between an object and a concept under which it falls or a kind to which it belongs, and also the various relations which an object may bear to intervals on quantitative scales of different sorts (addressing height, velocity, quantum number, and so on). There are also more transient relations, for example the relation between an object and your visual field. It will turn out that each of these relations between an object and a cell corresponds to one or other of the ways in which, when we make a true judgment, we bring about a corresponding partition of the reality to which our judgment corresponds. And while partitions, and the cells by which they are constituted, are artefacts of our cognition, when once a given partition exists, it is, for each cell in the partition and for each object in reality, an objective matter whether or not that object is located in that cell.

19. Towards a Formal Theory

Let variables $x, x', x_1, y, \ldots$ range over objects, $z, z', z_1 \ldots$ over cells, and $A, A', A_1 \ldots$ over partitions. The cells in a partition may have sub-cells. Thus for example the cell *rabbit* is a sub-cell of the cell *mammal* in a partition of the animal kingdom. The cell *Florida* is a sub-cell of the cell *United States* in the standard geopolitical partition *G* of the surface of the globe. We write:

$$z \subseteq_A z'$$

as an abbreviation for: z is a sub-cell of the cell z' in the partition A.

$\subseteq_A$ is reflexive, transitive and antisymmetric. It defines a partial order on the totality of cells in the partition A, by analogy with the usual set-theoretic subset relation. We stipulate further that it satisfies the finite chain condition to the effect that if $\ldots \subseteq_A z_1 \subseteq_A z_0$, then there is some n such that $z_n = z_{n+1} = \ldots$. An example of such a finite chain is your address (The Oval Office, The White House, 1600 Pennsylvania Avenue NW, Washington, DC 20500, USA).

We can define the property of being a minimal cell within a partition in the obvious way as follows.

DMC $\quad MC_A(z) =: z \subseteq_A z \wedge \neg\exists z'(z' \subseteq_A z \wedge z' \neq z)$,

where '$z \subseteq_A z$' is merely a convenient way of rendering: z is a cell in the partition A.

The counterparts of sets within our present framework are then those partitions which can be identified as the mereological fusions of minimal cells in the sense defined, so that minimal cells play the role played by singletons in Lewis's *Parts of Classes*. The corresponding partitions then have minimal cells which represent a jointly exhaustive and pairwise disjoint tiling of the pertinent domain of objects, and a cell z in such a partition A safisfies the following:

MC $\quad \exists z_1 \ldots \exists z_n(MC_A(z_1) \wedge \ldots \wedge MC_A(z_n) \wedge z = z_1 \cup_A \ldots \cup_A z_n)$,

where '$\cup_A$' symbolizes mereological fusion of cells within a partition A.

MC does not hold of partitions in general however. This is because the latter, as artefacts of our cognition, may be incomplete. Thus we can imagine a partition of the animal kingdom containing a cell labelled *mammal*, and other cells labelled *rabbit*, *dog*, etc., which is yet not such as to represent a complete accounting of all the species of mammal which exist.

In addition, partitions do not in general satisfy the usual set-construction principles of union, intersection and complement. However, they satisfy restricted versions of these principles along the lines set out in Smith (1991). Thus the union $z \cup_A z'$ of two cells in a partition A, which we are to conceive as a $\subseteq_A$-minimal cell satisfying the condition that it contains both z and z', is not in general defined. (Consider, again, the partition G, and take z = *Florida*, z' = *Zambia*.) Where it is defined

it will be, for most naturally occurring partitions, unique. (As applied to *Cyprus* and *Malta*, for example, it currently yields the unique output: *British Commonwealth*.) But it is not in general unique. (Cyprus and Malta are both candidates to join the European Union.)

$\cup_A$ is commutative, but it is not associative. That is to say $(z \cup_A z') \cup_A z''$, even where it is uniquely defined, is not in every case identical to: $z \cup_A (z' \cup_A z'')$.[9] We do however have the familiar equivalence of $z \subseteq_A z'$ and $z \cup_A z' = z'$.

Regarding intersections, we first of all define what it is for two cells of a partition A to overlap in A, as follows:

Do $z_1 \, o_A \, z_2 := \exists z(z \subseteq_A z_1 \wedge z \subseteq_A z_2)$.

The intersection of two cells in A is then defined as any $\subseteq_A$-maximal cell which is included as sub-cell within them both. If two cells overlap in a partition A, then these cells do indeed possess an intersection within A. Once again, intersection as thus defined is commutitative but it is not in general unique or associative.

Regarding complements, we set $-_A z$ to be a $\subseteq_A$-maximal cell which does not overlap with z. The complement of a cell, too, is not in general defined, and even where it is defined it is not in general unique. Consider, once again in relation to our geopolitical partition G, the question as to the $\subseteq_A$-maximal political entity which does not contain Florida as part.

That the complement of a cell is not in general defined goes hand in hand with the fact that there is no analogue of the empty set in the theory of partitions. That is to say, there is no cell which is empty *per se*. Even where a partition has a unique and exhaustive maximal cell, there is no complement of this maximal cell within that partition. Each partition will however characteristically contain cells which are empty *per accidens*—because they have no objects located in them (as a chessboard will contain squares empty of pieces, and as a hotel may, on any given night, contain rooms empty of guests). *Dodo* is an empty cell in one standard partition of the animal kingdom.

It will be clear that, while what we are offering is intended as a generalization of Lewis's (1991) conception, our theory is in many respects weaker than his mereologized set theory. Thus it is (for present purposes) without anything like the possibility of cells of higher order within which cells of lower order would be located. Our hierarchy of partitions is thus, like Schroeder's theory of

9. Consider a partition with cells as follows: $\{z\}$, $\{z'\}$, $\{z''\}$, $\{v\}$, $\{w\}$, $\{z, z', v\}$, $\{z', z'', w\}$, $\{z, z', z'', w\}$, $\{z, z', z'', v, w\}$.

manifolds, truncated at the first type, the sacrifice in mathematical power being compensated for by a gain in ontological realism.

20. Objects in Cells

We write:

$$L_A(x, z)$$

for 'object x is located in cell z in partition A.'[10] Objects in and of themselves satisfy the standard axioms of mereology, for example as set out in Simons (1987). Here we are concerned with what happens when objects are viewed through cells and partitions.

The following appears to be an attractive axiom governing L:

ALo $\quad L_A(x, z) \wedge L_A(x, z') \Rightarrow z \, o_A \, z'$

An object is never in two cells which do not overlap. (Here '$p \Rightarrow q$' abbreviates: it is not possible for p to be true and q false.) From this it follows in turn that if an object is in two distinct cells within a partition, then these cells are not both minimal, and they possess an intersection-cell.

For some partitions, which we can call *distributive*, if object x is a part of object y, and if y is located in a cell z, then x is also located in that cell:

Ddist $\quad dist(A) := \forall x \forall y \forall z(x \leq y \wedge L_A(y, z) \Rightarrow L_A(x, z))$,

where '$\leq$' abbreviates: 'is a proper or improper part of' understood according to the usual axioms of classical extensional mereology (Simons 1987). Distributive partitions satisfy the following principle, according to which, if two objects are located in two different cells, then the sum of these objects is located in the sum of these cells:

$$L_A(x, z) \wedge L_A(x', z') \Rightarrow L_A(x + x', z + z').$$

Spatial partitions are always distributive in the sense specified. If John is in Salzburg and Mary is in Salzburg, then their sum is in Salzburg and so, too, are all their bodily parts. A set, on the other hand, is a simple example of a non-distributive partition, and a partition generated by kinds or concepts, too, is non-distributive. A partition recognizing cats does not *ipso facto* recognize parts of cats. Moreover, if Bruno is a cat and Tibbles is a cat, then the mereological sum of Bruno and Tibbles is not itself a cat.

We can define the notion of recognition that is at work here as follows:

---

10. Our idea of location is inspired on the one hand by Casati and Varzi (1999) and on the other hand by Omnès (1994), whose theory is summarized in Smith and Brogaard (2002).

D∈ $x \in A := \exists z(\mathrm{L}_A(x, z))$

That an object x is *recognized* by a partition A means: x is located in some cell z in A.

If x is located in a cell z in A and if y is a part of x that is recognized by A, then y is located in z:

$$\mathrm{L}_A(x, z) \wedge y \leq x \wedge y \in A \Rightarrow \mathrm{L}_A(y, z)$$

Suppose John is recognized by a nominal partition A consisting of a single cell labeled 'John,' so that John $\in A$. This, as we have seen, is consistent with its being the case that a whole family F of crisp aggregates $f_1, \ldots, f_n$ are such that for each i, $f_i \in A$. This is because A does not care about the small (molecule-sized) differences between the different f_i. The cell *John* captures *all* the aggregates f_i which are almost identical to John.

21. Judgments and Partitions

Objects as they exist in nature stand to each other in various relations, they have hooks of various sorts which link them together; these include common boundaries (both fiat and bona fide) and relations of dependence and of functional or causal association. The operator of mereological fusion, when properly handled, preserves these inter-object relations, and it thus preserves the order and location of objects which fall within its grasp: if two objects are linked together in nature, then they are linked together also within their mereological fusion.

A set (class) is a mereological fusion of singletons, and mereological fusion preserves order and location. How can it be, then, that the elements within a set can be permuted at will and the set remain identical? The answer is that the set is built up mereologically not out of elements but out of singletons, and the latter are mere homeless *somethings*, outside time and space. The singleton operator has the effect of stripping away the various sorts of linkages which obtain between the objects to which it is applied and also of setting them apart from their surroundings and from time and change.

Partitions are distinct from sets and fusions however in that they are not constituted out of the objects that are located in their cells. Rather, they belong to the level of our theorizing and classifying activity. They can remain the same even though the objects towards which they are directed are subject at more fine-grained levels to constant change. Some partitions, now, are like sets in the sense that they will apprehend the objects which are located in their respective cells independently of order or arrangement or linkage or time. Other partitions, however, will inherit from mereology the ability to comprehend their objects in ways which map the different kinds of

relations that obtain among them. The cells in such partitions project their objects not in isolation, but rather in tandem with other objects located in related cells within the same partition. Consider, for example, those two-celled partitions which capture the relations between a part and its whole or between a substance and its accident. Such partitions apply to pairs of entities in reflection of specific relations in which the latter stand to each other. John and Mary, before they wed, are not, but through marriage they become, located in a two-celled partition of the type: *married pair*. Yet other two-celled partitions, for example the partition captured by our use of paired demonstratives such as *this* and *that*, *here* and *there*, or *left* and *right*, or *first* and *second,* apply to pairs of objects only in reflection of our ways of relating to them intentionally. A three-celled partition might for example capture the way in which, in an action of kissing or congratulating, two objects become bound together by a third object—a relational event—in which the one occurs as agent, the other as patient.

Some partitions are such as to possess cells which are indeterminate in number, so that there is a sense in which they capture the corresponding objects in reality in indeterminate fashion. Consider the way in which, in quite different (more or less scientific) sorts of contexts, we have partitions of biological reality which divide the latter (perhaps in a variety of slightly different ways) into *biological species*. The latter are built up out of cells having labels like: *the species cat*, *the species rabbit*, and so forth. These in turn are contained as sub-cells within larger cells such as *the sub-kingdom vertebrate, the class mammal*, *the phylum chordata*, and so forth. The reality located in the cell labeled *the species cat* is at any given time the mereological fusion of all whole, live cats. As seen through the lens of the given partition this total fusion is parceled out into individual cats (and not for example into parts or aggregates of cats). This parceling out is however effected in a completely indeterminate fashion, which is to say: in such a way that the partition does not know (or care) how many cats there are within its orbit nor where these cats are located. The partition will in addition trace over all the individual differences between all the different cats which fall within its scope. In this way the partition is able to keep track of all the cats in the world as forming a whole (a species) that is identical from one moment to the next even in spite of the fact that individual cats are born and die.

22. Partitions and Truth

To see how true judgments are associated with partitions of different sorts in different contexts we consider the following cases, which we have selected at random.

You make a judgment to the effect that event *e* occurs before event *e′*. A two-celled nominal partition is then associated with your judgment, containing cells which recognize *e* and *e′*, respectively. The judgment is then supertrue if and only if, in whatever way we select from the parcels of reality acknowledged by the first cell, the reality selected occurs before all the parcels of reality acknowledged by the second cell.

Consider a judgment to the effect that Harvey is a rabbit. Here two partitions are involved. The first is a nominal partition of reality containing a cell which recognizes Harvey. The second is a partition, drawn from some standard larger partition of the animal kingdom into *hare*, *kangaroo*, and so on, which contains the single cell *rabbit.* The given judgment then asserts, roughly, that whatever is located in the cell *Harvey* in the first partition is located in the cell *rabbit* in the second partition. An analysis along similar lines applies also to judgments such as: 'John is male,' 'John is in Salzburg,' 'John is over six feet tall,' and so on.

A case like 'John is kissing Mary' is more complex. Here (very roughly) three partitions are involved: [1] a partition of the space of those human beings known to the judger, which can be assumed for our purposes to contain just two cells labeled *John* [1a] and *Mary* [1b]; [2] a partition of the event kingdom containing a cell labeled *kissing*; [3] a three-celled partition whose cells are linked together in such a way that the parcels of reality foregrounded in [1a], [2] and [1b] are themselves linked together as *agent*, *action* and *patient*, respectively. The given judgment then asserts (again very roughly) that whatever is located in the John and Mary cells in partition [1] and in the *kissing* cell in partition [2] is at the same time located in the corresponding cells of the relational event partition [3].

The case of 'John kissed Mary' is still more complex. We are to imagine the given sentence as having been uttered by a judger who has no closer acquaintance with the many prior occasions on which Mary was being kissed by John. As noted in Smith (1999) this case differs crucially from the present-tense perceptual judgment 'John is kissing Mary' in a way that makes it more troublesome for the truthmaker theory. This is because the relevant family of truthmakers now lies outside the immediate environment of the judger and there is no one appropriate portion of reality—containing some one suitably qualified kissing event—that here stands forth as being able to perform the truthmaking role. The problem can be solved, however, by calling in aid the notion of an indefinitely-membered partition which we introduced above. The judger must employ a partition which is such as to trace over the differences between the different events which would otherwise serve to play the truthmaking role. Here again the judger imposes upon reality a nominal partition

[1] with cells [1a] *John* and [1b] *Mary*. As far as *kissing* is concerned, he must now direct himself not to a single scene that is contemporaneous with his own judging activity, but rather to the whole history of relations between John and Mary extending back in time. He now imposes [2*] an indefinitely-membered partition upon this history which projects in indeterminate fashion upon the many pertinent kissing events therein contained. [2*] knows neither how many such events have taken place nor where or when they have occurred. Hand in hand with this, the judger then imposes, for each of the indefinitely many *kissing* portions of reality projected by [2*], a partition of type [3], once again involving members: *agent*, *action* and *patient.* The first and last members then point in each case to the enduring objects John and Mary captured by [1a] and [1b], while each successive *action* member points to a portion of reality picked out by a corresponding *kissing* term. In sum, the judgment 'John kissed Mary' projects onto reality in indeterminate fashion all the many kissing events involving John and Mary which have ever taken place, and it traces over the differences between them.

The two judgments 'John is kissing Mary' and 'John kissed Mary' illustrate a difference between specific and generic judgments that is analogous to the *de dicto*/*de re* distinction in the literature on belief. Specific judgments project onto specific portions of reality—which are characteristically present to the judger and form part of the context of his judgment. Generic judgments, in contrast, partition the world in a more diffuse way, and in a way that will characteristically involve portions of reality not present to the judger. Examples of specific judgments are: I am hungry, that bird is flying high, the tables look like they need dusting. Examples of generic judgments are: Armadillos exist, there is life on distant galaxies, there are neutrinos passing through my body.

But there are also mixed cases. You point to the gum wrapper you see before you in the street and you judge: The person who dropped that gum wrapper is a careless slob. Or suppose Jones is murdered at a dinner party. You gather all the guests into the library and say: I don't yet know who the murderer is. I don't even have a suspect. But I have deduced that the murderer, whoever he is, is the one and only person in the library with a copy of the pantry key on his person. (Dennett, 1982) The object of your judgment is present to you; but your judgment is nonetheless of the generic sort. The partition labeled *the murderer* relates to its object in indeterminate fashion even though the relevant total object is simultaneously presented to you visually through partitions which are determinate in nature.

23. Truthmakers are Necessitators

If a given object x exists, and if the existence of x entails that a judgment p is true, then we say that x necessitates p. In symbols:

DN $\quad x \,\mathrm{N}\, p := \mathrm{E!}x \wedge (\mathrm{E!}x \Rightarrow p)$

Whenever x makes p true, then x necessitates p.

It is not always particular events or pluralities of events that are truthmakers for our empirical judgments. Consider the judgment that John is generous. The truthmaker of this judgment is, in first approximation, a particular aspect of John, something like a generosity trope. John's generosity is a trope that only he can have, for tropes do not migrate. The existence of this trope thus necessitates that 'John is generous' is true.

On more careful consideration, however, we see that the thesis that there is some one item in reality which secures the truth of 'John is generous' must be a simplified reading of the matter in hand. Rather we must say that the judgment in question imposes upon the reality around John and his life a partition whose corresponding *generosity* cell comprehends (once again in indeterminate fashion) a wide range of phenomena—all of those actions, attitudes, habits and dispositions which taken together are labeled 'John's generosity'—instanced at different times in his life. This cell serves to unify together under a single heading those of John's actions and attitudes which are relevant to the issue of his generosity. The factual material which is John's life is thereby comprehended in a quite specific way and in a quite specific direction, so that out of a congeries of disparate states, events and dispositions there is demarcated a unity of a certain sort, called 'generosity.' (Compare the way in which, out of the many diverse islands and promontories between Sweden and Germany there is demarcated a unity of a special sort, called 'Denmark.')

Similar kinds of comprehending and demarcating are at work in almost all our judgments. They may be at work already when we comprehend a certain congeries of psychological and physiological events in the neighborhood of John and Mary as a case of *kissing*. More nuanced varieties of comprehending and demarcating are at work in our judgments when we make distinctions of verbal aspect, for example between: 'John has kissed Mary,' 'John used to kiss Mary,''John was kissing Mary,' 'John resumed kissing Mary,' and the like.

Yet even when all of this has been said, one might still wonder why it is necessary to appeal to anything like tropes (including variously chopped about and unified events, states, and dispositions) in order to account for the ways in which subject-predicate sentences are made true. The reason is precisely that a truthmaker is a necessitator. John himself does not necessitate that he is generous. In a different possible world he might exist and yet not be generous at all. As Mulligan *et al.* (1984)

and Armstrong (1997) have argued, a truthmaker relation should not be an entity that merely in some circumstantial fashion brings it about in this actual world at this actual time that a given statement is true (as a thief might bring it about that you are disembarrassed of your wallet). If John himself were the truthmaker of 'John is generous,' then it would follow that John could not have been different in the generosity stakes than he in fact is. Similarly, John may of course play a role in making it true that John is kissing Mary, but he cannot do this alone.

On one option the truthmaking relation would be identified with the relation of necessitation—an exclusively world-to-mind complement of the semantic mind-to-world relation of projection. Here, however, we favor a usage according to which, while necessitation is indeed the heart of the truthmaker relation, truthmakers are necessitators of certain special kinds. For there are *malignant necessitators*—entities which are such that their existence entails the truth of a given judgment but for reasons which are skew to those involved in the the relation of making true. Examples might be: the successive necessitating acts of will of a Malebranchian God. If there are necessary truths, then every contingently existing object is a necessitator for all such truths. Smith (1999) sought to exclude malignant necessitators by stipulating that a truthmaker for a given judgment is a necessitator for that judgment which is in addition such as to satisfy a *projection constraint*. Roughly, it must be a necessitator which is in addition such as to fall within the totality of what that judgment is about.[11] In what follows we show how the theory of partitions yields a very natural formulation of a projection constraint along these lines.

24. Acknowledgment

Our judgments, as we have seen, come along with partitions of reality of various sorts, whose type, granularity and scope depend upon the contexts in which our judgments are made. Examples have been given in the previous section. The judgment 'John is kissing Mary,' for example, comes along with partitions like the ones listed under [1], [2] and [3] above. This relation between judgment and partitions is a complex one, and we have of course presented in the above only a scattering of

---

11. More precisely he introduced a relation of projection, the dual to the relation of necessitation, as follows:

$$x \mathrm{P} p := p \wedge (p \Rightarrow \mathrm{E}!x)$$

If a judgment p is true, and if p entails the existence of a given object x, then we say that *x falls within the projection of p*.

He then proposed the following definition of truthmaking in terms of necessitation and projection:

$$x \models p := x \mathrm{N} \exists y(x \leq y \wedge y \mathrm{P} p)$$

Or in other words: 'x is a truthmaker for p' means: x necessitates that there is some y containing x as part and falling within the projection of p.

examples which seem intuitively acceptable in relation to the specific cases mentioned. Given our general thesis, however, according to which there are standardly many (whole clouds of) truthmakers associated with any given judgment, even this opportunistic approach will suffice for present purposes.

We first of all introduce the relation of *acknowledgment*, which is to serve as our analysis of the relation of aboutness between a judgment and its objects. Each token act of judgment will then have associated with it some specific repertoire of partitions, along the lines described informally above.

This having been said, we now set:[12]

$x\ \mathrm{A}_1\ p :=$ x is recognized by some partition that is associated with the judgment p.

We then define, for $n \geq 1$,

$$x\ \mathrm{A}_{n+1}\ p := \exists yz(x{=}y{\cup}z \wedge y\ \mathrm{A}_1\ p \wedge z\ \mathrm{A}_n\ p)$$

We can now write:

DA $\qquad x\ \mathrm{A}\ p := \exists n\ (x\ \mathrm{A}_n\ p)$

x is acknowledged by p, in the simplest case, just in case p is true and there is a partition A such that A is associated with p and x is recognized by A. In more complex cases, x is acknowledged by p just in case x is a mereological sum of parts each of which is acknowledged by p.

A now satisfies the basic requirement of an adequate theory of truthmaking isolated in our discussion of the problem of granularity in the above, namely that it does not satisfy:

*If x A p and $y \leq x$, then y A p

For although John himself is acknowledged by 'John exists' in standard contexts, this is not the case of the molecules in John's ear.

The two relations of necessitation and acknowledgment are independent. From x N p we cannot infer x A p (take x to be your rabbit Harvey, and p to be the judgment 'Rabbit DNA exists'[13]); and from x A p we cannot infer x N p (take x to be Bruno and p to be the judgment 'Bruno is in your living room'). Neither N nor A will suffice alone to yield an account of the truthmaker relation.

---

12. Our formulations here are provisional only. In a more adequate treatment we shall need to distinguish the separate roles performed by the variables p, q, … by admitting, in addition to propositional variables, also variables ranging over both *token sentences* and *token acts of judgments.*

13. Or suppose (for the sake of argument) that the existence of a certain taxi receipt entails the prior occurrence of a certain taxi ride. Then the receipt necessitates the truth of 'this taxi ride occurred'. Clearly, however, there are contexts in which the receipt is not acknowledged by a judgment of this form. (We might think of the receipt as a *proof* that the ride occurred. The similarities between proofs and truthmakers have been expounded by Sundholm 1994.)

Rather, what makes a given judgment true is any object which both necessitates and is acknowledged by *p*.

D|= $x \models p :=: x \mathrm{N} p \wedge x \mathrm{A} p$.

One problem faced by the resulting theory is that it rests on the still only informally specified relation of 'association' between a judgment and its partitions. Another problem turns on the fact that we can no longer prove the otherwise attractive:

T|=|= *If $x \models p$, then $x \models (x \models p)$.

(If *x* makes *p* true, then it is *x* which makes this true: truthmaking is its own reward.)

T|=|= comes at a price, however, since it implies that, if the truthmaker relation holds in one instance, then it holds in an infinite number of instances. The realm of *truthbearers* thus becomes infinitely large, and this runs counter to our general policy, in the above, of conceiving episodic judging acts as the bearers of truth. Only judgments actually made (that is to say, if you read it right: all judgments) have truthmakers, on our present dispensation, since only judgments actually made are such that the corresponding objects in reality fall within the scope of partitions, and only partitions are in a position to do the job of carving up reality in the ways required to yield truthmakers. This is, as we shall see, no threat to the objectivity of truth. The truthmaker relation is precisely that: a relation. It holds between judgments and portions of reality. A portion of reality does not stand in this relation in and of itself, but only when a judgment comes along, with its associated partitions, to cast it into the light.

A more serious objection to D|= is the following. Suppose *p* is of the form '$q \wedge r$,' and suppose $x \mathrm{N} p$, but that *x* is acknowledged by only one conjunct of *p*. Take *x* = Restall's refrigerator, *q* = 'Restall's refrigerator exists,' and *r* = Fermat's last theorem. Then trivially:

i. $x \mathrm{A} q$,

from which we infer:

ii. $x \mathrm{A} q \wedge r$.

From DN we then derive immediately, for the given *q* and *r*:

iii. $x \mathrm{N} q \wedge r$,

whence, from ii. and iii. and the definition D|=, we have:

iv. $x \models q \wedge r$,

from which finally we can infer:

v. $x \models r$,

or in other words: Restall's refrigerator is a truthmaker for Fermat's last theorem. (Compare Restall 1996)

Only the steps from i. to ii. and from iv. to v. are non-trivial; yet both seem to be supported on intuitive grounds. For it seems reasonable, surely, to suppose that if x is recognized by a partition associated with a judgment q, then this same partition will be associated also with any judgment of the form '$q \wedge r$.' And similarly it seems reasonable to infer from the premise that x makes $q \wedge r$ true to the conclusion that x makes r true.

One way to resolve this problem is to see D|= as holding only in relation to judgments expressed by sentences which are logically simple. We would then define the more general truthmaker relation |=* as follows:

D|=* x |=* p :=

for p logically simple: x |= p

for p of the form $q \vee r$, where q and r are logically simple: x |= q or x |= r,

for p of the form $q \wedge r$, where q and r are logically simple: x |= q and x |= r,

and so forth,

along the lines set forth already in Mulligan *et al.* (1984).

This blocks the move from iii. to iv., and it seems to do so in the right way. That is, it ensures that the A and N components of the truthmaker relation flow in the appropriate way through the pertinent logical parts of each judgment that is in the market for being made true. The argument from i. to v. can be questioned also at other points, however. Above all, it is not clear that the move from iv. to v. is trivial. Thus no such move is sanctioned by the principles for truthmaking set forth in Smith (1999). Here, however, it is the move from i. to ii. upon which we should concentrate, and then we can note that the principle

*If x A p, then x A $p \wedge q$,

which sanctions this move, clearly cannot be of unrestricted validity. For it, too, would imply that the realm of truthbearers is infinitely large. The principle:

A$\wedge$ If x A p and x A q, then x A $p \wedge q$,

on the other hand, is acceptable, but this principle does not allow us to construct a Restall-type argument.

Under what conditions, then, can we go from x A p to x A $p \wedge q$? Given the tenor of our arguments on contexts and judgeability in the above, it will be clear that it is via an appeal to some sort of

relevance that this question must be answered. One might then define what it is for a judgment *q* to be relevant to a judgment *p* in terms of: *p* and *q* share the same context.[14] Since judgments are associated with partitions *in their contexts*, however, this approach would almost certainly bring us back to a principle like A∧.

We consider here one final objection against definitions of the truthmaker relation along the lines of D|=. Suppose that *x* is a malignant necessitator for *p.* This means that *p* is true and *x* is such as to necessitate *p*, but that *x* falls outside the scope of what *p* is about. Our truthmaker definition seems to sanction *x*'s becoming qualified as a truthmaker for *p* merely by dint of *p*'s being judged by a judger who employs a sufficiently weird partition.

Suppose, for example, that Malebranchianism is true, but that this fact is known only to Mary. Whenever Mary makes any empirical judgment, she thereby brings certain extra partitions to bear upon reality in which God's successive necessitating acts of will are set into relief. Let us bite this bullet. God's acts of will are indeed, under this scenario, truthmakers for Mary's judgments (when the latter are true). When Mendeleev made his prediction, somewhere around 1869, to the effect that there exists a hitherto unknown element which he called eka-aluminium, he was using a strange partition of the elements based upon the arrangements of cards in his favorite game of solitaire. The element *gallium* was later found to fit his predictions almost exactly, and his strange partition is nowadays what we call the Periodic Table.

25. A Coda on Realism and the Objectivity of Truth

That scattered portion of the world that is made up of rabbits, that which is made up of rabbit stages, and that which is made up of undetached rabbit parts, are all three just the same scattered portion of the world. The only difference, as Quine sees the matter, 'is in how you slice it.' (1969, p. 32). There are, we can now conclude, two sorts of slicings: the bona fide and the fiat. Both kinds of slicing are represented in our partitions. For even though the cells of the latter are entirely fiat in nature, some of them are coordinated with bona fide demarcations on the side of objects in reality and some of them merely with fiat demarcations which we ourselves have introduced into reality in our various dealings with nature.

14. As Smith (1991) shows, there then results a version of relevance logic that is similar to Parry's (1933) system of analytic implication.

Different philosophers have different views as to which slicings are bona fide and which are fiat. Quine himself holds a view which implies that the metaphysical distinctions between continuants, stages and undetached parts belong in the realm of fiat slicings. Since reference is behaviorally inscrutable as concerns such distinctions, Quine concludes that there is no fact of the matter which they might reflect—no fact of the matter on the side of the objects themselves as these exist before we address them in our language. (It is as if God, in regard to these distinctions, had behaved like the governments of Austria, Germany and Switzerland with regard to their mutual borders around the area of Lake Constance.)

Notice that this is not an epistemological thesis. Quine must hold that even an omniscient being would be in the same predicament as you or me as concerns referential inscrutability. Continuants, parts and stages do not differ from each other in virtue of any corresponding (bona fide) differences on the side of the entities in reality. Rather they differ from each other in the way in which, when asked to count the number of objects in the fruit bowl, you can say either: one orange, or: two orange-halves, or: four orange-quarters, and so on—and you will give the right answer in each case. The distinctions in question are merely the products of our purely fiat partitions of one and the same reality.

But note that Quine is being too hasty when he asserts that there is *no* fact of the matter as concerns the reality to which we are related when using singular referring terms. For it follows from his own doctrine that it is a fact of the matter that this reality is intrinsically undifferentiated as far as metaphysical distinctions and categories are concerned. This is just the other side of the coin from the fact that the corresponding boundaries are entirely fiat in nature.

Quine indeed comes close to a view according to which all boundaries on the side of objects in reality are of the fiat sort. Objects of reference, for him, can comprise any content of some portion of space time, however heterogeneous, disconnected and gerrymandered this may be. For Lewis, on the other hand, whose perspective on these matters we find more congenial:

> Among all the countless things and classes that there are, most are miscellaneous, gerrymandered, ill-demarcated. Only an elite minority are carved at the joints, so that their boundaries are established by objective sameness and difference in nature. Only these elite things and classes are eligible to serve as referents. (Lewis 1984, p.227)

Elite things and classes are in our terms the things and classes captured by those partitions which track bona fide boundaries and relations in reality. It is the job of science to move us in the direction of partitions of this sort. Even when science has completed this job, however, there will still be room

for partitions of the lesser sort, partitions which track boundaries—for example the boundary of Tibbles' tale, or of Mary's kiss, or of the No Smoking Section of your favorite restaurant—which exist only as a result of our acts of fiat.

True judgments of many sorts will thus have truthmakers which are at least partly fiat in nature—and, if our remarks on vagueness in the above are correct, then this will hold not least of our judgments about ourselves. But once again: this is in any case no threat to the objectivity of truth, and nor does it imply that we are completely free in the sorts of partitions we manufacture. This is because, whether our truthmakers are of the bona fide or the fiat type, the relevant portions of the world (including all their causal powers) do of course exist, even before we come along to make our judgments—just as the territory of Denmark did of course exist even before the Danes came along to claim it as their own.

References

Armstrong, D. M. 1997. *A World of States of Affairs*, Cambridge University Press, Cambridge.

Casati, R. and Varzi, A. C. 1999. *Parts and Places. The Structures of Spatial Representation*, Cambridge, MA, and London: MIT Press, Bradford Books.

Crimmins, M. 1992. *Talk about Beliefs*, Cambridge: MIT Press.

Dennett, D. C. 1982. "Beyond Belief," in A. Woodfield (ed.), *Thought and Object*, Oxford: Clarendon Press, 1-95.

Donnellan, K. 1966. "Reference and Definite Descriptions," *Philosophical Review* 75, 281–304.

Evans, G. 1982. *The Varieties of Reference*, J. McDowell, ed, Oxford: Clarendon Press.

Fine, K. 1975. "Vagueness, Truth and Logic," *Synthese*, 30, 265–300.

Heller, M. 2000. "Temporal Overlap is not Coincidence," *The Monist*, 83, 362–380.

Kriegel, U. 1998. "Wide Sense or Narrow Reference?," unpublished paper.

Lewis, D. 1984. "Putnam's Paradox," *Australasian Journal of Philosophy* 62, 221–236. Reprinted in *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press, 1999, 56–77.

Lewis, D. 1991. *Parts of Classes,* Oxford: Blackwell.

Lewis, D. 1993. "Many, But Almost One," in J. Bacon, K. Campbell & L. Reinhardt, *Ontology, Causality and Mind: Essays in Honour of D. M. Armstrong*, Cambridge: Cambridge University Press, 1993. Cited as reprinted in *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press, 1999, 164–182.

Lewis, D. 1996. “Elusive Knowledge,” *Australasian Journal of Philosophy* 74, 549–567.

Lowe, E. J. 1995. “The Problem of the Many and the Vagueness of Constitution,” *Analysis* 55, 179–82.

Mulligan, K. 1997. “How Perception Fixes Reference,” in A. Burri (ed.), *Sprache und Denken / Language and Thought*, Berlin/New York: de Gruyter, 122–138.

Mulligan, K., Simons, P. and Smith, B. 1984. “Truth-Makers,” *Philosophy and Phenomenological Research* 44, 287–321.

Mulligan, K. 1999. "Perception, Predicates and Particulars,” in D. Fisette (ed.), *Consciousness and Intentionality: Models and Modalities of Attribution*, Dordrecht/Boston/London: Kluwer, 163-194.

Omnès, R. 1994. *The Interpretation of Quantum Mechanics*, Princeton: Princeton University Press.

Parry, W. T. (1933), “Ein Axiomensystem für eine neue Art von Implikation (analytische Implikation),” *Ergebnisse eines mathematisehen Kolloquiums,* 4*,* 5–6.

Quine, 1969. “Ontological Relativity,” in *Ontological Relativity and Other Essays,* New York, Columbia University Press, 26–68.

Restall, G. 1996. “Truthmakers, Entailment and Necessity,” *Australasian Journal of Philosophy*, 72, 331-340.

Richard, M. 1990. *Propositional Attitudes*, Cambridge: Cambridge University Press.

Salmon, N. 1986. *Frege’s Puzzle*, Cambridge, MA: MIT Press.

Saul, J. M. 1999. “Substitution and Simple Sentences,” *Analysis* 57, 102–108.

Schiffer, S. 1977. “Naming and Knowing,” *Midwest Studies in Philosophy* 2, 28–41.

Sider, T. 1995. “Three Problems for Richard’s Theory of Belief Ascription,” *Canadian Journal of Philosophy* 25, 487–514.

Simons, P. 1987. *Parts. An Essay in Ontology*, Oxford: Clarendon Press.

Smith, B. 1991. “Relevance, Relatedness and Restricted Set Theory,” in G. Schurz and G. J. W. Dorn, eds., *Advances in Scientific Philosophy. Essays in Honour of Paul Weingartner*, Amsterdam/Atlanta: Rodopi, 45–56

Smith, B. 1995. “On Drawing Lines on a Map,” in A. U. Frank and W. Kuhn (eds.), *Spatial Information Theory. A Theoretical Basis for GIS*, Berlin/Heidelberg/New York: Springer, 475–484.

Smith, B. 1999. “Truthmaker Realism,” *Australasian Journal of Philosophy* 77: 3, 274–291.

Smith, B. 2001. “Fiat Objects”, *Topoi*, 20, 131–148.

Smith, B and Brogaard, B. 2002. "Quantum Mereotopology," *Annals of Mathematics and Artificial Intelligence*, 36, 153–175.

Smith, B. and Varzi, A. C. 2000. "Fiat and Bona Fide Boundaries," *Philosophy and Phenomenological Research* 60: 2, 401–420.

Soames, S. 1988. "Direct Reference, Propositional Attitudes, and Semantic Content," in N. Salmon and S. Soames (eds.), *Propositions and Attitudes*, New York: Oxford University Press, 196–239.

Sundholm, G. 1994 "Existence, Proof and Truth-Making: A Perspective on the Intuitionistic Conception of Truth," *Topoi*, 13, 117–126.

Unger, P. 1980. "The Problem of the Many," *Midwest Studies in Philosophy* 5, 411–67.

van Fraassen, B. C. 1966. "Singular Terms, Truth-Value Gaps, and Free Logic," *Journal of Philosophy* 63, 481–95.

Varzi, A. 2000. "Mereological Commitments," *Dialectica*, 54, 283–305.