

DAVID LEWIS

SCOREKEEPING IN A LANGUAGE GAME*

EXAMPLE 1: PRESUPPOSITION¹

At any stage in a well-run conversation, a certain amount is presupposed. The parties to the conversation take it for granted; or at least they purport to, whether sincerely or just “for the sake of the argument”. Presuppositions can be created or destroyed in the course of a conversation. This change is rule-governed, at least up to a point. The presuppositions at time t' depend, in a way about which at least some general principles can be laid down, on the presuppositions at an earlier time t and on the course of the conversation (and nearby events) between t and t' .

Some things that might be said require suitable presuppositions. They are acceptable if the required presuppositions are present; not otherwise. “The king of France is bald” requires the presupposition that France has one king, and one only; “Even George Lakoff could win” requires the presupposition that George is not a leading candidate; and so on.

We need not ask just what sort of unacceptability results when a required presupposition is lacking. Some say falsehood, some say lack of truth value, some just say that it’s the kind of unacceptability that results when a required presupposition is lacking, and some say it might vary from case to case.

Be that as it may, it’s not as easy as you might think to say something that will be unacceptable for lack of required presuppositions. Say something that requires a missing presupposition, and straightway that presupposition springs into existence, making what you said acceptable after all. (Or at least, that is what happens if your conversational partners tacitly acquiesce – if no one says “But France has *three* kings!” or “Whadda ya mean, ‘*even* George?’”) That is why it is peculiar to say, out of the blue, “All Fred’s children are asleep, and Fred has children.” The first part requires and thereby creates a presupposition that Fred has children; so the second part adds nothing to what is already presupposed when it is said; so the second part has no conversational point. It would not have been peculiar to say instead “Fred has children, and all Fred’s children are asleep.”

Journal of Philosophical Logic 8 (1979) 339–359. 0022–3611/79/0083–0339\$02.10.
Copyright © 1979 by D. Reidel Publishing Co., Dordrecht, Holland, and Boston, U.S.A.

I said that presupposition evolves in a more or less rule-governed way during a conversation. Now we can formulate one important governing rule: call it the *rule of accommodation for presupposition*.

If at time t something is said that requires presupposition P to be acceptable, and if P is not presupposed just before t , then – *ceteris paribus* and within certain limits – presupposition P comes into existence at t .

This rule has not yet been very well stated, nor is it the only rule governing the kinematics of presupposition. But let us bear it in mind nevertheless, and move on to other things.

EXAMPLE 2: PERMISSIBILITY²

For some reason – coercion, deference, common purpose – two people are both willing that one of them should be under the control of the other. (At least within certain limits, in a certain sphere of action, so long as certain conditions prevail.) Call one the *slave*, the other the *master*. The control is exercised verbally, as follows.

At any stage in the enslavement, there is a boundary between some courses of action for the slave that are permissible, and others that are not. The range of permissible conduct may expand or contract. The master shifts the boundary by saying things to the slave. Since the slave does his best to see to it that his course of action is a permissible one, the master can control the slave by controlling what is permissible.

Here is how the master shifts the boundary. From time to time he says to the slave that such-and-such courses of action are impermissible. Any such statement depends for its truth value on the boundary between what is permissible and what isn't. But if the master says that something is impermissible, and if that would be false if the boundary remained stationary, then straightway the boundary moves inward. The permissible range contracts so that what the master says is true after all. Thereby the master makes courses of action impermissible that used to be permissible. But from time to time also the master relents, and says to the slave that such-and-such courses of action are permissible. Or perhaps he says that some of such-and-such courses of action are permissible, but doesn't say just which ones. Then the boundary moves outward. The permissible range expands, if need be

(and if possible), so that what the master says is true. Thereby the master makes courses of action permissible that used to be impermissible.

The truth of the master's statements about permissibility – one aspect of their acceptability – depends on the location of the boundary. The boundary shifts in a rule-governed way. The rule is as follows; call it the *rule of accommodation for permissibility*.

If at time t something is said about permissibility by the master to the slave that requires for its truth the permissibility or impermissibility of certain courses of action, and if just before t the boundary is such as to make the master's statement false, then – *ceteris paribus* and within certain limits – the boundary shifts at t so as to make the master's statement true.

Again, this is not a very satisfactory formulation. For one thing, the limits and qualifications are left unspecified. But more important, the rule as stated does not say exactly how the boundary is to shift.

What if the master says that some of such-and-such courses of actions are permissible, when none of them were permissible before he spoke. By the rule, some of them must straightway become permissible. Some – but which ones? The ones that were closest to permissibility beforehand, perhaps. Well and good, but now we have a new problem. At every state there is not only a boundary between the permissible and the impermissible, but also a relation of comparative near-permissibility between the courses of action on the impermissible side. Not only do we need rules governing the shifting boundary, but also we need rules to govern the changing relation of comparative near-permissibility. Not only must we say how this relation evolves when the master says something about absolute permissibility, but also we must say how it evolves when he says something – as he might – about comparative near-permissibility. He might say, for instance, that the most nearly permissible courses of action in a class A are those in a subclass A' ; or that some courses of action in class B are more nearly permissible than any in class C . Again the rule is a rule of accommodation. The relation of comparative near-permissibility changes, if need be, so that what the master says to the slave is true. But again, to say that is not enough. It does not suffice to determine just what the change is.

Those were Examples 1 and 2. Examples of what? I'll say shortly; but first, a digression.

SCOREKEEPING IN A BASEBALL GAME

At any stage in a well-run baseball game, there is a septuple of numbers $\langle r_v, r_h, h, i, s, b, o \rangle$ which I shall call the *score* of that game at that stage. We recite the score as follows: the visiting team has r_v runs, the home team has r_h runs, it is the h th half (h being 1 or 2) of the i th inning; there are s strikes, b balls, and o outs. (In another terminology, the score is only the initial pair $\langle r_v, r_h \rangle$, but I need a word for the entire septuple.) A possible codification of the rules of baseball would consist of rules of four different sorts.

(1) *Specifications of the kinematics of score.* Initially, the score is $\langle 0, 0, 1, 1, 0, 0, 0 \rangle$. Thereafter, if at time t the score is s , and if between time t and time t' the players behave in manner m , then at time t' the score is s' , where s' is determined in a certain way by s and m .

(2) *Specifications of correct play.* If at time t the score is s , and if between time t and time t' the players behave in manner m , then the players have behaved incorrectly. (Correctness depends on score: what is correct play after two strikes differs from what is correct play after three.) What is not incorrect play according to these rules is correct.

(3) *Directive requiring correct play.* All players are to behave, throughout the game, in such a way that play is correct.

(4) *Directives concerning score.* Players are to strive to make the score evolve in certain directions. Members of the visiting team try to make r_v large and r_h small, members of the home team try to do the opposite.

(We could dispense with roles of sorts (2) and (3) by adding an eighth component to the score which, at any stage of the game, measures the amount of incorrect play up to that stage. Specifications of correct play are then included among the specifications of the kinematics of score, and the directive requiring correct play becomes one of the directives concerning score.)

Rules of sorts (1) and (2) are sometimes called *constitutive rules*. They are said to be akin to definitions, though they do not have the form of definitions. Rules of sorts (3) and (4) are called *regulative rules*. They are akin to the straightforward directives “No smoking!” or “Keep left!”.

We could explain this more fully, as follows. Specifications of sorts (1) and (2) are not themselves definitions of “score” and “correct play”. But they are consequences of reasonable definitions. Further, there is a systematic way to construct the definitions, given the specifications. Suppose we wish to define the *score function*: the function from game-stages to septuples of numbers that gives the score at every stage. The specifications of the kinematics of score, taken together, tell us that the score function evolves in such-and-such way. We may then simply define the score function as that function which evolves in such-and-such way. If the kinematics of score are well specified, then there is one function, and one only, that evolves in the proper way; and if so, then the score function evolves in the proper way if and only if the suggested definition of it is correct. Once we have defined the score function, we have thereby defined the score and all its components at any stage. There are two outs at a certain stage of a game, for instance, if and only if the score function assigns to that game-stage a septuple whose seventh component is the number 2.

Turn next to the specifications of correct play. Taken together, they tell us that correct play occurs at a game-stage if and only if the players’ behavior at that stage bears such-and-such relation to score at that stage. This has the form of an explicit definition of correct play in terms of current behavior. If current score has already been defined in terms of the history of the players’ behavior up to now, in the way just suggested, then we have defined correct play in terms of current and previous behavior.

Once score and correct play are defined in terms of the players’ behavior, then we may eliminate the defined terms in the directive requiring correct play and the directives concerning score. Thanks to the definitions constructed from the constitutive rules, the regulative rules become simply directives to strive to see to it that one’s present behavior bears a certain rather complicated relation to the history of the players’ behavior in previous stages of the game. A player might attempt to conform to such a directive for various reasons: contractual obligation, perhaps, or a conventional understanding with his fellow players based on their common interest in enjoying a proper game.

The rules of baseball could in principle be formulated as straightforward directives concerning behavior, without the aid of definable terms for score and its components. Or they could be formulated as explicit definitions of the score function, the components of score, and correct play, followed by directives in which the newly defined terms appear. It is easy to see why neither of these methods of formulation has found favor. The first method would pack the entire rulebook into each directive; the second would pack the entire rulebook into a single preliminary explicit definition. Understandably averse to very long sentences, we do better to proceed in our more devious way.

There is an alternative analysis – the baseball equivalent of operationalism or legal realism. Instead of appealing to constitutive rules, we might instead claim that the score is, by definition, whatever some scoreboard says it is. Which scoreboard? Various answers are defensible: maybe the visible scoreboard with its arrays of light bulbs, maybe the invisible scoreboard in the head umpire's head, maybe the many scoreboards in many heads to the extent that they agree. No matter. On any such view, the specifications of the kinematics of score have a changed status. No longer are they constitutive rules akin to definitions. Rather, they are empirical generalizations, subject to exceptions, about the ways in which the players' behavior tends to cause changes on the authoritative scoreboard. Under this analysis, it is impossible that this scoreboard fails to give the score. What is possible is that the score is in an abnormal and undesired relation to its causes, for which someone may perhaps be blamed.

I do not care to say which analysis is right for baseball as it is actually played. Perhaps the question has no determinate answer, or perhaps it has different answers for formal and informal baseball. I only wish to distinguish the two alternatives, noting that both are live options.

This ends the digression. Now want to propose some general theses about language – theses that were exemplified by Examples 1 and 2, and that will be exemplified also by several other examples.

CONVERSATIONAL SCORE

With any stage in a well-run conversation, or other process of linguistic interaction, there are associated many things analogous to the components

of a baseball score. I shall therefore speak of them collectively as the *score* of that conversation at that stage. The points of analogy are as follows.

(1) Like the components of a baseball score, the components of a conversational score at a given stage are abstract entities. They may not be numbers, but they are other set-theoretic constructs: sets of presupposed propositions, boundaries between permissible and impermissible courses of action, or the like.

(2) What play is correct depends on the score. Sentences depend for their truth value, or for their acceptability in other respects, on the components of conversational score at the stage of conversation when they are uttered. Not only aspects of acceptability of an uttered sentence may depend on score. So may other semantic properties that play a role in determining aspects of acceptability. For instance, the constituents of an uttered sentence – subsentences, names, predicates, etc. – may depend on the score for their intension or extension.

(3) Score evolves in a more-or-less rule-governed way. There are rules that specify the kinematics of score:

If at time t the conversational score is s , and if between time t and time t' the course of conversation is c , then at time t' the score is s' , where s' is determined in a certain way by s and c .

Or at least:

... then at time t' the score is some member of the class S of possible scores, where S is determined in a certain way by s and c .

(4) The conversationalists may conform to directives, or may simply desire, that they strive to steer certain components of the conversational score in certain directions. Their efforts may be cooperative, as when all participants in a discussion try to increase the amount that all of them willingly presuppose. Or there may be conflict, as when each of two debaters tries to get his opponent to grant him – to join with him in presupposing – parts of his case, and to give away parts of the contrary case.

(5) To the extent that conversational score is determined, given the history of the conversation and the rules that specify its kinematics, these rules can be regarded as constitutive rules akin to definitions. Again,

constitutive rules could be traded in for explicit definitions: the conversational score function could be defined as that function from conversation-stages to n -tuples of suitable entities that evolves in the specified way.

Alternatively, conversational score might be operationally defined in terms of mental scoreboards – some suitable attitudes – of the parties to the conversation. The rules specifying the kinematics of conversational score then become empirical generalizations, subject to exceptions, about the causal dependence of what the scoreboards register on the history of the conversation.

In the case of baseball score, either approach to the definition of score and the status of the rules seems satisfactory. In the case of conversational score, on the other hand, both approaches seem to meet with difficulties. If, as seems likely, the rules specifying the kinematics of conversational score are seriously incomplete, then often there may be many candidates for the score function, different but all evolving in the specified way. But also it seems difficult to say, without risk of circularity, what are the mental representations that comprise the conversationalists' scoreboards.

It may be best to adopt a third approach – a middle way, drawing on both the alternatives previously considered. Conversational score is, by definition, whatever the mental scoreboards say it is; but we refrain from trying to say just what the conversationalists' mental scoreboards are. We assume that some or other mental representations are present that play the role of a scoreboard, in the following sense: what they register depends on the history of the conversation in the way that score should according to the rules. The rules specifying the kinematics of score thereby specify the role of a scoreboard; the scoreboard is whatever best fills this role; and the score is whatever this scoreboard registers. The rules specifying the kinematics of score are to some extent constitutive, but on this third approach they enter only in a roundabout way into the definition of score. It is no harm if they underdetermine the evolution of score, and it is possible that score sometimes evolves in a way that violates the rules.

RULES OF ACCOMMODATION

There is one big difference between baseball score and conversational score. Suppose the batter walks to first base after only three balls. His behavior would be correct play if there were four balls rather than three. That's just

too bad – his behavior does not at all make it the case that there *are* four balls and his behavior *is* correct. Baseball has no rule of accommodation to the effect that if a fourth ball is required to make correct the play that occurs, then that very fact suffices to change the score so that straightway there are four balls.

Language games are different. As I hope my examples will show, conversational score does tend to evolve in such a way as is required in order to make whatever occurs count as correct play. Granted, that is not invariable but only a tendency. Granted also, conversational score changes for other reasons as well. (As when something conspicuous happens at the scene of a conversation, and straightway it is presupposed that it happened.) Still, I suggest that many components of conversational score obey rules of accommodation, and that these rules figure prominently among the rules governing the kinematics of conversational score.

Recall our examples. Example 1: presupposition evolves according to a rule of accommodation specifying that any presuppositions that are required by what is said straightway come into existence, provided that nobody objects. Example 2: permissibility evolves according to a rule of accommodation specifying that the boundaries of the permissible range of conduct shift to make true whatever is said about them, provided that what is said is said by the master to the slave, and provided that there does exist some shift that would make what he says true. Here is a general scheme for rules of accommodation for conversational score.

If at time t something is said that requires component s_n of conversational score to have a value in the range r if what is said is to be true, or otherwise acceptable; and if s_n does not have a value in the range r just before t ; and if such-and-such further conditions hold; then at t the score-component s_n takes some value in the range r .

Once we have this scheme in mind, I think we will find many instances of it. In the rest of this paper I shall consider some further examples. I shall have little that is new to say about the individual examples. My interest is in the common pattern that they exhibit.

EXAMPLE 3: DEFINITE DESCRIPTIONS³

It is not true that a definite description “the *F*” denotes *x* if and only if *x* is the one and only *F* in existence. Neither is it true that “the *F*” denotes *x* if and only if *x* is the one and only *F* in some contextually determined domain of discourse. For consider this sentence: “The pig is grunting, but the pig with floppy ears is not grunting” (Lewis). And this: “The dog got in a fight with another dog” (McCawley). They could be true. But for them to be true, “the pig” or “the dog” must denote one of two pigs or dogs, both of which belong to the domain of discourse.

The proper treatment of descriptions must be more like this: “the *F*” denotes *x* if and only if *x* is the most salient *F* in the domain of discourse, according to some contextually determined salience ranking. The first of our two sentences means that the most salient pig is grunting but the most salient pig with floppy ears is not. The second means that the most salient dog got in a fight with some less salient dog.

(I shall pass over some complications. Never mind what happens if two *F*’s are tied for maximum salience, or if no *F* is at all salient. More important, I shall ignore the possibility that something might be highly salient in one of its guises, but less salient in another. Possibly we really need to appeal to a salience ranking not of individuals but rather of individuals-in-guises – that is, of individual concepts.)

There are various ways for something to gain salience. Some have to do with the course of conversation, others do not. Imagine yourself with me as I write these words. In the room is a cat, Bruce, who has been making himself very salient by dashing madly about. He is the only cat in the room, or in sight, or in earshot. I start to speak to you:

The cat is in the carton. The cat will never meet our other cat, because our other cat lives in New Zealand. Our New Zealand cat lives with the Cresswells. And there he’ll stay, because Miriam would be sad if the cat went away.

At first, “the cat” denotes Bruce, he being the most salient cat for reasons having nothing to do with the course of conversation. If I want to talk about Albert, our New Zealand cat, I have to say “our other cat” or “our New Zealand cat”. But as I talk more and more about Albert, and not any more about Bruce, I raise Albert’s salience by conversational means. Finally,

in the last sentence of my monologue, I am in a position to say “the cat” and thereby denote not Bruce but rather the newly-most-salient cat Albert.

The ranking of comparative salience, I take it, is another component of conversational score. Denotation of definite descriptions is score-dependent. Hence so is the truth of sentences containing such descriptions, which is one aspect of the acceptability of those sentences. Other aspects of acceptability in turn are score-dependent: non-triviality, for one, and possibility of warranted assertion, for another.

One rule, among others, that governs the kinematics of salience is a rule of accommodation. Suppose my monologue has left Albert more salient than Bruce; but the next thing I say is “The cat is going to pounce on you!” If Albert remains most salient and “the cat” denotes the most salient cat, then what I say is patently false: Albert cannot pounce all the way from New Zealand to Princeton. What I have said requires for its acceptability that “the cat” denote Bruce, and hence that Bruce be once again more salient than Albert. If what I say requires that, then straightway it is so. By saying what I did, I have made Bruce more salient than Albert. If next I say “The cat prefers moist food”, that is true if Bruce prefers moist food, even if Albert doesn’t.

The same thing would have happened if instead I had said “The cat is out of the carton” or “The cat has gone upstairs”. Again what I say is unacceptable unless the salience ranking shifts so that Bruce rises above Albert, and hence so that ‘the cat’ again denotes Bruce. The difference is in the type of unacceptability that would ensue without the shift. It is trivially true, hence not worth saying, that Albert is out of the carton. (“The carton” denotes the same carton as before; nothing has been done to raise the salience of any carton in New Zealand.) It may be true or it may be false that Albert has gone upstairs in the Cresswells’ house in New Zealand. But I have no way of knowing, so I have no business saying that he has.

We can formulate a *rule of accommodation for comparative salience* more or less as follows. It is best to speak simply of unacceptability, since it may well be that the three sorts of unacceptability I have mentioned are not the only sorts that can give rise to a shift in salience.

If at time t something is said that requires, if it is to be acceptable, that x be more salient than y ; and if, just before t , x is no more salient than y ; then – *ceteris paribus* and within certain limits – at t , x becomes more salient than y .

Although a rule of accommodation, such as this one, states that shifts of score take place when they are needed to preserve acceptability, we may note that the preservation is imperfect. It is not good conversational practice to rely too heavily on rules of accommodation. The monologue just considered illustrates this. Because “the cat” denotes first Bruce, then Albert, then Bruce again, what I say is to some extent confusing and hard to follow. But even if my monologue is not perfectly acceptable, its flaws are much less serious than the flaws that are averted by shifts of salience in accordance with our rule of accommodation. Confusing shifts of salience and reference are not as bad as falsity, trivial truth, or unwarranted assertion.

(It is worth mentioning another way to shift comparative salience by conversational means. I may say “A cat is on the lawn” under circumstances in which it is apparent to all parties to the conversation that there is some one particular cat that is responsible for the truth of what I say, and for my saying it. Perhaps I am looking out the window, and you rightly presume that I said what I did because I saw a cat; and further (since I spoke in the singular) that I saw only one. What I said was an existential quantification; hence, strictly speaking, it involves no reference to any particular cat. Nevertheless it raises the salience of the cat that made me say it. Hence this newly-most-salient cat may be denoted by brief definite descriptions, or by pronouns, in subsequent dialogue: “No, it’s on the sidewalk.” “Has Bruce noticed the cat?” As illustrated, this may happen even if the speaker contradicts my initial existential statement. Thus although indefinite descriptions – that is, idioms of existential quantification – are not themselves referring expressions, they may raise the salience of particular individuals in such a way as to pave the way for referring expressions that follow.)

EXAMPLE 4: COMING AND GOING⁴

Coming is a movement toward a point of reference. Going is movement away from it. Sometimes the point of reference is fixed by the location of speaker and hearer, at the time of conversation or the time under discussion. But sometimes not. In third-person narrative, whether fact or fiction, the chosen point of reference may have nothing to do with the speaker’s or the hearer’s location.

One way to fix the point of reference at the beginning of a narrative, or to shift it later, is by means of a sentence that describes the direction of

some movement both with respect to the point of reference and in some other way. "The beggars are coming to town" requires for its acceptability, and perhaps even for its truth, that the point of reference be in town. Else the beggars' townward movement is not properly called "coming". This sentence can be used to fix or to shift the point of reference. When it is said, straightway the point of reference is in town where it is required to be. Thereafter, unless something is done to shift it elsewhere, coming is movement toward town and going is movement away. If later we are told that when the soldiers came the beggars went, we know who ended up in town and who did not.

Thus the point of reference in narrative is a component of conversational score, governed by a rule of accommodation. Note that the rule must provide for two sorts of changes. The point of reference may simply go from one place to another, as is required by the following text:

When the beggars came to town, the rich folk went to the shore. But soon the beggars came after them, so they went home.

But also the point of reference is usually not fully determinate in its location. It may become more or less determinate, as is required by the following:

After the beggars came to town, they held a meeting. All of them came to the square. Afterwards they went to another part of town.

The first sentence puts the point of reference in town, but not in any determinate part of town. The second sentence increases its determinacy by putting it in the square. The initial fixing of the point of reference is likewise an increase in determinacy – the point of reference starts out completely indeterminate and becomes at least somewhat more definitely located.

EXAMPLE 5: VAGUENESS⁵

If Fred is a borderline case of baldness, the sentence "Fred is bald" may have no determinate truth value. Whether it is true depends on where you draw the line. Relative to some perfectly reasonable ways of drawing a precise boundary between bald and not-bald, the sentence is true. Relative

to other delineations, no less reasonable, it is false. Nothing in our use of language makes one of these delineations right and all the others wrong. We cannot pick a delineation once and for all (not if we are interested in ordinary language), but must consider the entire range of reasonable delineations.

If a sentence is true over the entire range, true no matter how we draw the line, surely we are entitled to treat it simply as true. But also we treat a sentence more or less as if it is simply true, if it is true over a large enough part of the range of delineations of its vagueness. (For short: if it is *true enough*.) If a sentence is true enough (according to our beliefs) we are willing to assert it, assent to it without qualification, file it away among our stocks of beliefs, and so forth. Mostly we do not get into any trouble this way. (But sometimes we do, as witness the paradoxes that arise because truth-preserving reasoning does not always preserve the property of being true enough.)

When is a sentence true enough? Which are the "large enough" parts of the range of delineations of its vagueness? This is itself a vague matter. More important for our present purposes, it is something that depends on context. What is true enough on one occasion is not true enough on another. The standards of precision in force are different from one conversation to another, and may change in the course of a single conversation. Austin's "France is hexagonal" is a good example of a sentence that is true enough for many contexts, but not true enough for many others. Under low standards of precision it is acceptable. Raise the standards and it loses its acceptability.

Taking standards of precision as a component of conversational score, we once more find a rule of accommodation at work. One way to change the standards is to say something that would be unacceptable if the standards remained unchanged. If you say "Italy is boot-shaped" and get away with it, low standards are required and the standards fall if need be; thereafter "France is hexagonal" is true enough. But if you deny that Italy is boot-shaped, pointing out the differences, what you have said requires high standards under which "France is hexagonal" is far from true enough.

I take it that the rule of accommodation can go both ways. But for some reason raising of standards goes more smoothly than lowering. If the standards have been high, and something is said that is true enough only under lowered standards, and nobody objects, then indeed the standards are shifted down. But what is said, although true enough under the lowered

standards, may still seem imperfectly acceptable. Raising of standards, on the other hand, manages to seem commendable even when we know that it interferes with our conversational purpose. Because of this asymmetry, a player of language games who is so inclined may get away with it if he tries to raise the standards of precision as high as possible – so high, perhaps, that no material object whatever is hexagonal.

Peter Unger has argued that hardly anything is flat. Take something you claim is flat; he will find something else and get you to agree that it is even flatter. You think the pavement is flat – but how can you deny that your desk is flatter? But “flat” is an *absolute term*: it is inconsistent to say that something is flatter than something that is flat. Having agreed that your desk is flatter than the pavement, you must concede that the pavement is not flat after all. Perhaps you now claim that your desk is flat; but doubtless Unger can think of something that you will agree is even flatter than your desk. And so it goes.

Some might dispute Unger’s premise that “flat” is an absolute term; but on that score it seems to me that Unger is right. What he says is inconsistent does indeed sound that way. I take this to mean that on no delineation of the correlative vagueness of “flatter” and “flat” is it true that something is flatter than something that is flat.

The right response to Unger, I suggest, is that he is changing the score on you. When he says that the desk is flatter than the pavement, what he says is acceptable only under raised standards of precision. Under the original standards the bumps on the pavement were too small to be relevant either to the question whether the pavement is flat or to the question whether the pavement is flatter than the desk. Since what he says requires raised standards, the standards accommodately rise. Then it is no longer true enough that the pavement is flat. That does not alter the fact that it *was* true enough *in its original context*. “The desk is flatter than the pavement” said under raised standards does not contradict “The pavement is flat” said under unraised standards, any more than “It is morning” said in the morning contradicts “It is afternoon” said in the afternoon. Nor has Unger shown in any way that the new context is more legitimate than the old one. He can indeed create an unusual context in which hardly anything can acceptably be called “flat”, but he has not thereby cast any discredit on the more usual contexts in which lower standards of precision are in force.

In parallel fashion Unger observes, I think correctly, that “certain” is an

absolute term; from this he argues that hardly ever is anyone certain of anything. A parallel response is in order. Indeed the rule of accommodation permits Unger to create a context in which all that he says is true, but that does not show that there is anything whatever wrong with the claims to certainty that we make in more ordinary contexts. It is no fault in a context that we can move out of it.

EXAMPLE 6: RELATIVE MODALITY⁶

The “can” and “must” of ordinary language do not often express absolute (“logical” or “metaphysical”) possibility. Usually they express various relative modalities. Not all the possibilities there are enter into consideration. If we ignore those possibilities that violate laws of nature, we get the physical modalities; if we ignore those that are known not to obtain, we get the epistemic modalities; if we ignore those that ought not to obtain – doubtless including actuality – we get the deontic modalities; and so on. That suggests that “can” and “must” are ambiguous. But on that hypothesis, as Kratzer has convincingly argued, the alleged senses are altogether too numerous. We do better to think of our modal verbs as unambiguous but relative. Sometimes the relativity is made explicit. Modifying phrases like “in view of what is known” or “in view of what custom requires” may be present to indicate just which possibilities should be ignored.

But sometimes no such phrase is present. Then context must be our guide. The boundary between the relevant possibilities and the ignored ones (formally, the accessibility relation) is a component of conversational score, which enters into the truth conditions of sentences with “can” or “must” or other modal verbs. It may change in the course of conversation. A modifying phrase “in view of such-and-such” does not only affect the sentence in which it appears, but also remains in force until further notice to govern the interpretation of modal verbs in subsequent sentences.

This boundary may also shift in accordance with a rule of accommodation. Suppose I am talking with some elected official about the ways he might deal with an embarrassment. So far, we have been ignoring those possibilities that would be political suicide for him. He says: “You see, I must either destroy the evidence or else claim that I did it to stop Communism. What else can I do?” I rudely reply: “There is one other possibility – you can put the public interest first for once!” That would be

false if the boundary between relevant and ignored possibilities remained stationary. But it is not false in its context, for hitherto ignored possibilities come into consideration and make it true. And the boundary, once shifted outward, stays shifted. If he protests “I can’t do that”, he is mistaken.

Take another example. The commonsensical epistemologist says: “I *know* the cat is in the carton – there he is before my eyes – I just *can’t* be wrong about that!” The sceptic replies: “You might be the victim of a deceiving demon”. Thereby he brings into consideration possibilities hitherto ignored, else what he says would be false. The boundary shifts outward so that what he says is true. Once the boundary is shifted, the commonsensical epistemologist must concede defeat. And yet he was not in any way wrong when he laid claim to infallible knowledge. What he said was true with respect to the score as it then was.

We get the impression that the sceptic, or the rude critic of the elected official, has the last word. Again this is because the rule of accommodation is not fully reversible. For some reason, I know not what, the boundary readily shifts outward if what is said requires it, but does not so readily shift inward if what is said requires that. Because of this asymmetry, we may think that what is true with respect to the outward-shifted boundary must be somehow more true than what is true with respect to the original boundary. I see no reason to respect this impression. Let us hope, by all means, that the advance toward truth is irreversible. That is no reason to think that just any change that resists reversal is an advance toward truth.

EXAMPLE 7: PERFORMATIVES ⁷

Suppose we are unpersuaded by Austin’s contention that explicit performatives have no truth value. Suppose also that we wish to respect the seeming parallelism of form between a performative like “I hereby name this ship the *Generalissimo Stalin*” and such non-performative statements as “Fred thereby named that ship the *President Nixon*”. Then we shall find it natural to treat the performative, like the non-performative, as a sentence with truth conditions. It is true, on a given occasion of its utterance, if and only if the speaker brings it about, by means of that very utterance, that the indicated ship begins to bear the name “Generalissimo Stalin”. If the circumstances are felicitous, then the speaker does indeed bring it about, by means of his utterance, that the ship begins to bear the name. The performative sentence

is therefore true on any occasion of its felicitous utterance. In Lemmon's phrase, it is a sentence verifiable by its (felicitous) use.

When the ship gets its name and the performative is verified by its use, what happens may be described as a change in conversational score governed by a rule of accommodation. The relevant component of score is the relation that pairs ships with their names. The rule of accommodation is roughly as follows.

If at time t something is said that requires for its truth that ship s bear name n ; and if s does not bear n just before t ; and if the form and circumstances of what is said satisfy certain conditions of felicity; then s begins at t to bear n .

Our performative sentence does indeed require for its truth that the indicated ship bear the name "Generalissimo Stalin" at the time of utterance. Therefore, when the sentence is felicitously uttered, straightway the ship bears the name.

The sentence has other necessary conditions of truth: the ship must not have borne the name beforehand, the speaker must bring it about that the ship begins to bear the name, and he must bring it about by uttering the sentence. On any felicitous occasion of utterance, these further conditions take care of themselves. Our rule of accommodation is enough to explain why the sentence is verified by its felicitous use, despite the fact that the rule deals only with part of what it takes to make the sentence true.

A similar treatment could be given of many other performatives. In some cases the proposal may seem surprising. "With this ring I thee wed" is verified by its felicitous use, since the marriage relation is a component of conversational score governed by a rule of accommodation. Is marriage then a *linguistic* phenomenon? Of course not, but that was not implied. The lesson of performatives, on any theory, is that use of language blends into other social practices. We should not assume that a change of conversational score has its impact only within, or by way of, the realm of language. Indeed, we have already seen another counterexample: the case of permissibility, considered as Example 2.

EXAMPLE 8: PLANNING

Suppose that you and I are making a plan – let us say, a plan to steal some plutonium from a reprocessing plant and make a bomb of it. As we talk,

our plan evolves. Mostly it grows more and more complete. Sometimes, however, parts that had been definite are revised, or at least opened for reconsideration.

Much as some things said in ordinary conversation require suitable pre-suppositions, so some things we say in the course of our planning require, for their acceptability, that the plan contain suitable provisions. If I say "Then you drive the getaway car up to the side gate", that is acceptable only if the plan includes provision for a getaway car. That might or might not have been part of the plan already. If not, it may become part of the plan just because it is required by what I said. (As usual the process is defeasible. You can keep the getaway car out of the plan, for the time being at least, by saying "Wouldn't we do better with mopeds?") The plan is a component of conversational score. The rules governing its evolution parallel the rules governing the kinematics of presupposition, and they include a rule of accommodation.

So good is the parallel between plan and presupposition that we might well ask if our plan simply *is* part of what we presuppose. Call it that if you like, but there is a distinction to be made. We might take for granted, or purport to take for granted, that our plan will be carried out. Then we would both plan and presuppose that we are going to steal the plutonium. But we might not. We might be making our plan not in order to carry it out, but rather in order to show that the plan needs better security. Then plan and presupposition might well conflict. We plan to steal the plutonium, all the while presupposing that we will not. And indeed our planning may be interspersed with commentary that requires presuppositions contradicting the plan. "Then I'll shoot the guard (I'm glad I won't really do that) while you smash the floodlights." Unless we distinguish plan from presupposition (or distinguish two levels of presupposition) we must think of presuppositions as constantly disappearing and reappearing throughout such a conversation.

The distinction between plan and presupposition is not the distinction between what we purport to take for granted and what we really do. While planning that we will steal the plutonium and presupposing that we will not, we might take for granted neither that we will nor that we won't. Each of us might secretly hope to recruit the other to the terrorist cause and carry out the plan after all.

One and the same sentence may require, and if need be create, both

provisions of the plan and presuppositions. "Then you drive the getaway car up to the side gate" requires both a getaway car and a side gate. The car is planned for. The gate is more likely presupposed.

Princeton University

NOTES

* I am doubly grateful to Robert Stalnaker: first, for his treatment of presupposition, here summarized as Example 1, which I have taken as the prototype for parallel treatments of other topics; and second, for valuable comments on a previous version of this paper. I am also much indebted to Stephen Isard, who discusses many of the phenomena that I consider here in his 'Changing the Context' in Edward L. Keenan, ed., *Formal Semantics of Natural Language* (Cambridge University Press, 1974). Proposals along somewhat the same lines as mine are to be found in Thomas T. Ballmer, 'Einführung und Kontrolle von Diskurswelten', in Dieter Wunderlich, ed., *Linguistische Pragmatik* (Athenäum-Verlag, 1972), and Ballmer, *Logical Grammar: with Special Consideration of Topics in Context Change* (North-Holland, 1978).

An early version of this paper was presented to the Vacation School in Logic at Victoria University of Wellington in August 1976; I thank the New Zealand–United States Educational Foundation for research support on that occasion. The paper also was presented at a workshop on pragmatics and conditionals at the University of Western Ontario in May 1978, and at a colloquium on semantics at Konstanz University in September 1978.

¹ This treatment of presupposition is taken from two papers of Robert Stalnaker: 'Presuppositions', *Journal of Philosophical Logic* 2 (1973), 447–457, and 'Pragmatic Presuppositions', in Milton K. Munitz and Peter K. Unger, eds., *Semantics and Philosophy* (New York University Press, 1974).

² This treatment of permissibility is discussed more fully in my paper 'A Problem about Permission', in Esa Saarinen *et al.*, eds., *Essays in Honour of Jaakko Hintikka* (Reidel).

³ Definite descriptions governed by salience are discussed in my *Counterfactuals* (Blackwell, 1973), pp. 111–117; and in James McCawley, 'Presupposition and Discourse Structure', in David Dinneen and Choon-Kyu Oh, eds., *Syntax and Semantics*, Vol. 11 (Academic Press, 1979). A similar treatment of demonstratives is found in Isard, *op. cit.*

Manfred Pinkal, 'How to Refer with Vague Descriptions' (presented at the Konstanz colloquium on semantics, September 1978) notes a further complication: if some highly salient things are borderline cases of *F*-hood, degree of *F*-hood and salience may trade off.

Indefinite descriptions that pave the way for referring expressions are discussed in Charles Chastain, 'Reference and Context', *Minnesota Studies in the Philosophy of Science* 7 (1975), 194–269, and in Saul Kripke, 'Speaker's Reference and Semantic Reference', *Midwest Studies in Philosophy* 2 (1977), 255–276.

⁴ See Charles Fillmore, 'How to Know Whether You're Coming or Going', in Karl Hyldgaard-Jensen, ed., *Linguistik 1971* (Athenäum-Verlag, 1972), and 'Pragmatics and

the Description of Discourse', in Siegfried J. Schmidt, ed., *Pragmatik/Pragmatics II* (Wilhelm Fink Verlag, 1976).

⁵ See the treatment of vagueness in my 'General Semantics', *Synthese* 22 (1970), 18–67. For arguments that hardly anything is flat or certain, see Peter Unger, *Ignorance* (Oxford University Press, 1975), pp. 65–68. For another example of accommodating shifts in resolution of vagueness, see the discussion of back-tracking counterfactuals in my 'Counterfactual Dependence and Time's Arrow', *Noûs* 13 (1979).

⁶ See Angelika Kratzer, 'What "Must" and "Can" Must and Can Mean', *Linguistics and Philosophy* 1 (1977), 337–355. The accessibility semantics considered here is equivalent to a slightly restricted form of Kratzer's semantics for relative modality.

Knowledge and irrelevant possibilities of error are discussed in Alvin I. Goldman, 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73 (1976), 771–791.

⁷ See J. L. Austin, 'Performative Utterances', in his *Philosophical Papers* (Oxford University Press, 1961) for the original discussion of performatives. For treatments along the lines here preferred, see E. J. Lemmon, 'On Sentences Verifiable by Their Use', *Analysis* 22 (1962), 86–89; Ingemar Hedenius, 'Performatives', *Theoria* 29 (1963), 1–22; and Lennart Aqvist, *Performatives and Verifiability by the Use of Language* (Filosofiska Studier, Uppsala University, 1972). Isard (*op. cit.*) suggests as I do that performative utterances are akin to other utterances that 'change the context'.