

## 3

## Showing and Meaning

I have begun to characterize self-expression by situating it within the larger framework of signaling, which can occur in the absence of any intentions to send a signal or consciousness of doing so. Self-expression does require what is expressed to be a cognitive, affective, or experiential state, but it is not true that all such states are expressed intentionally. Nevertheless, one intuitively familiar form of self-expression occurs in cases that philosophers refer to as ‘speaker meaning’ (née ‘non-natural meaning’), and such cases do involve intention. No current theory of speaker meaning is, however, accurate, so in this chapter I’ll offer a new one. In particular, I lay the groundwork for an account of the relation between self-expression and speaker meaning, paying particular attention to the extent to which and the ways in which intentions are involved. I first (3.1) distinguish among three species of showing. Then (3.2) on that basis I further develop the signaling model of communication outlined in Chapter 1, explaining three ways in which signals of the sort that concern us in this book can show what they signal. I next (3.3) argue, contrary to widely held opinion, that speaker meaning does not require intentions to produce effects on an audience, much less intentions to produce effects by means of recognition of those intentions. We shall see instead that it’s necessary and sufficient for speaker meaning that one overtly show something, or overtly show that something is so, or overtly show one’s commitment to a content in a certain way. In 3.4 I develop this notion of overtness on which that characterization of speaker meaning depends. In 3.5 I defend a view of many familiar speech acts as handicaps in the technical sense of that term. This will in turn motivate a generalization of the notion of speaker meaning achieved up to that point. I also consider, in 3.6, some alternative conceptions of speaker meaning, and show why they are inadequate. These ideas will lay the groundwork for Chapter 4, in which certain forms of self-expression emerge as species of speaker meaning.

## 3.1. Three ways of showing

Our inquiry will make heavy use of the notion of showing, which comes in at least three forms. First of all, I might show my courage by acting bravely. My brave behavior is good evidence of my courage. Or I might, by means of extensive calculations, show that there is a black hole in the center of the Milky Way. In these cases I don’t make what I show perceptible; I certainly couldn’t make the black hole perceptible, and it is not clear what it could mean to perceive courage. Rather, in these cases I provide compelling, though not necessarily conclusive, evidence for a conclusion that could be grasped even by someone with no capacity for vision or other sensation. A grammatical tag for this category is *showing-that*. Because my brave behavior (calculations, etc.) is good evidence of my courage (the existence of the black hole, etc.), an appropriately situated thinker aware of that evidence is in a position to know of my courage (of the black hole, etc.). Showing—that thus enables propositional knowledge.

Just as a horse can be led to water without being made to imbibe, so too enabling propositional knowledge is not the same thing as guaranteeing that all interested parties will come to know what has been shown. I might prove a theorem for my class, showing them that, say,  $a^2 + b^2 = c^2$ , where  $a$ ,  $b$ , and  $c$  are three sides of a right triangle, without everyone’s cottoning on. A week later one of my students might finally get it, remarking, “Oh! Last week he showed us that  $a^2 + b^2 = c^2$ .” I just didn’t follow the reasoning at the time.” Making knowledge available doesn’t guarantee its transmission. On the other hand, my running through the proof doesn’t make knowledge available to everyone who watches. If the aforementioned horse happens to have witnessed the proof, I won’t have shown him the truth of Pythagoras’ theorem; the same goes for my newborn daughter. In order to show you that something is so, you need, at the very least, the conceptual resources required to believe the proposition in question. I can therefore show my student that  $a^2 + b^2 = c^2$  without, in so doing, showing my horse or daughter that this is so even if they’re watching.

A second form of showing makes a thing perceptible. I show my bruise, and thereby enable others to see that bruise. Although it is most natural to speak of showing in visual terms, showing is not limited to vision: one can show someone a rough texture (you’d need to feel the texture)

or a coyote's howl (you'd have to hear it). Just as what I show you in the propositional-knowledge enabling case requires appropriate conceptual resources on your part, what I show you in this perceptual case depends on your perceptual capacities and your position in the environment. If you had electroreception like a hammerhead shark, I could show you the electrical activity in the body of a fish hiding under the sand. In that case you'd not only perceive the fish, you'd "electroreceive" it. Likewise, even if there are mice in the field, I don't show you them from an airplane passing two hundred yards above the field. On the other hand if you had the visual acuity of a hawk, I might well do so. Let us put this perceptual-knowledge enabling form of showing under the rubric of *showing- $\alpha$* , where ' $\alpha$ ' is a singular term referring to a perceptible object or affair.

Finally, I might also *show how* something looks, feels, sounds, and so on. Apply friction to a scratch-and-sniff picture of a skunk. You won't thereby smell any skunk, but if your nose is functioning properly, you will learn how skunks smell. By accurately painting Mary's profile you will show how Mary looks in profile—what she looks like from that angle—thereby enabling me to know how Mary's profile looks. If I then acquire the knowledge that has been made available, I can later manifest it by reliably discriminating the Mary-like profiles from the rest. Similarly, the trepidation in my voice might enable you to know how my anxiety feels if you are sufficiently empathetic. If you are sufficiently empathetic, then hearing my voice may enable you to imagine feeling my trepidation. If you can do that, then you know how I feel. Showing-how can provide qualitative knowledge for those with appropriate sensory capacities. It can also enable empathy for those with the capacity for empathy. The above three forms of showing—showing-that, showing- $\alpha$ , and showing-how—enable propositional knowledge, perceptual knowledge, and either experiential knowledge or empathy, respectively.

Very often a single experience will involve all three forms of showing. My location and sensory capacities will enable me to see the cloud; doing so might show me that a storm is brewing, as well as what a cumulonimbus looks like. On the other hand, such phenomena as blindsight suggest that an event might show  $\alpha$  without showing me how  $\alpha$  looks: analogous things may be possible for other senses. We have, further, already mentioned that one can attain propositional knowledge of a state of affairs without perceiving it. For these reasons, clarity is best served by keeping our three

forms of showing distinct even though in typical cases they arrive bundled together.

One thread that unites the above three forms of showing is knowledge: Evidence enables those who are shown the things mentioned above, and who are in the right circumstances (being empathetic, being in the right perceptual location, possessed of the right conceptual resources or background knowledge, etc.) to know some fact, some object of perception, or how some emotion, mood, or experience feels. Showing is thus a stronger relation than indication, in two ways. First, showing, unlike indication, is a "success" notion: One can only show facts (showing that), or real things (showing  $\alpha$ ), or how something appears or feels (showing how), whereas one can indicate that something is so when it is not, or indicate an object that is not, or indicate how something appears or feels that does not appear or feel that way. Another thread unifying the three types of showing is that each of the forms of knowledge it enables can be made available by design. That is to say that each of the three forms of showing can be the content of a signal. Just as the height of the tiger's scratch marks shows its size, so too the intensity of my cry, when I'm neither faking nor exaggerating, shows the extent of my terror. A sincere speech act of mine might show my belief; the cry might also show how my terror feels. I develop these points in the next section.

### 3.2. Showing what's within, part i

Recalling terminology of Chapter 1, a *signal* is any feature of an entity that conveys information (including misinformation) and that was designed for its ability to convey that information. That information might pertain to how things were, are, will be, or ought to be. The design in question might be due to the work of intelligent agents, or be the product of evolution by either artificial or natural selection. When a signal succeeds in conveying the information for which it was designed, communication takes place. This is the *signaling model of communication*. I argued in Chapter 1 that this signaling model can accommodate the use of codes as suggested by the code model, and inference as suggested by the inferential model, as special cases of a more general pattern. In addition, I have been developing a position

that both articulates and clarifies what in Chapter 1 we called the extended senses model of communication.

For any organism and characteristic ecology, there will be a cost involved in sending a signal. The warningly colored tree frog expends some resources in producing bright colors. However, those resources appear to be negligible in relation to the resources required to produce any other skin color. For that reason, the tree frog only pays what is called an *efficacy cost* in producing its warning signal (Maynard Smith and Harper 2004, p. 15). Mimics, who wear bright skin color without being noxious, pay the same efficacy cost to produce the same signal. It should be plausible, however, that if too many mimics are found in a given environment, predators will start to ignore the warning signal and start eating all brightly colored frogs whether they are noxious or not. The result is that, when mimicry is possible, and organisms can gain from mimicry without incurring any substantial cost, signals are liable to lose their credibility.

This conclusion might make one wonder if any signaling system will be unstable, always liable to overthrow by mutations or other tricksters that produce deceptive signals. The answer is that some signaling systems seem to have developed that are resilient against such threats. We know of three conditions under which this stability emerges, where the third is a special case of the second. One occurs when the organisms in question share the same preference ordering. We have evidence that the African honeyguide bird (*Indicator indicator*) has been helping humans to find beehives for at least 20,000 years. These birds gain the attention of humans with a distinctive chirp, and the humans follow them to the beehive. The humans then extract the honey, and the honeyguides eat the larvae and wax from the nest (Isack and Reyer 1989). Neither party to this interaction would have an incentive to signal “dishonestly”, and that’s good reason to believe that the honeyguides’ signals are and will remain reliable.

Another source of signaling stability is found when the signals in question are indices—signals that can only be faked with great difficulty due to limitations on the organism. We have already seen this in the example of the tiger’s scratch-marks. Another case is funnel-web spiders, *Agelenopsis aperta*, who find themselves in contests over webs. Two spiders will vibrate on a disputed web. Reichert (1978, 1984) found that if two contesting spiders differ in weight by 10 per cent or more, the lighter spider retreats 90 per cent of the time rather than fighting. A losing spider can, in addition, be

made into a winner by placing a weight on its back. This strongly suggests that vibrating on a web is a spider’s signal of its size. Further, in the absence of scientists placing weights on their backs, funnel-web spiders can’t fake these signals. The “vibrating game” thus exploits an index of spider size. In so doing, its players not only signal their size; they also show it. It is, I suggest, precisely by showing their size that spiders have happened upon a stable signaling system.

A special case of an index is a handicap—which is a signal that is more costly than is required just for sending a signal with that content. Intuitively, when a signal is a handicap, only those who are “honest” can afford to pay the price required to send them. As we observed in Chapter 1, the ostentatious feathers of the male peacock don’t just signal his viability—they show it because in spite of carrying such a handicap, he survives: the fact that he has survived up to now with all that baggage shows his viability. Another example of a handicap is found in the male stalk-eyed fly, *Cyrtodiopsis dalmanni*. In choosing mates, females prefer males with long eye-stalks (Wilkinson and Reillo 1994; David, *et al* 1998). This is in spite of the fact that these appendages make the male slower and easier to spot for predators, and are costly to produce. Observing that these appendages provide evidence of the male’s viability—since he’d have to be unusually fit to survive with them—Maynard Smith and Harper (2004, pp. 33–4) hypothesize that the eye-stalks are also signals of the male’s viability. A signal, such as the peacock’s tail or the stalk-eyed fly’s stalks, whose cost goes beyond its efficacy cost carries a *strategic cost* (Maynard Smith and Harper 2004, p. 14). My suggestion now is that when a signal carries a sufficiently high strategic cost, it not only signals some property of the organism; it also shows that property. Further, it is its ability to show this property that vouchsafes the stability of this signaling strategy.

Signaling systems, then, can become reliable by virtue of the confluence of interests of the communicators, or by exploiting indices, of which a special case is a handicap. Indices, I have suggested, gain their power by their ability to show what they index, and handicaps show what they do in a peculiar way. (Figure 3.1 depicts relations among such concepts as cue, signal, showing, index, and handicap.) Because of the intimate connection between reliable signaling and knowledge, we are now in a position to see that reliably signaling a state of affairs enables others who are appropriately




Figure 3.1: Relations among cues, signals, showing, indices, and handicaps

situated to know of that state of affairs or an aspect of it—either by knowing that it obtains, by knowing how it feels, or simply by perceiving it.

We've now arrived at a picture of reliable signaling as, *inter alia*, making knowledge available to others. Recalling our characterization of self-expression from Chapter 2 as signaling and showing some aspect of one's point of view, this implies that in expressing ourselves we make knowledge of our point of view available to others. To make that picture plausible we will need to explain more fully what it is to show one's cognitive, affective, or experiential state, and that will be an aim of the coming chapters. To make the picture informative, we will need to explain how it relates to other notions that are common fare for students of language, communication, emotion, and experience. In particular, the picture we've arrived at appears far removed from the kind of signaling that for many readers will seem ubiquitous in our own species, namely, the kind that exploits "speaker meaning". A half-century-old tradition in the philosophy of language, for instance, has it that the kind of "meaning" carried by

a tiger's scratch marks, a fly's eye stalks, or a peacock's tail feathers is profoundly distinct from that carried by a person's intentional honking of a car horn or, for that matter, her utterance of words. It may, further, seem that in describing speaking as a way of signaling, we only scrape the surface of the subtleties with which linguists and philosophers have wrestled for many decades. I'll argue, however, that this signaling picture is the beginning of wisdom about many aspects of verbal communication. Establishing this will span a few chapters, and our first step in that project will be a reconsideration of the notion of speaker meaning just mentioned.

### 3.3. Grice's ladder

Philosophers have spent great effort developing accounts of various kinds of communicative act. In this section I discuss H.P. Grice's well known approach. Grice couches this discussion in terms of the notion of meaning rather than that of communication, since the latter requires that information be not only sent but also received by its recipient, whereas the former does not and is thereby simpler. (A person might mean something without being understood, or even heard; in such a case we have an instance of meaning but not of communication.) The notion of meaning nevertheless contains many of the core elements needed for our study. As we proceed, we will find ourselves taking issue with the bulk of Grice's major contentions about the nature of meaning. It may for this reason be surprising that we are discussing his work at all. The reason is that although Grice's views about meaning are flawed in many ways, they are clues to a more adequate account that most effectively emerges once we see how he and other predecessors have erred.

In his influential 1957 article, Grice distinguished between two senses of 'mean'. One sense is exemplified by remarks such as "Those clouds mean rain," and "Those spots mean measles." The notion of meaning in play in such cases Grice dubs 'natural meaning'. Grice suggests that we may distinguish this sense of 'mean' from another sense of the word more relevant to communication, exemplified in such utterances as:

in saying, "You make a better door than a window", George meant that you should move,

and  
 in gesticulating that way, Salvatore means that there's quicksand over there,  
 and  
 in saying, "Look!", Alfonso means the man with the red hat,  
 and  
 in pointing in that direction, Mary means the one on the top shelf.

Grice used the term 'non-natural meaning' for this sense of 'mean', and in more recent literature this jargon has been replaced with the term 'speaker meaning'. (This terminology is misleading because according to philosophers' usage, an act can be one of speaker meaning with no sounds uttered or even any inscriptions made. For instance two hunters with no common language might communicate with pantomime, so that when one acts out the path of attack he means, in the sense of speaker meaning, that the other is to approach the mammoth from behind. In spite of the misleading nature of the jargon of speaker meaning I shall retain it rather than introduce new nomenclature.) Although he does not discuss the point, Grice's idea of speaker meaning can be expressed in either of two forms, exemplified above. One is that in which a speaker means *that* something is the case. Here what is meant is a proposition, and that proposition may be being put forth with the force of assertion or one of its cousins such as conjecture, prediction, or supposition. Call this *propositional speaker meaning*. (In the course of this chapter, this notion will fission into what we shall call 'factual speaker meaning' and 'illocutionary speaker meaning'.) Another form that speaker meaning may take has to do with an agent having something in mind. When Alfonso says 'Look!', pointing at an eagle passing overhead, he means a particular object without, as such, saying anything about the object in question. For this reason he is not asserting, conjecturing, or performing any other illocutionary act involving a proposition. Call this sort of case *objectual speaker meaning*. Both propositional and objectual speaker meaning are familiar from everyday communicative life.<sup>1</sup>

<sup>1</sup> Speaker meaning is usually construed as taking only propositional objects. Even when the force of an utterance is not within the assertion family (which includes assertion, conjecture, supposition, prediction, and so forth), but is rather imperative or interrogative, the content of that utterance is normally taken to be propositional. We nevertheless have a robust intuitive feel for objectual speaker meaning, and we can take objectual speaker meaning seriously without turning the clock back to a

After distinguishing between natural and (what we shall hereafter call) speaker meaning, Grice asks what features would enable a situation to be a case of speaker meaning. Grice first considers the suggestion that it is sufficient for speaker meaning that a person do something that influences the beliefs of an observer. This is clearly inadequate, however, since in putting on a coat I might, unbeknownst to me, lead an observer to conclude that I am going for a walk. Yet in such a case it is not plausible that I mean that I am going for a walk in the sense germane to speaker meaning. Nor do I mean any object, such as myself, my walking, or the prospect of my walking. (I might of course *mean to go* for a walk, but intending to do something is not, by itself, enough for speaker meaning either.) Thus performing an action that influences someone's beliefs is not a sufficient condition for speaker meaning, be it objectual or propositional.

Might performing an action with an intention, successful or not, of influencing someone's beliefs be sufficient for speaker meaning? It is not. I leave Smith's handkerchief at the crime scene to make the police think that Smith is the culprit. Here I leave the handkerchief where I do for the sake of influencing the beliefs of the police. However, whether or not I am successful in getting the authorities to think that Smith is the culprit, in this case it is not plausible that I mean that Smith is the culprit. Similarly, it is not plausible that in leaving the handkerchief where I do, I mean Smith, or Smith's guilt, or any other object. Accordingly, performing an action with an intention of influencing someone's beliefs is not sufficient for either objectual or propositional speaker meaning.

Perhaps what is missing in the handkerchief example is the element of overtness. This suggests another criterion, namely that of performing an action with the, or an, intention of influencing someone's beliefs, while

crude semantic theory typical of eighteenth- and nineteenth-century empiricism, according to which the paradigmatic unit of meaning is the noun.

One striking example of the use of objectual speaker meaning in contemporary literature is in Chuck Palahniuk's *Choke* (2001). Here the narrator remarks how his friend Denny has found a newspaper ad apparently placed by his parents, "Next to me, Denny reads: 'Free to good home, twenty-three-year-old-male, recovering self-abuser, limited income and social skills, house trained.' Then he reads a phone number. It's his phone number. It's my folks, dude, that's their phone number," Denny says. 'It's like they're hinting.' He found this left on his bed last night. Denny says, "They mean me." (p. 122)

intending that this very intention be recognized. Grice contends that even here we do not have enough for speaker meaning. He offers the example of Herod, who presents Salome with St John's severed head on a charger, intending that she discern that St John is dead and intending that this very intention of his be recognized. Grice observes that in so doing Herod is not *telling* Salome anything, but is instead deliberately and openly letting her know something. Grice concludes that Herod's action is not a case of speaker meaning either. The problem is not that Herod is not using words; we have already considered hunters who mean things wordlessly. The problem seems to be that to infer what Herod intends her to, Salome does not have to take his word for anything. She can see the severed head for herself if she can bring herself to look. By contrast, in its central uses, telling requires a speaker to intend to convey information (or alleged information) in a way that relies crucially upon taking her at her word. Grice appears to assume that at least for the case in which what is meant is a proposition (rather than a question or an imperative), speaker meaning requires a telling in this central sense. What is more, this last example is a case of performing an action with an intention of influencing someone's beliefs, even while intending that this very intention be recognized; yet it is not a case of telling. Grice infers that it is not a case of speaker meaning either.

Grice holds that for speaker meaning to occur, not only must one (a) intend to produce an effect on an audience, and (b) intend that this very intention be recognized by that audience, but *also* (c) one must intend this effect on the audience be produced at least in part by their recognition of the speaker's intention. I shall return in a moment to his reasons for introducing this third element. Before doing so, however, let's pause over these cases in which one overtly manifests or displays an object, situation or state of affairs. In a similar spirit to that of Grice, Clark (1996, p. 129) remarks that it would be unnatural to describe a person, Ernest, who throws open a window to show Matthew the rain outside, as meaning that it is raining outside. Let us assume that as in the Herod–Salome case, Ernest's action is overt. Thus rather than, say, causing the window to be open without enabling Matthew to discern that Ernest is the one who opened it, Ernest opens the window in full view of Matthew, where it is also clear that Ernest can see Matthew, that Matthew is aware that Ernest can see him, and so on. Even with the example thus clarified it may not seem

entirely colloquial to describe Ernest as meaning that it is raining outside (or some objectual analogue). This, however, could be due to the fact that so describing the situation seems roundabout in contrast simply to describing Ernest as showing the rain outside. It would be a bit like describing a person as trying to swim across a river when he in fact is succeeding in doing so; it is true that he is trying to swim across, but this understates the situation and so can be misleading. Likewise, the awkwardness of describing Ernest as meaning that it is raining outside may be due to the fact that for many conversational purposes we could more informatively describe him as showing that it is.

To circumvent this distracting feature, then, consider a case in which what is meant is not obvious in what is shown. Adverting to what is meant will then not be apt to mislead. Accordingly, consider a variation of Clark's case in which Ernest overtly opens the window to reveal threatening skies. He shows the looming weather, and yet Matthew, not appreciating the impending storm, might be puzzled, asking, "What's your point?" Alice might intercede, saying, "He means there's a tornado over there. Can't you see the funnel cloud?"

In this last case we have an act that shows an impending tornado that can at the same time be described as the agent's either meaning that tornado (objectual speaker meaning), or that the tornado is impending (propositional speaker meaning). More precisely, a single act can be both a case of speaker-meaning that P (or speaker-meaning  $\alpha$ , for the objectual case), and a case of showing that P (or showing  $\alpha$ ). Hence showing does not preclude speaker meaning. A variation of the Herod–Salome case makes an analogous point. Suppose that Salome is seriously nearsighted and does not trust her vision to determine whether what she sees is a severed head, or John wearing a charger-shaped necklace with the rest of his body occluded. She might accordingly not only wonder, "What is Herod showing me—a dead John or a partly occluded John?" She might equally wonder, "What does he mean? That John is dead or merely that this shackle is too tight?" The window example did not employ a telling that crucially involved taking someone at their word, for Ernest showed Matthew the impending storm even if Matthew did not immediately grasp the significance of what he had seen. Similarly, this revised Herod–Salome case does not involve a telling in this sense, for here again Salome does not need to take Herod's word for anything to determine that John is indeed no longer. She need

only look closer. After doing so, she has her answer: "Ah. He means John is dead."<sup>2</sup>

Other cases point to a similar conclusion. Imagine that A asks B for a game of squash, to which B replies by showing an extensively bandaged leg that clearly incapacitates him for sports (Schiffer 1972, p. 56). Here it seems natural to say that B means that he cannot play. Yet here B presents the leg with the, or an, intention of influencing A's beliefs, while intending that this intention be recognized. Here it does not seem that he need have any more complex intention than this, such as an intention to produce a belief at least in part by recognition of this intention. Again imagine a mathematics teacher who, at the end of a demonstration of the Pythagorean Theorem, utters, "And so,  $a^2 + b^2 = c^2$ ." The teacher has shown that  $a^2 + b^2 = c^2$  with the intention of influencing her audience's beliefs while intending that this intention be recognized. At the same time it is colloquial to describe her as meaning that  $a^2 + b^2 = c^2$ .

<sup>2</sup> Grice might also be being misled by the assumption that little of communicative interest lies between the criteria he proposes for natural meaning and those he proposes for speaker meaning. In articulating natural meaning, Grice gives five conditions that must be satisfied:

1. For natural meaning, 'x means that p' and 'x means that p' entail p. (One cannot consistently say, "Those spots mean measles, but he hasn't got measles.")
2. For natural meaning, 'x means that p' does not entail that something was meant by x.
3. For natural meaning, 'x means that p' does not entail that somebody meant something by x.
4. For natural meaning, 'x means that p' cannot be reformulated as 'x means that "p"'.<sup>3</sup>
5. For natural meaning, 'x means that p' can be reformulated as 'the fact that x occurred means that p'.

He goes on to argue that in cases of speaker meaning all five of these conditions are violated. However, Grice gives us no reason to believe that these five conditions stand or fall together. Instead he simply asserts, "I do not want to maintain that all our uses of 'mean' fall easily, obviously and tidily into one of the two groups I have distinguished; but I think that in most cases we should be at least fairly strongly inclined to assimilate a use of 'mean' to one group rather than to the other" (p. 215).

Return now to the case of the bandaged leg. Letting x be 'Jones presents his heavily bandaged leg' and p be 'Jones is unable to play squash', we may see that while condition 1 above is met, conditions 2 and 3 are not. Jones's presenting his heavily bandaged leg means that he cannot play' does entail that something was meant by Jones's presenting his heavily bandaged leg. Jones's presenting his heavily bandaged leg means that he cannot play' does entail that somebody (namely Jones) meant something by Jones's presenting his heavily bandaged leg.

Here then is a case showing that Grice's five conditions do not stand or fall together, but the case is by no means *recherché*. I conjecture that Grice failed to see that cases satisfying some but not all of the five conditions are relatively common, with that background. In seeing that the Herod case and its ilk do not violate all five conditions as do his showcase examples of speaker meaning, he inferred that they must merely be cases of natural meaning. We are now in a position to see that this was a false step. (See Denkel 1992, and Wharton 2002, for further discussion of intermediate cases such as those we have considered.)

Many theorists take Grice's main result in 'Meaning' to be a treatment of speaker meaning that involves a certain complex communicative intention. After (in our view incorrectly) rejecting the Herod case as exemplifying speaker meaning, Grice asks us to consider the difference between Z showing Mr. X a photo of Mrs. X and Mr. Y in a love embrace, and Z presenting Mr. X with a drawing depicting the same situation. As one might expect, Grice denies that Z's presenting the photo is a case of speaker meaning, and this position need not detain us. More important is his analysis of the drawing case. For Grice, if the presenting of the drawing is to be a case of speaker meaning, then Z must be doing so not just with the intention of making Mr. X believe there is an affair, but must also intend to make him believe this by means (at least in part) of Mr. X's recognition of this very intention. That is, Grice holds that Z must intend that Mr. X reason, on some level: "Z is trying to tell me something, namely that things are as his picture depicts them. So, since Z is evidently trustworthy, I should conclude that that is how things are." The intention to produce a belief or other attitude by means (at least in part) of recognition of this very intention has come to be called a *reflexive communicative intention*. Figure 3-2 summarizes the progression of analyses and counterexamples alleged to motivate the requirement of reflexive communicative intentions.<sup>3</sup>

It is important, perhaps crucial, to communication that we take the words of others, often though of course not always, at face value. As we will see more clearly in Section 3.4, the possibility of such speech acts as telling and promising would be hamstringing if we did not do so. On the other hand, we should resist the temptation to overstate this point, by for instance taking it to be essential to speaker meaning, either propositional or objectual. Cases of speaker meaning do not always involve the intention that others rely upon our word; this has already been shown by the examples of the bandaged leg and the proof of the geometrical theorem. The same goes for a person who takes an unpopular stand, espousing her beliefs without any intention that her audience will come to agree with her. Indeed, the Pythagorean theorem example shows that in some cases we decidedly do not want others to rely on our word. The geometry teacher would be disappointed if any of her students came to believe that

<sup>3</sup> Our formulation thus far leaves it open whether the apparent reflexivity is eliminable. Grice took it to be, whereas I will argue below that it is not.

"Moon Over Miami"	Performing an action with (a) the intention to produce a belief in an audience, (b) the intention that this effect be achieved at least in part by the audience's recognition of the speaker's intention, and (c) the intention that the audience be aware of her intention (a) and (b)
The River Runs takes into the house that A is thinking of buying and has loose a big fat river rat. He knows that A is watching him and knows that A believes that S is unaware that he, A, is watching him.	Performing an action with (a) the intention to produce an effect on an audience, (b) the intention that this very intention be recognized by that audience, and also (c) the intention that this effect on the audience be produced at least in part by their recognition of the speaker's intention.
Hered shows Salome St. John's head on a charger with the overt intention of making Salome believe that St. John is dead; he intends that she discern that St. John is dead and also that this very intention of his be recognized.	Performing an action with the intention of influencing someone's beliefs while intending that this very intention be recognized.
I leave Smith's handkerchief at the scene of the crime with the intention of making the police believe Smith is the culprit. Still, it's not plausible that I mean that Smith is the culprit.	Performing an action with an intention, successful or not, of influencing someone's beliefs.
In putting on a coat I might lead an observer to conclude that I am going for a walk, yet this is not a case of speaker meaning.	Performing an action that influences the beliefs of an observer.

Figure 3.2: Grice's Ladder

theorem in part because she intended them to. Rather, she wants them to believe the theorem on the strength of the proof that has been given. Reflexive communicative intentions are out of place here, yet speaker meaning remains.

Reflexive communicative intentions are not necessary for speaker meaning. In fact, speaker meaning can occur without a speaker intending to produce *any* beliefs in an audience. A framed suspect might mean that she is innocent in saying, "I am innocent!", yet be fully aware that no one will believe her and perhaps, being realistic, not intending to convince anyone. She might not even intend her interrogators to believe that *she* believes she is innocent, since she might know that they are certain she is lying. Or, gazing into my newborn daughter's eyes I might say, "All things valuable

are difficult as they are rare," meaning what I say, without having the slightest intention to produce beliefs or other attitudes in her or in anyone else. Again, in the film *Sleeper*, Woody Allen's character Miles Monroe comes across, while exploring alone, a genetically modified chicken the size of a small house. Miles remarks, "That's a big chicken." In saying this he does not seem to be intending to produce an effect on anyone, himself included.<sup>4</sup>

Following a suggestion of Schiffer (1972, p. 15), Strawson (1970, p. 7), and Bennett (1976, p. 271), Avramides proposes in response to these kinds of case that the speaker is addressing himself, intending in particular to produce a certain cognitive effect in himself.<sup>5</sup> While it may be that in the newborn daughter and *Sleeper* cases the speaker is addressing himself, it neither follows, nor does it seem true, that in those cases the speaker is intending to produce any cognitive effect in himself. Certainly we sometimes address ourselves in order to produce a cognitive effect: 'I can do it!' as I sprint up the steep road, or '945-6743, 945-6743' as I try to internalize a phone number I just got out of the phone book. However, I already believe that all things valuable are difficult as they are rare. In fact it is a belief I have held firmly since encountering it in Spinoza two decades ago, and I actively believe it as I reflect upon the number of diapers I will have changed by the time I am forty. As a result it is quite unclear what cognitive effect I might be trying to produce in myself in saying what I do. Again, Miles Monroe does not need to produce in himself, or strengthen

<sup>4</sup> This discussion is indebted to Davis 1992a. We observe here also that Armstrong 1971 quite reasonably offers an account of speaker meaning in terms of objectives rather than intentions, his reason being that the latter notion is narrower than the former. One who intends a certain result must believe that the thing aimed at is within her power, while one who has that result as an objective need not do so. Presently we shall show an affinity between Armstrong's position and that offered below. However, just replacing 'intention' with 'objective' in Grice's account will not deal with the cases we have considered. It is not part of my objective to produce an effect in my newborn daughter in uttering the last line of Spinoza's *Ethica*. Similarly, it need be no part of the objective of the Framed suspect in maintaining her innocence to produce effects on her interrogators. Instead she may say what she does in order to make public, for anyone who may be concerned with the matter, her avowal of innocence. Her objective is simply to establish a pattern of consistently maintained innocence.

<sup>5</sup> Avramides' discussion here is confusing because the first responds to a case, due to Harman, of a person maintaining a proposition in full knowledge that no one will believe him, with the words, 'I think that in Harman's case the speaker is not really speaking to an audience at all' (p. 64). But then two pages later Avramides writes, "The misleading thing about Harman's case is that there appears to be an audience present. The speaker, however, does not really address his utterance to those present.... If this is true, why not say that in such cases the speaker intends his audience to be himself....' (p. 66). I shall take Avramides to hold the view that in these cases the speaker *does* have an audience, namely himself.

or activate, the belief that the chicken before him is big. His eyes have already done that for him. Likewise it is far from clear what belief the framed suspect might be trying to produce or strengthen or activate in herself as she maintains her innocence. The suspect knows perfectly well that she has never set foot in the part of town in which the crime was committed, and that she has no idea how to use the garotte with which the victim was killed.

It might, alternatively, be suggested that in the above cases the speaker addresses a virtual audience. This idea can take one of two forms. First of all, it might be suggested that the speaker is imagining addressing someone, not unlike the way in which a child might address an imaginary friend. Yet I certainly don't feel as if I am addressing an imaginary person when I utter the Spinozistic dictum, and I doubt that the framed suspect feels this way either. We can likewise easily construe Miles Monroe as not imagining that he is addressing anyone. On the other hand, the "virtual audience" proposal might take the form of a suggestion that in these cases, it is the speaker's intention that were someone in a position to understand and appreciate his utterance, the remark would produce an effect on that audience, for instance belief (Hyslop 1977). This proposal may well be correct, and we need not take issue with it. The reason is that one can have this counterfactual intention without intending to produce any effects on an audience. It is consistent to think both (a) were someone in a position to understand and appreciate an utterance, it would have an effect on them, and (b) no one is in a position to understand and appreciate the utterance. Likewise, one could (a) refrain from intending one's utterance to have an effect on anyone, and (b) intend that were someone in a position to understand and appreciate one's utterance, the remark would produce an effect on them. A counterfactual recasting of the notion of audience-directed intentions does not salvage the idea that speaker meaning requires audience-directed intentions.

We have shown difficulties in each of two ways of defending the idea that speaker meaning requires intentions to produce effects on an audience. Our responses to these defenses might raise the question why in the cases we have considered the speakers said anything at all rather than keeping their mouths shut. In the interrogation example the answer, as we saw above, is simply that the suspect wants it on record that she has maintained

her innocence. In the *Sleeper* example the answer is also clear: sometimes we are struck with a thought that it is natural to express. This is not different in principle from the way in which we are possessed by an emotion that it is natural to express. Indeed the two phenomena often overlap. In the *Sleeper* example, the speaker is not only expressing his belief that the chicken is large, but seems also to be expressing a sense of awe or wonder at the enormity of the fowl. My utterance of the last line of Spinoza's *Ethics* expresses my sense of both the burdens and pleasures of parenting, and I can find satisfaction in thus expressing myself. One need look for no deeper explanation than this.<sup>6</sup>

Speaker meaning does not require intentions to produce beliefs in an audience, even oneself, to say nothing of intentions to produce beliefs by means of recognition of one's intention. What then, is speaker meaning, and how could Grice and others have gone so wrong in defining it in terms of audience-directed intentions?

### 3.4. Intention: nests and hierarchies

After Grice formulated his reasons for imputing reflexive communicative intentions to producers of speaker meaning, among philosophers a cottage industry sprang up concerned with whether even these intentions were sufficient to produce the phenomenon in question. Those convinced that reflexive communicative intentions are necessary for speaker meaning soon agreed that they are not sufficient in light of cases such as one devised by Strawson 1964:

*The River Rat:* Homebuyer is inspecting a house for possible purchase, and his friend—call him Friend—is concerned to convince him that the home is rat infested. Friend arrives at the house at a time when he knows that Homebuyer is inside, and although he knows he is being watched, skulks around to make Homebuyer think that he, Friend, believes he is acting unobserved. Friend has a river rat that he places in a salient position for Homebuyer to see. He intends for Homebuyer to see the rat and reason as follows: "Although the rat display was rigged, Friend would not have put it there unless he believed that the house really

<sup>6</sup> The satisfaction found in these cases of self-expression is over and above the well-documented way in which expression of emotions can have a therapeutic effect, often with positive repercussions for bodily health. For further discussion see Pennabaker 1990.

is rat infested; hence Friend, who is reliable and honest, must intend me to believe that the house is rat infested."

Friend thus has the following complex of intentions. In so placing the rat, he intends Homebuyer to believe that the house is rat infested; further, he intends to produce this effect in Homebuyer by means of his recognition of Friend's intention. Friend thus meets Grice's conditions for reflexive communicative intentions. It's also colloquial enough to describe Friend as trying to convey to Homebuyer that the house is rat-infested. However, it does not seem that in placing the rat under the conditions we have described, Friend means that the house is rat infested.

Some proponents of Grice's program of construing speaker meaning in terms of communicative intentions responded to this case by adding another clause to the definition of speaker meaning.<sup>7</sup> In addition to (a) the intention to produce a belief in an audience, and (b) the intention that this effect be achieved at least in part by the audience's recognition of the speaker's intention, we require further that the speaker intend (c) that her audience be aware of her intentions (a) and (b). Unfortunately, examples that satisfy these three clauses, without, however, seeming to be cases of speaker meaning, have been found.<sup>8</sup> Trying to rule out these further cases with a more complex analysis threatens to make the Gricean approach to speaker meaning psychologically unrealistic.

Other authors have suggested that what unites the various cases we have observed to make trouble for the Gricean approach is that they involve a kind of covertness, or at least a lack of overtness. What is missing in the handkerchief case is the intention, on the part of the man planting the handkerchief, that his intentions be publicly available. Similarly, part of the intent of Friend in the River Rat case is to be covert (he wants Homebuyer to be unaware that he knows he is watching him), and this is what seems to prevent that case from being one of speaker meaning in spite of the complex communicative intentions involved.

However, the foregoing subtleties might seem moot in light of our realization that speaker meaning does not require intentions to produce beliefs or other attitudes. What at first seemed an ascent of Grice's ladder may now appear a fool's errand. Did we get off on the wrong foot in judging

that, since the overcoat and handkerchief examples do not exemplify speaker meaning, some more complex congeries of intentions must do the job? I suggest that the search for ever more complex intentions such as would be sufficient for speaker meaning conflated two things: (1) intentions to produce cognitive effects in some audience, and (2) intentions that one's intentional state be *manifest*, that is, publicly accessible, but not necessarily in fact discerned by anyone. One can have the latter intention while harboring no intention whatever that anyone come to be aware of one's own state of mind. Instead one need only intend that the intentional state be "out there": there to be discerned by anyone concerned to look whether or not anyone ever does.<sup>9</sup>

Intentions to produce beliefs in an audience are of course commonly found in the company of speaker meaning. Our discussion thus far shows that nevertheless it can manage without them. What it cannot eschew is overtness. Recall the above case in which a person dons an overcoat, leading others to conclude that she plans to go for a walk. Here there is no speaker meaning. However, imagine a revision of the case in which Hermione is arguing with a group of friends about the weather. Everyone else has been warning her of the dangers of going outdoors in the storm. She now dramatically swings open the closet door, dons an overcoat with a flourish, and grabs an umbrella while beadyly meeting the gaze of her friends. Here it seems clear that she means that she is intending to go for a walk, and I suggest that what differentiates this from the earlier case in which someone's intention to go for a walk is discernible is that Hermione is making her periphrastic intention overt.<sup>10</sup>

<sup>9</sup> There are two notions of reason, one requiring, the other not requiring, a subject's awareness of the reason in question. In accord with the latter notion, I might, unbeknownst to me, have a reason to believe that one of my employees is embezzling funds even if I am not aware of that reason because I have not noticed the Byzantine character of the balance sheets he has submitted. Ridding a proposal of Armstrong's according to this "objective" notion of a reason allows us to see an affinity between it and the approach taken here. Armstrong suggests (1971, p. 433) that a speaker S who means that P has the objective that the audience should have reason to believe that S believes that P. On the objective notion of a reason, S's objective need not include producing an effect on any audience. It need only include producing evidence, there for anyone to see, that S has a certain belief.

<sup>10</sup> Grice 1957 considers a case that is superficially similar to this one, but which is less clearly an instance of speaker meaning. In that example a father leaves lying around on the floor for the mother to see the bits of the china that their child has broken. Here it is doubtful whether the father means anything in leaving the china on the floor. As Grice describes the case, the mother might be open to any of a variety of reasons why the china is on the floor: perhaps the father didn't notice it, or had other more pressing issues to attend to, or was too furious to clean it up. This is why the example

<sup>7</sup> For a thorough review of the literature in the wake of Strawson, see Vlach 1981.

<sup>8</sup> One such case is the "Moon Over Miami" example of Schiffer 1972.

What, then, is overtness? To be overt an action of mine need not actually be discerned by anyone; I could flagrantly show contempt for someone, and if they are too distracted or obtuse to recognize what I am doing this will make the action no less overt. Intending that an action of mine be overt should likewise not require intentions to produce effects on others. Instead, it must, very roughly, be an intention that everything of relevance be "out in the open". One way of articulating this thought is with a hierarchy of conditions having the following structure:

- (a) I intend to show my contempt,
  - (b) I intend that my intention (a) be manifest,
  - (c) I intend that my intention (b) be manifest,
- and so on.

It is, however, difficult to accept this as a description of what is going on when I, for instance, flagrantly show my contempt for someone. The reason is that it is unlikely that I would harbor such an unending hierarchy of intentions. How then may we articulate the wanted notion of overtness? Another approach capturing the wanted notion of overtness found in an important class of cases of speaker meaning is in terms of an action done intending that (a) something be publicly discernible, and (b) this intention itself be publicly discernible as well. This is an intention part of whose content, the part expressed in clause (b), refers, *inter alia*, to the very intention whose content it is, with the result that there is no finite way of writing down the content of this intention from which self-reference has been elided. That does not, however, imply that to grasp this intention a thinker must grasp a thought that is in any sense infinitely "long". Instead the thought can be expressed, and thus grasped, quite succinctly.<sup>11</sup> Consider an analogy. I might be regaining consciousness from an accident that, as it was about to occur, seemed certain to kill me. As I contemplate my good fortune before opening my eyes to assess the damage I might think, 'This

is only superficially similar to the case of Hermione in the text. Hermione leaves little doubt about the intention with which she is acting. This is to say that Hermione overtly displays her intention, whereas the father does not make overt the intentions with which the china has been left lying around. Accordingly, while both the china and Hermione examples involve showing rather than telling, the awkwardness of describing the former example as a case of speaker meaning does not undermine the contention that Hermione means that she intends to go for a walk.

<sup>11</sup> Here I am indebted to Harman 1974, 1977, and to Recanati 1986.

thinking is miraculous."<sup>12</sup> The content of this thought refers to the thought token (a particular thinking with a spatiotemporal location, or at least a temporal location, and a content), and says of it that it is miraculous. It will, then, be true just in case that very thought token is miraculous. The fact that the content of this thought token cannot be finitely rewritten in a way that eschews all reference to the thought token itself, does not upset the fact that this content has determinate truth conditions, and a determinate meaning.

"Self-referential" thoughts do not, as such, appear incoherent.<sup>13</sup> If their content can be anchored by a particular thinking, perceiving, intending, and so on, there need be no difficulty in principle in grasping that content. This point is applicable to our elucidation of speaker meaning, for as we have seen, one way of meaning that P, namely that associated with (but not exhaustive of) propositional speaker meaning, is overtly to make the fact that P manifest. That is, in such a case one intentionally makes P manifest, intending as well that this very intention to make P manifest, itself be manifest. That suggests an elucidation of one form of speaker meaning along the following lines:

*Factual Speaker Meaning:* Where P is an actual state of affairs, S factually speaker-means that P iff

1. S performs an action A intending that
2. in performing A, it be manifest that P, and that it be manifest that S intends that (2).

P might be the fact that it is windy outside, or that S believes that it will rain, with A being the uttering of certain words or a non-conventional action such as the throwing open of curtains to reveal a looming storm. In the Herod case the P in question is the fact that St John is dead, and the action is Herod's presenting of St John's head on the charger. In the case of an irritated person scowling intentionally, the P in question is that he is irritated, and the A is the irritated person's production of the angry scowl. As with the case of the thought, 'This thinking is miraculous', Factual Speaker Meaning requires, in condition (1), that an intention (as opposed

<sup>12</sup> The example is due to Peacocke 2003.

<sup>13</sup> Some authors, such as Harman, hold that all intentions are self-referential. We need not commit ourselves to this strong claim in order to hold that intentions germane to speaker meaning are self-referential.

to a thought) be tokened. The content of that intention is then articulated in condition (2), which refers, *inter alia*, to the intention-token whose content it is. That intention-token will, in turn, be satisfied (as opposed to being true) just in case condition (2) holds. Once again, P might be manifest without anyone being aware of this fact. As a result one can intend that P be manifest without intending to produce effects on others.<sup>14</sup> In light of our account of Factual Speaker Meaning, then, one can mean something without intending to produce effects on any audience, and one can mean something in the course of making some state of affairs manifest.

Speaker meaning, as we have seen, also comes in an objectual variety, in which an agent means an object rather than that something is so. This has been exemplified in such cases as a person referring to someone with words, as in, 'Look over there!', or without, as in their simply pointing to or gazing pointedly at the thing in question. After observing such an act we may ask, "Which one do you mean?" if we are still unsure. Such objectual meanings must, like their factual analogues, be overt, and we now know how to elucidate this requirement. Where  $\alpha$  is a perceptible object,

*Objectual Speaker Meaning*: S objectually speaker-means  $\alpha$  iff

1. S performs an action A intending
2.  $\alpha$  to be manifest, and for it to be manifest that s/he intends (2).

Instead of the factual construal offered above, Herod's action may be construed as done with the intention of making St John's severed head manifest, while intending that this very intention be manifest as well. In the case of the overt scowl, the  $\alpha$  in question might be the agent's anger,

<sup>14</sup> Some authors have been skeptical of the very possibility of self-referential intentions. (Such skepticism is not to be confused with skepticism of the doctrine that all intentions are self-referential. One may accept that self-referential intentions are possible while remaining neutral on the question whether all intentions are self-referential.) Thus for instance Seibel 2003 writes, "the content of [the self-referential intention] contains an element which refers to the intention itself. But what does that element look like?... How does it single out the intention and nothing but it? By identifying features, i.e., properties which are exclusively possessed by the intention? But what could be these features?" Intentions admit of the same act/object dichotomy as do many other intentional states, and like other intentional events, are spatiotemporally located. On the modest assumption that no two intentions have identical spatiotemporal coordinates, we may then use such coordinates to individuate intentions. One might still wonder what the content is of an intention whose content refers to that very intention, which itself comprises both an intending (a state or act) and a content. One answer may be given in terms of an analogue of truth conditions applicable to intentions, namely satisfaction conditions. Just as the thought, 'This thinking is miraculous' will be true just in case that thinking is, indeed, miraculous, so too, the intention, 'This intention shall be manifest' will be satisfied just in case that intention is, indeed, manifest.

the agent intending that his intention to make his anger manifest be itself manifest. In fact, particularly when no conventional language is used, it may be difficult to know whether the case is objectual or factual speaker meaning. When an indicative sentence is uttered, that is strong evidence that the speaker's meaning is not objectual. Failing that evidence, however, not only may it be difficult to tell, we need not assume that there must be a fact of the matter whether the speaker's meaning is factual or objectual. It may well, that is, be indeterminate whether the speaker's meaning is one or the other. If that is so, then in such cases, a choice between an objectual and a factual way of representing what the speaker means, if such a choice must be made, is presumably to be made on pragmatic grounds.<sup>15</sup>

Both Factual Speaker Meaning and Objectual Speaker Meaning pertain to cases in which the speaker takes what she means to obtain (for the factual case) or exist (for the objectual case). Otherwise it is difficult to see how she could intend to make what she means to be manifest. Clearly we can also mean propositions that we do not think are true, as when we lie, and evidently we can also mean objects that we do not believe to exist, as in make-believe. We can also mean propositions whose truth we are not in a position to demonstrate. These are reasons why the above conditions are not formulated as necessary for speaker meaning generally. If our exploration of speaker meaning in this chapter has been accurate, however, it would be a mistake to infer, from the premise that *some* forms of propositional speaker meaning do not involve intentions to manifest facts, to the conclusion that *none* ever do. Likewise it would be a mistake to infer, from the premise that some forms of objectual speaker meaning do not involve intentions to display or manifest objects, to the conclusion that none ever do. We are now, however, in a position to build on what we have learned to develop a general account of speaker meaning.

### 3.5. Speech acts and handicaps

Factual speaker meaning usually occurs in the context of a speech act. Let me make clear that speech acts are to be distinguished from acts of

<sup>15</sup> In Chapter 6 we will return to issues of indeterminacy in ascribing attitudes to agents.

speech. When I test a microphone, utter lines on stage, or practice a speech in the shower, I am performing acts of speech but no speech acts. By contrast, a speech act is any act that can be performed by, under the right circumstances, speaker-meaning that one is doing so. I can raise a question ('I ask you what time it is'), make a statement ('I state that it is 5 p.m. '), issue a command ('I command you to make the appointment on time') by saying that I am doing so in such a way as to speaker-mean it. This is why questioning, stating, and commanding are speech acts. I can also perform one of these acts by speaker-meaning a content in a certain way but without saying that I am doing so. I can assert that snow is white without saying that I am asserting that snow is white. I can do that simply by uttering the words, 'Snow is white', meaning this as an assertion.

Since I cannot scare or persuade you by saying that I am doing so, scaring and persuading are not speech acts. Although speech acts are only a small portion of the acts that we are able to perform, some such acts are nevertheless quite momentous: so long as I have the authority and other conditions are in place, I can excommunicate, appoint, bequeath, or even declare war by speaker-meaning that I am doing so. In this section I will argue that central features of certain speech acts can be explained in terms of their being handicaps in the sense of that term introduced in Chapter 1.

How can the biologist's notion of a handicap be so much as relevant to the elucidation of speech acts? To see why it is, consider that it's a bit of common sense that speech acts often, though not always, have expressive dimensions. It is essential to a wide variety of speech acts that when performed in all propriety they express a state of thought or feeling.<sup>16</sup> A sincere assertion expresses one's belief, a sincere promise one's intention, a sincere apology one's regret. Further, one of the dicta we formulated in Chapter 2 has it that expression is a species of showing: another has it that it is a species of showing one's state of mind, state of heart, or state of experience. It follows that one who sincerely asserts P shows her belief that P, one who sincerely promises to do something shows her intention of doing so, and so forth. We have also suggested, in Section 3.2, that this notion of showing comes in at least three forms: showing that something is so, showing  $\alpha$  (where  $\alpha$  is an object of

perception), and showing how something feels, be it an experience, emotion, or mood. Beliefs, intentions, and the like are not, apparently, possible objects of perception, nor is it plausible that there is a way that a belief or intention feels. If we are to show our belief by means of a sincere speech act, it will evidently have to be by showing that we harbor the belief that we have expressed. Yet how can a mere speech act provide enough evidence for the presence of a belief in order to show it?

Consider a well known conception of assertion. In *Speech Acts*, John Searle characterizes assertion of P as, "an undertaking to the effect that P represents an actual state of affairs" (1969, p. 66). Because undertakings are more naturally thought of as undertakings to action than as having propositional objects, it is not entirely clear what such an undertaking consists in. Minimally, we may construe an undertaking to the effect that P represents an actual state of affairs as putting the assertor in a position such that she is correct on the issue of P if P obtains, and incorrect on the issue of P otherwise. Assertion, however, does more than put the assertor at risk of being either correct or incorrect on the issue of the asserted proposition. After all, one who guesses that P is also correct on the issue of P if P obtains, and incorrect on that issue otherwise. Instead, when a speaker asserts that P, it is within the rights of an addressee to reply with the challenge, "How do you know?" In response to such a challenge the assertor should either offer reasons of her own, or defer to another ostensibly reliable authority ("I read it in the *Times*," "Susan told me."). If she cannot do either of these things she should retract the challenged assertion.<sup>17</sup>

Compare assertion with the case of conjecture: If I offer P as a conjecture, then it is inappropriate to reply with the challenge, "How do you know that?" For instance, I might conjecture that a black hole inhabits the center of the Milky Way. As with the case of assertion, what I say is right or wrong on the issue of P depending on whether P is true. However, it is inappropriate to reply to my black hole conjecture with a challenge to show that I know what I say to be true. A legitimate challenge to a conjecture would instead come in the form of showing that what I say

<sup>16</sup> Green forthcoming a develops the point in further detail.

<sup>17</sup> We need not assume that the normative notions invoked here are reducible to moral norms or to norms of theoretical or practical rationality. They may be, but we need not settle that issue here.

is demonstrably wrong, or at least very unlikely. In that case I would be obliged to retract that conjecture. Likewise, a conjecture should be backed with some justification or other; failing that it should be put forth as a guess rather than a conjecture. However, if you do put it forth as a conjecture, the required justification need not be as compelling as that for an assertion.

Assertions, conjectures, suggestions, guesses, presumptions, and the like are cousins sharing the property of commitment to a propositional content. They differ from one another in the norms by which they are governed. In addition, in assertions, conjectures, suggestions, and presumptions, the speaker incurs a vulnerability—not just a liability to being in error, but also a mandate to defend what she has said if appropriately challenged. These liabilities to error and injunctions to defend what has been said put the speaker at risk of losing credibility in the community in which she has a reputation.

Assertions, conjectures, suggestions, and presumptions, and the like thus carry a cost. Why should speakers bother to pay it? Presumably speakers are willing to incur such costs because the information they contribute to the group can help achieve the group's aims—finding food, prevailing over rivals, securing shelter, and so forth. In addition, an individual might want to be the first one to provide such information in order to enhance her or his credibility within the group: saying something unobvious and useful to the community can enhance one's prestige. Yet in the absence of some mechanism for vouchsafing honesty, these speech acts will be prone to abuse by those who take assertion lightly in the hope of scoring some epistemic points. After all, if I say enough, I'll eventually get something right. This temptation threatens to make assertors, conjecturers, and the like on the whole less worthy of our belief: liars, and those who say things on insufficient evidence, threaten to undermine the efforts of sincere speakers in the way that brightly and dishonestly colored frogs threaten the credibility of signals sent by those frogs who are both brightly colored and noxious. In chapter 2 of *Utilitarianism*, Mill emphasizes the value of making statements when he refers to, “the trustworthiness of human assertion, which is not only the principal support of all present social well-being, but the insufficiency of which does more than any one thing that can be named to keep back civilization, virtue, everything on which human happiness on the largest scale depends” (Mill 2002, p. 23). That's pretty high praise.

Mill may be overstating the point, but he is responding to a legitimate anxiety.

I propose that many linguistic communities have found ways of circumventing this danger of the erosion of speakers' credibility. Recall that the long eye stalks of the male stalk-eyed fly show its viability. The magnitude of the spider's vibration on its web show its size. So too, the liabilities that an assessor undertakes by incurring the commitments characteristic of that speech act give strong evidence that the assertion is both sincere and justified. Certainly not conclusive evidence: many lie, and many believe things on insufficient evidence. Nevertheless, by asserting something sincerely you perform an act that, because of the handicap it involves, *shows* your belief. So, however, does a sincere conjecture. In addition, the assertion, precisely because of the strict standards to which it is held, shows your belief *as justified at a level appropriate for knowledge*, while the sincere conjecture does not. (The sincere conjecture shows your belief as having some justification, but not sufficient justification for knowledge.)

We have discussed speaker meaning in objectual and factual forms. However, when I mean that P as an assertion (rather than as a conjecture, etc.) I do not therefore mean that I am asserting that P. I might just mean P. Yet, the fact that I am putting forth P as an assertion rather than a conjecture is part of what I mean. This is why it would be appropriate for someone who doesn't know the force of my utterance, to ask, “Do you mean that as an assertion, a conjecture, a guess, or what?” How, then, may we understand the contribution that illocutionary force makes to speaker meaning? We may do so as follows. I can mean P *as an assertion*, rather than in some other way. What makes it the case that I mean P as an assertion is that in so doing I undertake a certain set of responsibilities and liabilities, namely those that we have outlined above as distinctive of assertion as against other speech acts. What makes it the case that I mean P as a conjecture is that in so doing I undertake a distinct set of responsibilities and liabilities. And so on. My meaning P with one force rather than another is not a matter of *what* I mean but rather *how* I mean it. We are now in a position to offer a characterization of speaker meaning distinctive of speech acts with propositional contents:

*Illocutionary Speaker Meaning*: S illocutionarily speaker-means that P  $\varphi$ 'ly, where  $\varphi$  is an illocutionary force, iff

1. S performs an action A intending that
2. in performing A, it be manifest that S is committed to P under force  $\varphi$ , and that it be manifest that S intends that (2).

One illocutionarily speaker-means that P assertorically if one performs an action intending that in so doing it be manifest that one is committed to P assertorically, and that this itself be manifest. One illocutionarily speaker-means that P as a conjecture if one performs an action intending that in so doing it be manifest that one is committed to P as a conjecture, and that this itself be manifest. These characterizations are not circular because 'illocutionarily speaker-meaning P as an assertion', 'illocutionarily speaker-meaning P as a conjecture', and so on may be fleshed out in terms of conversational norms. One way to make manifest that you are committed to P in a certain way is to make explicit that you are: use such words as "I assert that P." Parenthetical expressions work in this way as well, as in "P, I claim." Or you could let context make clear the force of your utterance, so that by uttering P free-standing, rather than embedding it in a larger sentence, you rely on context and the reasoning capacities of your addressees to discern that P is meant as an assertion rather than in some other way.<sup>18</sup>

Speaker meaning is a matter of overtly showing an object, overtly showing a state of affairs, or overtly showing one's commitment—both the modality of that commitment and its content. Each of the three conceptions of speaker meaning we have offered—Objectual Speaker Meaning, Factual Speaker Meaning, and Illocutionary Speaker Meaning—corresponds to one of these three forms of speaker meaning. For this reason we are now in a position to give a general account of speaker meaning, including the propositional, objectual and illocutionary cases:

<sup>18</sup> This definition of illocutionary speaker meaning only applies to speech acts with propositional contents. It does not apply to speech acts with interrogative contents (*where is John?*, *how many apples are in the bowl?*), or what I would call imperatival contents (*do shut the door*, *to eat the apples*). I believe that the account can be generalized to include such cases, but such a generalization would over-burden this chapter.

Also, just to forestall any worries about an impending regress: One who asserts P by saying that she is doing so most likely will utter some such words as, "I assert that P." What, then, is the status of the embedding sentence? It too will be being put forth with assertoric force. You could of course indicate that by saying that you are doing so, but this approaches the margins of precision. Most likely you will just let context make clear that you are putting forth the sentence, "I assert that P" as an assertion rather than in some other way.

*Speaker Meaning*: S speaker-means something just in case S either objectively speaker-means something, factually speaker-means something, or illocutionarily speaker-means something.

If the foregoing is correct, then we may understand speaker meaning as a species of signaling, whether or not that speaker meaning is sincere or justified. This framework will also enable us to see culture, human or otherwise, as standing in where biological fixity does not to create stable signaling systems. This is not to say that all cultures are sensitive to distinctions among asserting, conjecturing, guessing, and the like; some might well not bother with these subtleties. I do, however, suggest as an empirical hypothesis that in those cultures in which one can handicap oneself by means of a speech act, communicative stability is served.

### 3.6. Alternative accounts of speaker meaning

In a survey of literature up to 1980, Vlach 1981 shows that the theories propounded in Grice 1957, Grice 1969, Armstrong 1971, and Bennett 1976 founder on one or another version of the "proof" example, in which the geometry teacher intends her students to believe what she says, but not because of their recognition of her intention that they do so.<sup>19</sup> Vlach also shows that many more recent analyses of speaker meaning are objectionable for various reasons. Following a suggestion of Searle's, Vlach propounds an account of speaker meaning as follows: U means P by x iff U does x in the belief that he is thereby committing himself to P (1981, p. 382). This suggestion has the virtue of allowing for the possibility of speaker meaning in the absence of intentions to produce effects on an audience. It also captures a crucial feature of a significant class of speech acts, namely that they involve a form of commitment. However, the undertaking of commitment is not a necessary condition of every speech act. Speech acts are typically instances of speaker meaning,<sup>20</sup> but not all

<sup>19</sup> See, however, note 3 above for a qualification of this remark about Armstrong.

<sup>20</sup> I say typically here because there could be speech acts governed by a "strict liability" constraint. Thus for instance grasping a certain conventionally defined object could constitute agreement to enter into military service for a certain country. Arguably, this was precisely the case in England in the

involve the undertaking of commitment. For instance, it is hard to see how imperatives such as commands involve the undertaking of commitment. Likewise for greetings such as "Hello" or other of Austin's "behaviours" such as apologies and congratulations. Furthermore, while it seems clear that conversational implicature is a form of speaker meaning, one who conversationally implicates something is not thereby committed to the content thus implicated. If in answer to the question, 'Where is Mary?' asked by someone hoping to visit her, I answer, 'Somewhere in France', I may conversationally implicate that I am unable to be any more specific about her whereabouts. However, in so answering I have not committed myself to the proposition that I cannot be any more informative than I have been. This is in contrast to conventional implicature, in which a speaker's use of an expression commits her to a content by virtue of the conventional meaning of that expression even though that content is no part, or at least not an explicit part, of what she says.

Neale 1992 also denies that reflexive intentions to produce beliefs in an audience are necessary for speaker meaning, while adhering to the view that speaker meaning admits of a uniform characterization. He suggests the following account of speaker meaning: by uttering *x*, *U* meant that *P* iff for some audience *A*, (1) *U* uttered *x* intending *A* actively to entertain the thought that *P* (or the thought that *U* believes *P*), (2) *U* uttered *x* intending *A* to recognize that *U* intends *A* actively to entertain the thought that *P*, (3) *U* does not intend *A* to be deceived about *U*'s intentions (1) and (2). Observe first that this is a considerable weakening of the analysis, since it will follow from this account that under normal conditions, one who asserts "If *P*, then *Q*" will mean *P* and *Q*, and not just the conditional. Similarly, normally a person who asserts "It is not the case that *P*" will mean both *P* and its negation on Neale's account of speaker meaning, and one who recites *P* on stage will mean *P*. These consequences do not as such undermine Neale's proposal. However, that proposal runs afoul of the fact that we sometimes direct our remarks to those whom we know are not even comprehending what we say. We have already seen this in the

eighteenth century: grasping the 'King's shilling', whether one intended to or not, and whether or not one was aware of the consequences, constituted an agreement to enter into England's navy. That agreement might have been a speech act in spite of the fact that in so grasping the Queen's shilling one does not mean that one is willing to join her navy. See Hare 1989 and Green 1997 for further discussion, and I recur to the point in Section 4.3.

examples of my speaking to my newborn daughter, and of Miles Monroe remarking on the size of the chicken. Or I might, out of determination to fulfill the duties of my contract, carry through a complex lecture on a topic that I know is well over the heads of my students. Here when I state that the class *S* of all Borel sets coincides with the  $\delta$ -ring generated by the class *U* of all open sets, I exemplify speaker meaning. However, I do not meet Neale's condition (2), for I do not intend my audience to entertain the thought that the class *S* of all Borel sets coincides with the  $\delta$ -ring generated by the class *U* of all open sets. Indeed, because I am sure that none of my students is listening, I do not even intend that they entertain the thought that I believe this. As a result Neale, who waters down the Gricean requirement of reflexive communicative intentions about as much as possible while retaining its spirit, does not capture what is central to speaker meaning.

Davis (1992a, 1992b, 2003), as we have mentioned, is one source of our earlier arguments that speaker meaning does not require audience-directed intentions. His arguments apply to well known theories offered by Grice 1957, 1969, 1982, Schiffer 1972, and Avramides 1989, among others, and I shall not rehearse them here. His own position, developed most fully in Davis 2003, naturally provides an account of speaker meaning that eschews reflexive communicative intentions: *S* meant that *P* by producing *e* iff *S* performed an observable action as a direct and undisguised indication that he occurrently believes *P* (p. 57). 'Observable', because Davis holds that all speaker meaning requires making one's thought, or ostensible thought, public. 'Direct' because Davis wishes to focus his discussion on what he calls the 'exclusive' sense of 'meaning', as opposed to an 'inclusive' sense which includes all that one implies, conversationally, conventionally, or otherwise. Undisguised because of the following example: John and Mary are trying to fool George, and so Mary pretends that she is in great pain. John rushes to her aid, performing an observable action as a direct indication that he believes she is in pain. However, Davis observes, John does not here mean that Mary is in pain; instead he is merely pretending unintentionally to manifest that belief. Davis refers to this as a covert simulation of an unintentional manifestation of belief, and for brevity rules out cases of this sort by requiring that speaker meaning must be undisguised (ibid, pp. 54–6). Finally, 'occurrent' because Davis holds that one can only mean *P* if *P* expresses an occurrent thought.

One who asserts P expresses the belief that P, indeed expresses *her* belief that P if she is sincere, and so Davis gives us a necessary condition for assertion. Yet even leaving aside the non-indicative case, we can see that the condition is not necessary even for other indicative speech acts such as supposition or conjecture. Given his remarks about guessing (1992b, p. 226) Davis would likely treat one who puts forth P as a conjecture as expressing the belief that P is probably true. But that is not what I mean when I put forth P as a conjecture. Granted, I would not be inclined to conjecture that P if I did not think that P was at least reasonably likely, but this does not imply that this is what I mean when I utter, 'The cause of the perturbations in the planet's orbit is a distant white dwarf,' with the force of a conjecture. The difference comes out dramatically when we understand the kind of commitment undertaken in the forwarding of a conjecture. One who conjectures that P is apt to be right about P if indeed P is true; apt to be wrong about P if P is not true. This much cannot be said for one who means or believes that P is probably true. I can mean or believe that P is probably true without being shown wrong if indeed P turns out to be false. Probabilistic assessments do involve the undertaking of commitments, but the sorts of commitments they involve are quite different from those undertaken by such "qualitative" judgments as conjectures.<sup>21</sup> Davis's treatment also founders on examples of putting forth P with the force of a supposition for the sake of argument; in such a case the speaker means P, but she expresses neither belief nor any particular degree of belief that P.

Davis's account of speaker meaning requires that the thought in question be recurrent, but this is also too restrictive. Many things that we say, and mean, are said out of habit; in other cases things we mean are conceptually quite complex, and when we say them it is not clear that we need to entertain the thought expressed by the words that we use in order for our utterance to be a case of speaker meaning. I recently reproduced the proof that the class S of all Borel sets coincides with the  $\sigma$ -ring generated by the class U of all open sets, and just last Tuesday I was rehearsing both the proof and its conclusion. I still believe that conclusion, and I could formulate it verbally as well. Suppose I do formulate it verbally without consciously contemplating the theorem. This seems a clear case of speaker meaning in

spite of the fact that right now I lack the energy or concentration needed to entertain the thought that those words express.<sup>22</sup>

We may also note that not all speech acts characteristically express intentional states. In discussing Vlach's proposal we have already mentioned greetings. Austin's "exercitatives" are another case: it is difficult to see that excommunicating, for instance, expresses any intentional state. Davis's proposal is thus at best a sufficient condition for speaker meaning, and is therefore in no conflict with the sufficient condition that we have offered above in the form of our Sufficient Condition for Propositional Speaker Meaning.

Sperber and Wilson (1995) offer an account of communication with affinities to the view of speaker meaning propounded here. In particular, while Sperber and Wilson do not attempt to characterize speaker meaning, they do offer a view of communicative intentions from which a view of speaker meaning may be developed. On their official account, a fact is *manifest* to an individual just in case that individual is capable of either perceiving it or inferring it from other things she knows or thinks probable (p. 39).<sup>23</sup> A *cognitive environment* of an individual is the set of all facts that are manifest to him (p. 39). A cognitive environment is *shared* just in case it is manifest to more than one individual (p. 41). A *mutual cognitive environment* is a shared cognitive environment in which it is manifest which people share it (p. 41). Every fact or assumption in a mutual cognitive environment is what Sperber and Wilson call *mutually manifest* (p. 41). Further, behavior that makes manifest the intention to make something manifest is *ostensive behavior*, or just *ostension* (p. 49). Finally, to *communicate*

<sup>22</sup> Davis holds that even beliefs that I am not now consciously entertaining are occurring to me (2003, p. 326). For this reason he may not find the example concerning Borel sets compelling. However, such a view of occurrent belief seems untenable, implying as it does that all those beliefs I hold but am not now consciously entertaining are occurring to me. For more discussion of Davis see Green 2007a.

<sup>23</sup> I use this account of manifestess rather than the official one that Sperber and Wilson offer on page 39 of their 1995, because that official one does not square with the use to which they put it. Their official definition has it that a fact is manifest to an individual just in case that individual is capable of representing it mentally, and capable of accepting that representation as true or probably true (p. 39). This definition counts as manifest any fact that I am capable of thinking about and believing true or probably true, regardless of my evidence. Accordingly, if someone outside my range of vision is walking towards me, then on this definition, the fact that he is doing so is manifest to me. However, Sperber and Wilson make clear that this is not their view; rather, the fact in question becomes manifest only when someone or something draws my attention to the person walking towards me, or does something that makes that approaching person visible (p. 49). For this reason I shall allow Sperber and Wilson's gloss of the notion of manifestess to stand in as a definition.

<sup>21</sup> Green 1999c defends this point in further detail.

*intentionally by ostension* is to produce a stimulus with the aim of informing an audience of something, and intending to make it mutually manifest to audience and communicator that the communicator has this informative intention (p. 61).

An ostensive-communicative intention, then, for Sperber and Wilson aims at producing a cognitive effect on an audience, and we have seen in this chapter that this aim is not a necessary condition for speaker meaning. However, an intention to make mutually manifest to audience and communicator one's communicative intention does not require this; in just requiring mutual manifestness it demands only that one's intention be there for public view, so to speak. In this way it is like our notion of overttness. Yet it is unlike our notion of overttness in that it treats cases as ostensive intentional communication that our account of speaker meaning would not. An example is the case, considered above, of the river rat. In that example, Friend deposits a river rat in a house that he knows is being inspected for possible purchase by Homebuyer, and he does so in such a way as to make sure that Homebuyer observes him. However, Friend makes it look as if he thinks he is acting unobserved. In such a case we do not have speaker meaning, and our account of speaker meaning in 3.4 respects this fact. However, Sperber and Wilson's account of ostensive communication does not do so. To see why, observe that Friend (F) produces a stimulus aiming to achieve a cognitive effect in Homebuyer (H), namely the belief that the house is rat-infested (R). Further, F intends to make it manifest to H that he has this informative intention. After all, F intends H to be aware that he is trying to get him to believe that R. In fact, F intends to make it *mutually manifest* that he is trying to get H to believe that R.

Here is why: Let  $R_{F \rightarrow H}$  be the proposition that F is trying to get H to believe that R. We know that  $R_{F \rightarrow H}$  is part of F's cognitive environment, since surely F is either aware or capable of becoming aware of what he is trying to do. In addition, by the description of the case we know that  $R_{F \rightarrow H}$  is part of H's cognitive environment. So  $R_{F \rightarrow H}$  is part of F and H's shared cognitive environment. To show that  $R_{F \rightarrow H}$  is part of F and H's mutual cognitive environment, we need only show that it is manifest to H, and manifest to F, that  $R_{F \rightarrow H}$  is part of F and H's shared cognitive environment. (I consider just the case of H; the same reasoning will carry over to R.) This is easily done. H can surely see that  $R_{F \rightarrow H}$  is part of his own cognitive environment. Can H see that  $R_{F \rightarrow H}$  is part of F's cognitive

environment? Well, H can tell that F is trying to get him to believe that R; and presumably H knows that most people are aware of what they are trying to do. Consequently, H should be in a position to figure out that F is aware  $R_{F \rightarrow H}$ . We see then that the Sperber and Wilson approach treats the river rat example as a case of ostensive intentional communication, while the account of speaker meaning offered in 3.4 does not do so. That is good reason for preferring the present account to that of Sperber and Wilson, in spite of the fact that these authors correctly identify the intention to make facts manifest as an important element of communication.