

## Perfect Speaker Theory

AC Grayling

One thing we all agree is that meanings do not just fall out of the sky; they were and are at least in important part generated by what users of signs do with them. This suggests that the semantics of natural language is the product of the history of its pragmatics. It does not automatically follow, of course, that therefore the right way to go about giving a theory of meaning is in some sense to do it from the pragmatic end: but in the larger project I wish to suggest a reason for doing so, and I suggest that what increases the plausibility of this thought is evidence that certain problem cases in our understanding of language are amenable to an approach in which pragmatic considerations are given full emphasis.

I do this in this paper by giving a brief sketch of part of a theory I call Perfect Speaker Theory (PST), which, despite its grand name, is a modest attempt to provide a methodological perspective from which to comment on some familiar debates in the philosophy of language and mind. For present purposes, it is the simple and natural way that PST deals with certain well-worn problems in these debates that prompts the suggestion that pragmatic considerations should be treated as irreducibly central to meaning.

At this juncture I should state that by 'pragmatics' I mean what, over and above what is standardly assigned to theories of structure, reference, sense and satisfaction conditions, relates to the use of expressions of the language on given occasions, where what the speaker intends to convey, and the means the speaker employs in conveying it on the occasion, are especially in focus. I want to say that among pragmatic considerations questions of force and point (for my purposes especially point) are central. These facets of linguistic behaviour include, essentially, considerations about interaction between speakers and their audiences.

I found, in working at these ideas, that they afford independent reason for agreeing with some of what Grice says. I can only claim some agreement, because Grice is professedly inclined to accept two things which PST is not, namely, the distinctness of meaning and use, and the privileged role of a notion of truth in any account of natural discourse. PST's commitment to arguing for the dominance of pragmatic considerations is allied to a view that notions of truth (which in my view is not one insubstantial thing but a number of different substantial things for which the expression 'truth' is a homophonic dummy) play roles of a different kind (I argue for this view elsewhere.) A major tension between PST and Grice's theory is that because the former has it that the crux in meaning is point, which is to be explained in terms of speakers' intentions to mean something on an occasion, conventional meaning is to be characterised as the dry residue of speakers' meanings, agreed in the language community under constraints of publicity and stability (another view I argue for elsewhere). Now Grice has given considerable attention to showing that speakers' intentions and conversational maxims are insufficient for a full account. PST aims to suggest that from these resources a satisfying account of meaning can be brought into view without having to import considerations from outside the reach of pragmatics, suitably construed.

There are many debates about language which can be exploited to show how PST pushes us in the indicated direction, but I will advert to just four very familiar examples, and a fifth less familiar one. The problem about using familiar examples is that the literature on them is large, so with no space to review it I have to be unceremonious; and also, everyone has a favourite view about these cases, so doubtless what follows will seem too swift. The cases are (1) the proper understanding of natural language analogues of logical constants, (2) presupposition-failures for certain uses of verbs of doing and trying, where the appropriateness of certain locutions comes into question because some implied condition for their appropriate utterance fails to obtain; (3) some questions about referential uses of definite descriptions, and (4) a certain application of the Twin Earth story. This looks like a heterogeneous collection, but I submit it as a virtue of PST that it identifies some patterns in the aetiology of problems arising in relation to them.

And finally (5) I suggest that PST provides a way of approaching the question whether, for any given natural language, there is something that counts as 'the language', or whether a natural language is a vaguely bounded family of idiolects, perhaps as numerous as their speakers. But first it is necessary to define a Perfect Speaker (PS), and to explain why he is so called. A PS of his language is one who so uses it that whenever he makes an assertion (and mutatis mutandis for other kinds of utterance) he:

(1) expresses his intended meaning as fully as, if not more fully than, his audience needs in the circumstances;

(2) expresses his intended meaning as exactly as, if not more exactly than, his audience etc; and

(3) is as epistemically cautious as the circumstances do or might require, if not more so, with respect to the claims made or presupposed in or by what he says.

Together these conditions make the PS determining or overdetermining in respect of the point and epistemology of his utterances. Note that they do not cover the same ground as Grice's maxims – although there is some overlap – and they do not pretend to be regulative for ordinary communication. Together they make a PS of his language one who is in practice absolved a certain duty, namely, the responsible speaker's duty to stand ready to clarify, qualify or defend what he says if called upon to do so by hearers who have and are exercising anything recognisable as normal competence with the language in use. The PS is absolved the duty of restatement because he is by definition never in a position to have to fulfil it, as long as his audience is as specified.

The PS is of course an idealisation, and one which immediately prompts questions. So it is important to be clear about what kind of idealisation he is. Let us distinguish him from two other possible kinds of idealised speaker whom I shall call the Ideal Speaker and the Optimal Ordinary Speaker. (Strictly, it is only the former who is an idealisation, because there might in fact be Optimal Ordinary Speakers). The Ideal Speaker (IS) satisfies the first two conditions, but for a third has 'is omniscient'. Thus the IS is god or relevantly godlike. One immediate difference between an IS and a PS as I define the latter is that the language of an omniscient speaker would have to be apt for the expression of everything, whatever that means, whereas the language of a speaker who suffers the finitary predicament (is finite in knowledge and powers) cannot be guaranteed to be apt for the expression of everything. So not only does a PS differ from an IS in being subject to finitary constraints, but his language (which is an ordinary natural language after all) carries the mark of that finitude also. This might make it seem that the PS should not be so described, but rather that he should be thought of as an Optimal Ordinary Speaker (OOS), i.e., an ordinary speaker who simply is as careful and precise as he can be, and that is all. We might be trying to be such speakers when we do philosophy or law; so OOS—or 'optimal speaking' by ordinary speakers – might be relatively commonplace. But there are significant differences between OOSs so considered, and the PS I require for my model. One is that an OOS, an ordinary speaker doing his best, must be allowed at times to be in states where the beliefs and intentions that determine the content of what he says are not transparent to himself. This means that an OOS may at times fail to satisfy (1) and (2), and at times all of (1) to (3). But by stipulation a PS is one who always satisfies all three. So his meaning, intentions and

beliefs are transparent to himself. Another and consequent reason is that an OOS has the duty of restatement, i.e., he would fail to be doing his best if he did not stand ready to clarify or defend what he said if asked by linguistically competent hearers to do so. But his obligations here are a function of the defeasibility of his attempts to do his best as a speaker. By my definition a PS is one to whom this duty is inapplicable for the reasons given.

So a PS is better than optimal; and he therefore comes between an IS and an OOS, although he is somewhat closer to the latter than the former. A second stab at a definition adds these comments to conditions (1) to (3).

It should be immediately added, if it is not already clear, that the PS is conceived as one whose interlocutors are always ordinarily competent speakers of the same language. He is not a citizen of a kingdom of PSs.

Some comments are needed on the conditions. The third, that the PS is epistemically cautious in the way described, suggests that he is governed by an ethics of epistemic caution which at least imposes an obligation not just to be sensitive, but to articulate or to be ready to articulate sensitivity, to the possibility of epistemic defeat of any claim presupposed or made on its occasion. This condition is vital to resolving certain problems, one of which I consider. But it is not a requirement that a PS should be regressively hedging about his claims, making explicit such protases to his remarks as 'if there is a world at all...' and the like. A PS is not by stipulation absurd, only pedantic. But it is a corollary of this duty of (so to speak) maximal epistemic caution that the PS be as well-informed as is required to fulfil conditions (1) to (2). This is not a demand that the PS be an IS, i.e., omniscient: it is rather that his use of the expressions of the language should be conformable to what I shall later describe as a Best Dictionary for the Language—that is, one which makes use of the best current theories of what use of the language's expressions constitutes talk of.

A comment on the PS himself is prompted by the first two conditions and his correlative freedom from the duty of restatement, and this comment in turn underlines a substantial point often made about a feature of language. In line with (1) to (2), the PS never indulges in metaphor or irony if there is a risk of misunderstanding—which even with normally competent hearers there often enough is; and he never indulges in ambiguity. He is, in short, in danger of being a bore. But his potential boringness is interesting in this respect: that it reveals one of the constraints imposed on him by language (there are plenty of others). For whereas by (1) and (2) he shuns ambiguity, he cannot avoid vagueness, except by stipulative means, which he will not anyway normally wish to resort to. He is of course minimally vague: a major use of certain resources in the language, namely adjectives and adverbs, is the reduction of vagueness, allowing both ordinary speakers and a PS to be maximally specific; maximally but not completely, because vagueness is a built-in feature of language upon which a good deal of its utility turns. But this does not generate an inconsistency with condition (2), which is that the PS expresses exactly his intended meaning, for one can exactly mean to say what cannot be expressed otherwise than by use of a vague expression. For example: suppose the PS says, 'X is bald'. That can be exactly what he intends to convey, independently of questions about the degree to which X has less hair on his head than Y to whom the PS does not apply this predicate. It is of course possible for a non-arbitrary stipulation to be made with respect to some vague expression which precisifies it relative to a certain purpose. One can say that a person n % of whose scalp has a covering of fewer than n hairs per some measure, is bald. Suppose registered bald people by law have to wear a white hat on sunny days. Then the legal instrument which enacts this law would have to be precise: trichological police would need a definition to work to. Just such precisifications in fact obtain in registration of blind people to whom welfare benefits are due.

But as we see, a PS would normally neither need nor desire to go for precisification of vague expressions, however non-arbitrary relative to a purpose: for their vagueness is exactly what from time to time he needs. The constraint they impose is not a limitation.

These comments together give us a second stab at a definition of a PS which will suffice for the present. I shall sometimes speak of the three conditions as rules which bind the PS in his practice. Another important feature of PS-hood, a corollary of the third condition, emerges as the model is applied. I turn to that in a moment. First I will just mention a reason for using the adjective 'perfect' in my label. It is because it is

informatively symmetrical with the use of the adjective made by Russell and others in their attempts to define a 'perfect language'. That was a programme aimed at specifying the underlying logical structure of natural language. In addition to the assumption that there is such a thing, there was a further, at least at the outset: that it admits of a uniquely correct representation. The ambition was to set out in algebraic description of logical form something which, like Leibniz's longed-for universal characteristic, would completely and unambiguously represent what is said by any natural languages. And this in turn was held to have exciting metaphysical potential, since the idea was that what there is can be read off from what the language says. Well: all I need say is that I propose we substitute idealisation of the speaker for idealisation of the language: instead of looking for the perfect language, let us try to describe a PS of ordinary language and see where it takes us. I think such a task justifies the assumption upon which it rests, namely, that it is not the language that says things, but its speakers.

I turn now to apply PST to cases. The suggestion is this: application of the theory reveals a certain pattern in what generates the problems, namely, a falling-short, either because of the usual vicissitudes of discourse in an imperfect world, or artificially by hypothesis, of a PS's conditions; and it therefore suggests a solution to them, which is to make explicit appeal to the point of utterance and/or to identify the epistemic deficit requiring remedy.

The first case I consider – and I consider all of them briefly – is that of the natural language analogues of the logical constants, classically interpreted. I assume familiarity with the standard examples, and just register them here. On a certain natural view, one central use of 'and' in English is to convey temporal succession, one central use of 'or' conveys the speaker's ignorance of the truth-value of the disjuncts, and one central use of conditional statements is to assert that accepting an antecedent is a ground for accepting its consequent, and normally that there is anyway uncertainty about whether the situation denoted or described by the antecedent obtains.

The difficulties felt about these natural uses of the natural language operators arise from their divergence from their formal analogues. Their content exceeds that of their formal analogues, and they are therefore more fruitful in implications. Suppose it is given that a particular disjunction is true. We thus know that at least one of the disjuncts is true. But we are not in a position to infer something of a different order, such as that an asserter of this disjunction is either ignorant of the disjuncts' truth-values, or might be dissembling, guessing, joking or some such. The additional content is provided by pragmatic considerations: those specified in terms of the speaker's intentions and certain contextual features.

Faced with questions about these divergences, our inclination is to look for mappings. Strawson, for example, held that the acceptability or truth of a conditional rests on whether acceptance of the antecedent is a ground for accepting the consequent, but that this is not sufficient for a conditional's truth or acceptability, for which the truth of the associated material conditional, entailed by the natural language conditional, is also required. From this it follows

that whenever an associated material conditional is false, its associated natural language conditional is false also, and that there is a problem about what to say of natural language conditionals whose associated material conditionals are true but in which it is obviously the case that the antecedent states no ground for acceptance of the consequent, for some such reason as, say, sheer irrelevance of one to the other.

The PST offers the following way with the matter. For a PS, the governing question concerns what the point is of choosing to say things one way rather than another. Consider the conjunction and disjunction cases. Let us accept that the most natural construction to place on someone's saying 'he jumped into the swimming pool and put on his trunks' is that the circumstances were such that the man referred to put on his trunks in the water. And let us likewise accept that if someone says 'he's either in Austria or Switzerland' that the speaker does not know in which of the two countries he is to be found. In the first case, if the reverse temporal order were meant, the situation would have to be regarded as non-standard, or the speaker as not fully competent. Either way, the point of the utterance is at risk of being obscured. In the second case, if the natural implication is false, a different explanation of the speaker's point offers: the speaker dissembles, jokes, or something like. But both in the account to be given of the natural thing to say about what the speaker intends to convey, and in the account to be given of the ways things can go wrong or differently, the key is the point of the utterance. A PS by his rules seeks to convey exactly his point, and so if he chooses to say 'he jumped into the swimming pool and put on his trunks' rather than 'he put on his trunks and jumped into the swimming pool,' or 'he put on his trunks while jumping into the swimming pool,' or any other variation of the temporal relation between the jumping-in and putting-on events, then that is what he means to say. So a PS means the standard implications of saying something a certain way to be present in saying it that way.

The problems (or supposed problems, as this intuitive move suggests) of the divergence between the natural language and formal cases are on this view purely an artefact of taking too seriously the syntactic similarity of 'and' to ampersand, 'or' to the 'if-then' to arrow. The meanings of the second in each pair are determined by their truth-tables, that is, by purely material considerations. But those of the first in each pair cannot be explained without reference to the pragmatic features which determine a speaker's choice of just this rather than that way of saying something: which is to say, his point. This is revealed simply enough by answers to the question, 'why would a PS say it that particular way?'

So the observation is that application of conditions (1) and (2) (the full and exact expression conditions) make choice of a way of saying something wholly deliberate, so that every natural construction to be placed on saying it that way is what saying it that way is intended to convey, so if – and we accept they do – natural language conjunctions, disjunctions and conditionals carry the implicatures identified, then if a PS did

not mean them to be taken in what he says, he would say it otherwise. (He would, say, qualify appropriately).

The same analysis in terms of point reduces the problem in the second case I consider, again briefly, namely that of what might be called condition-failure, although it turns out, as I shall claim, that this is a misnomer. The debate in Austin, Searle and Grice is as familiar as the foregoing, and needs as little setting up here as it did. The problem concerns what account should be given of what is implied by saying that—to take two familiar examples – someone omitted to do something or that someone tried to do something. Using those familiar examples from Grice, it is natural to take it that the sentence 'A omitted to turn on the light' implies that A might have been expected to do so, and 'A tried to turn on the light' that some difficulty afflicted A in this matter, and indeed, most likely, that he therefore failed to turn on the light. So it is natural to take as a condition of the appropriateness of saying 'A omitted to turn on the light', rather than saying he did not turn on the light, that in some way there was an obligation on him to do so which he did not fulfil, or that he had intended to do so but forgot, or chose not to, or the like. If none of these things are the case, then it is inappropriate to use the verb 'omit'. If therefore someone says that A omitted to turn on the light, an audience is thereby licensed to make the appropriate inference. *Mutatis mutandis* the same applies to the 'tried' case. Our inclination would be to put the point by saying that it is a condition of the use of 'omitted' and 'tried' in these cases that the implications hold.

The problem as standardly conceived relates to the truth-evaluation of sentences for which such conditions fail. Suppose it is false that there was an expectation that A would turn on the light, and A did not turn on the light. Then is one to say that the sentence 'A omitted to turn on the light' is false or truth-valueless? Well, consider the PS. According to his rules his choice of ways of saying what he does is governed by the requirement to make his point explicit. If it is standard to use 'omit' and 'try' in given cases because the point of doing so is given by these conditions for their appropriateness, then only if they obtain would the PS use them. This obstructs the alleged problem by stressing the directionality of the dependence: to use 'omit' or 'try' is to say that these conditions obtain. Here I shall restrict the notion of 'saying that' to: expressly conveys. ('Saying that' is not coterminous with asserts – for a liar 'says that' but does not assert – and we must be reasonable about what the utterer intends by way of implication of what he says: he does not consciously have to imply every implication of what he says, given that their number might be unmanageably large). The point of an utterance using these locutions is not given by contraposing, although what the contrapositive says is true (that if not-X then not-A omitted): so in determining choice of expression a PS, as the speaker who is expressly mindful of what one is taken to be saying in saying something in this particular way, uses the expressions in question only if he means to say that the conditions obtain. To say that 'A omitted to do X' is to say that (giving this idiom full weight) there was an expectation that he would. If there is no such expectation, there is at best no point in using 'omit', and at worse it deceives the audience. For the PS, accordingly, the problem cannot arise. I take this to mean that there is no problem and that a moral offers: that the property of evaluation is not truth-value but 'appropriateness', understood as parasitic on point.

The third case is also a straightforward one for the theory: it shows that the appearance of a difference between two kinds of uses of definite descriptions is an artefact of an underlying problem-generator which the theory exposes. Recall the familiar examples: 'Her husband is kind to her', where the man is not her husband, and 'the man drinking champagne is happy tonight', where the man is drinking fizzy water. And we are familiar with Donnellan's view. Now consider the PS in these cases. He intends to say something about someone, and in the example cases what he wishes to say is that someone is kind to her or is happy tonight. In order to fulfil this intention he has to pick out the target of his remark for his audience. The possibility that a description which he believes applies to the referent might get the audience's attention to the target without in fact being true of it is one which falls within the range of epistemic defeats to which the PS is required by his rules to be sensitive and, where doing so is germane, to articulate sensitivity. So he means, and therefore might say, 'The man whom I take to be her husband', 'the man whom I take to be drinking champagne'. Now the descriptions refer by attributing something to the referent by means of which the audience can pick him out. In such practice all descriptions are thus attributive—they attribute to the referents they pick out a property they are believed to have on the basis of which they are successfully identifiable. For this latter purpose, it is sufficient that on the occasion of use there is a presumption shared by speaker and audience that the belief could be true. This is different from Searle's account, given in terms of aspects.

Suppose, however improbably, that there are no other aspects which in the circumstances could direct the audience's attention to the man whom the speaker says is happy tonight. It would remain enough that it is a plausible belief for the speaker and audience to hold about him, and which the speaker can exploit for referential purposes, that he is drinking champagne.

It seems entirely natural to say that what the PS does in these cases is what all speakers are enthymematically doing anyway, as an entirely natural extrapolation of the cases shows: if the audience responds, 'he is not her husband, he is her lover', 'it is not champagne, it is Perrier water', the speaker replies, 'Oh well I thought he or it was such-and-such—but anyway you know what I mean'.

This does not defend in any way Russell's allegedly pure attributive theory, in which the form of the example is represented as 'the unique x to satisfy F, Gs', because the epistemic restraint built into specification of F imports something which is additional to the purely conventional aspect of the other expressions constituting F. We can best see what is at work by reminding ourselves of Kripke's suggestion. For Kripke, the distinction lies between semantic and speaker reference. For some idiolect, the semantic referent of an indexical-free designator is fixed by general intentions of speakers to pick out a given object by its means. Speakers' referents are given by specific intentions on specific occasions to refer to a given object. And again familiarly,

Kripke identifies simple and complex cases; in the former, a speaker's specific intention just is his general intention. In the latter, his specific and general intentions diverge, but as a matter of fact he believes that his specific intention determines the same object as his general intention. He might be right or wrong but still succeed in referring. On this view, D's 'attributive' case is nothing but the simple case, his 'referential' case nothing but the complex case.

On the PS view, however, all cases are to be understood as complex cases. Reflecting on the oddity of Kripke's distinction suggests why they have to be so. It is surely never the case, intentions to deceive apart, that speakers believe that their specific and general referential intentions diverge. Their choice of designating expression is governed by a desire to succeed in making reference. So as a matter of fact speakers are always in Kripke's complex case: they believe that their specific and general intentions coincide. Indeed, the distinction between the intentions is not one that speakers make from inside their referring practice: it is one that retrospectively offers in theoretical reflection on the fact that many beliefs are defeasible.

But the PS builds recognition of this fact of life into his practice. So his is expressly the complex case: his specific intention to refer is effected by articulably cautious attribution of certain properties to the referent by means of expressions whose conventional meaning is apt for the task.

Among the points worth noting here are the following. A familiar contrast between conventional meaning and speaker meaning lies in the offing of those remarks: it is just such a Gricean distinction from which Kripke takes his start, but this way with the outcome hints that any account of conventional meaning is, pace Grice, parasitic upon an account of speaker meaning: that the former is, in a sense worth specifying, the dry residue of agreements forced among speakers seeking success for their intentions to mean. This is a point which lies just off-stage in this paper.

The second is that what gave rise to the debate about a reference-attribute distinction is ellipsis: the thought is that making speakers express too little of what a model speaker—the PS—would say is exactly what generates the problem. And there is nothing philosophically problematic about ellipsis.

But now we need to bring in another feature of PS-hood, implied by the conditions but not so far exploited. This feature comes out in the final case I consider: a certain employment of Putnam's Twin Earth argument.

Putnam challenged two received assumptions, the first being that to know the meaning of an expression is to be in a certain sort of psychological state, and the second being that the meaning of an expression determines its extension. These assumptions entail that psychological states determine the extension of terms. Putnam's twin-earth case purports to show that this cannot be so, giving us his celebrated conclusion that meanings are not in the head. Of course this applies just when the psychological states in question are narrow, that is, understood concordantly with 'methodological solipsism', the view that psychological states supervene on intrinsic states of an individual considered independently of anything besides, in particular without relation to environmental factors causing or being effected by those states. Now on PS theory the crux for meaning is point: explanation of what an utterance means has to make essential reference to what speakers intend to convey. This means that a psychological state privileged by the theory, viz. 'intending to mean p', determines the extensions of expressions used. Putnam's twin-earth considerations seem to block this. Do they? I suggest not. The claim I think we justify by appeal to PS theory is this—and here we introduce the additional feature of PS-hood required by the conditions—: that 'intending to mean p' does indeed determine extension, under the two constraints (a) that the words used are used according to what I shall label 'the Best Dictionary', which reports the agreed, relatively stable senses of expressions employed as tokens of communicative exchange in the linguistic community, and which does so according to the best current theories of the linguistic community, and (b) that the PS's audience of normally competent ordinary speakers will so take it.

We get at this thought by noting something about how the twin-earth case is set up. The first thing to note is an apparently tangential point, namely that rejection of methodological solipsism does not by itself entail some kind of realism about the domain of application of the terms whose meanings are to be understood broadly. For so far all that we have identified as required in addition to a speaker's knowledge of meaning is the existence of other speakers whose interactions with him and one another constrain that knowledge. If this is not enough to serve as the broad context, then we need to find the reason in Putnam's twin-earth argument. Does his argument show that unless there is H<sub>2</sub>O and XYZ out there on earth and twin-earth respectively, Oscar and his twin could not respectively have meant 'water' and 'water'? That of course is not what Putnam sought to conclude: his argument was not aimed at proving the existence of the external world, but at showing that meanings cannot be individuated narrowly. So there are three assumptions at work in the twin earth case. The existence of the different waters—water H<sub>2</sub>O and water XYZ—is assumed, as is the qualitative identity of the twins' narrow states, as in the twins' woeful ignorance of views current about water in their worlds. Thus richly equipped, the argument has no difficulty in delivering its celebrated conclusion. But before we accept these premisses in our welcome of the conclusion, we should enquire whether the twin-earth tale could not be as well or better told, because more parsimoniously told, using just the idea that the meanings of expressions in a language are the agreed dry residue of speakers' meanings.

Any individual's problem is that he does not know everything that all other speaker's jointly know about the meaning of the expressions in his language. This is unsurprising: if some best and latest dictionary pooled a community's knowledge of meanings, it would be a rare individual whose linguistic knowledge came close. Such a dictionary would report the knowledge possessed by the completest speakers—OOSs—including, for example, chemists and hydroengineers. Now let us bring in the PS, who knows everything in the best dictionary. Note that the PS is not

omniscient, like the IS: he just knows the Best Dictionary, which is finite and fallible, but reports at their completest the linguistic community's dry residue agreements over what expressions can be used to do. In knowing the meaning of 'water' the PS therefore knows that, in the latest state of chemical theory on earth, it is stuff of molecular structure H<sub>2</sub>O. And let us also introduce the twin-earth PS, who is by definition equally up to date. Then he knows that 'water' according to twin-earth chemistry is XYZ. In this idealised state of knowledge of meanings— that is, where a PS has at his disposal the linguistic community's best joint knowledge—we find that when he says 'water' he intends to refer to water, that is H<sub>2</sub>O, or in the case of the twin-earth PS, then to water XYZ; and so in either case the PS's grasp of the expression's meaning determines its extension, and the psychological state in which his grasp of the meaning consists is broad. But this is not because it is related, causally or in some other way, to water, but rather to theories about water, because he is speaking in conformity with the Best Dictionary, that is, with the fullest shared knowledge of meanings, in accordance with the best current theories held by the linguistic community.

The trick in Putnam's thought experiment is that the people talking about water are ignorant, in the way people are apt to be, as to the best current theories about the stuff. So we who know something more about H<sub>2</sub>O than, ex hypothesi, they do, can see the point as to what else is needed for them to achieve a successful reference: namely, to intend to refer to just that stuff, and not something that cannot be distinguished from it when one's level of knowledge about it is suitably impoverished. Putnam's thought experiment does its work because it premisses that the speakers on earth and twin-earth should be identically ignorant in respect of what their referential intentions would be if they were ideally, or just more, knowledgeable: that is, if they were, or at least approximated the status of, PSs. Now the third condition in PS-hood was epistemic caution. The suggestion here is that it is a corollary of being epistemically cautious that one be as well informed as one can be for the purposes of satisfying (1)-(2). The PS, in obedience to (3), knows (or carries around and consults) the Best Dictionary. So when he refers, he does not do so under Putnamian epistemic privation of the kind suited to making the twin-earth case plausible.

I conclude now by drawing a couple of morals. The point of the PS model is claimed to be that by applying it to cases like the ones just canvassed, we get perspicuous accounts of what is being said and done, and they show that the problematic character of the canvassed cases is an artefact of failing to give full weight to considerations made salient by the pragmatics of Perfect Speaking. What helps with the problem cases is appeal to considerations of point and epistemic aptness; the PS is 'perfect' in his practice with respect to both—and in being so is such that the problems do not arise for him. This, my inference is, suggests that we should look to beliefs and intentions for the basis of a general account of meaning.

In the first two cases the problems were generated by incomplete determination of point in the formulation adopted by a putative ordinary speaker. In the second two cases the problems were generated by epistemic underdetermination in the cases; the putative ordinary speakers failed in achieving their intended targets of utterance precisely because of it.

A PS perspective on the cases brings a salutary reminder to our attention. It is that no non-idly employed sentence of natural language exists outside a pragmatic frame. For every non-idle use of a sentence the particularities—how things are in respect of the utterer, his intentions, his audience, the current state of the language, and the circumstances of its use – determine the meaning of what is said on that occasion.

It is these two thoughts together that suggest the third – at this juncture I do not claim they do more – namely that the conventional meanings of expressions in natural language are the precipitates of the linguistic community's tacit agreements to place the use made of certain signs under publicity and (relative) stability constraints, so that the ends of communication can be realised. These agreements are the conventions which dictionaries report. Any account of meaning so understood would have on this view to make essential reference to the pragmatic considerations—and central among them, point—which figure thus in its genesis.

Now it is widely held that—and I quote Searle—"meaning is more than a matter of intention, it is also a matter of convention" and in illustration he quotes Wittgenstein's remark in *Philosophical Investigations*, 'Say "it's cold here" and mean "it's warm here"'. There is no inconsistency between agreeing with this and accepting what the PS theory says is central. For one thing, conventional meanings are, so the theory seeks to suggest, the precipitate of intentions anyway: what an expression means is what, in effect, it is intended to mean by the linguistic community. For another, nothing in PS theory excludes the obvious, which is that the point of a speaker's utterances, even a PS's utterances, cannot be individuated independently of facts about the conventional meanings of the expressions he uses, the context, the speaker's attitudes, and whatever other pragmatic features are required besides. But what it does do is to say that point figures centrally.

By way of conclusion, and even more briefly than in the familiar problem cases mentioned, I suggest that among the other things PST can offer is a simple solution to the question whether there is such a thing as 'the language' for any natural language. The thought that there is such a thing is held by some to play the role that naive realism does in theories of perception; it is the dumbo view, to be replaced by more sophisticated views such as that there are as many personal (so to speak) paroles as speakers, and that what we too loosely call communication is in fact a form of translation or, more accurately, interpretation. A natural desire to respond that at least a notional 'the language' is required to provide a norm—which among other things can be invoked to explain such phenomena as (say) the differing divergences of idiolects from majority intelligibility—might be met with the riposte that such norms are in fact constituted not by something which is genuinely 'the language' but by the family of idiolects of an historically favoured class of speakers (the currently rich, the currently powerful, the people currently in charge of the culture. There could be—after the revolution, say, there might be – quite different folk in these roles, speaking a different family of idiolects). And so the debate might proceed. But one thing that would help to give it shape would be to offer

something that could count as a criterion of identity for 'the language' if there is such a thing. This is where PST comes in: for such a criterion is offered by saying: the language is what the PS speaks. And this genuinely does offer a normative conception across which mappings must fall if the very idea of an idiolect of some language is itself to make sense.