

tent”) to phenomena that we do not understand. Instead, we should seek to construct “bodies of doctrine” just as Newton did, “deferring” problems that lie beyond reach until we understand enough to raise them seriously (chemist Joseph Black; see reply to Poland). That’s the course I and others are trying to follow. I don’t see that Rey has raised any problems about it, apart from the huge gaps of understanding that are already familiar.

Rey ends by posing a dilemma: If my approach “requires intentionality and a CRTT,” then it’s my responsibility to say how I understand such “intentional idioms” as “representations,” “cognizing,” “information.” If it doesn’t require intentionality and the CRTT, then I have to say why we should regard grammatical competence “as genuinely part of the *mind*, and not merely as some non-mental system that a mind might happen to exploit.” On the second horn, I cannot comment. I have no idea what it takes to be “genuinely part of the *mind*,” or what he thinks “the mind” is, or how “a mind” might “exploit” a non-mental system. As for the first horn, I agree – in fact insist – that it’s my task to say what I mean by the terms I introduce, though not to introduce and explain intentionality and CRTT (in the sense that Rey may have in mind). As to how well I’ve done that, others can judge. The questions are not raised here, as far as I can determine.

Reply to Ludlow

Peter Ludlow opens by observing that I am not “hostile towards semantics.” That’s correct. Most of my work on language, from the outset, has been an effort to explore the form and meaning of expressions and the principles that determine them, core problems of semantics. The numbered examples of the comments are typical illustrations. It’s true that I prefer to call this work “syntax,” and think the same term might well be used for other investigations that keep to symbolic systems that are internal to the mind-brain, the faculty of language FL or others.

Virtually all contemporary reflections about language and mind take the expressions that have sound and meaning to be in a “public language” that is external to the mind-brain, “a kind of social object the structure of which has been established by convention,” say by the pronouncements of the French Academy (Ludlow). It is misleading to say that I regard “the French language, so understood, as devoid of scientific interest.” One reason is that the point of view is not specifically mine; it is standard in the empirical study of language. A more important reason is that the French Academy has had little to say about the French language in anyone’s sense of the term, informal or technical. It has not even attempted to describe the form and interpretation of expressions of French. It has not been concerned, for example, with the properties of French expressions (including very simple ones) discovered and studied in Richard

Kayne 1975, the work that laid the basis for much of contemporary inquiry into the syntax and semantics of the Romance languages, with far broader impact.

The French Academy issues pronouncements that can be understood by people who already know French; that is, people with a functioning FL that has attained a state that underlies behavior that one might choose (under obscure and variable circumstances) to call “speaking French.” The same is true of the *Oxford English Dictionary*, or the most elaborate traditional grammars. This is no criticism. Discovery of the elementary properties of French, or the meanings of words of English, is a challenging task. The French Academy and *OED* have different concerns.

The problem with the “E-language perspective” that Ludlow discusses is not simply the recourse (in some versions) to such factors as social convention and colors on maps. Rather, it does not address the basic properties of language, even the most rudimentary of them: that a language has an infinite variety of expressions, each a form–meaning pair. To play some role in theoretical explanation, an E-language approach has to provide us with at least some indication of the membership of this infinite class. The “public language” approach does not. Nor do other versions of an E-language perspective; Quine’s, for example, which takes a language to be an infinite class of well-formed objects. One problem is that natural language appears to lack any relevant notion of well-formedness (Chomsky 1955, and much subsequent work); variations along many dimensions have been a major topic of empirical inquiry, the study of weak and strong islands and their syntactic–semantic properties, for example. A more fundamental problem is that any such approach can be understood only insofar as it provides a finitary account of some infinite class of objects that it has in mind, thus shifting the inquiry to a course that Quine condemned as sheer “folly.” Some have used the term “E-language” to refer to actual utterances or texts. People are free to invent terminology as they like, but that was not the meaning I gave to the term when I introduced it; I-language approaches do not differ from others in their concern with such materials.

It is for such reasons that I doubt that there is an E-language perspective, and believe that even those who deny the fact are tacitly assuming the existence of I-languages, internal states of FL. It is only in these terms, as far as I can see, that one can make some sense of work that purports to adopt an E-language perspective, and evaluate its contributions to understanding language and mind.

Suppose we now explicitly adopt the I-language perspective (see Ludlow’s note 5). Suppose Jones’s I-language generates the expressions (1) ((i) is Ludlow’s example):

- (1) (i) the flaw in the argument is obvious, but it escaped John’s attention
- (ii) the fly in the bottle is trapped, but it will escape

(iii) John bought a book for Mary, but she already had it, so after reading it, he shredded it

Each of these expressions has a sound and a meaning; we can think of each E as a pair $\langle \text{PHON}(E), \text{SEM}(E) \rangle$ (in the terminology of Horwich and reply). Jones makes use of this information to produce and interpret utterances.

Suppose we want to account for such facts. Here's a simple way. The words of Jones's I-language map onto language-external objects (sounds, things, properties); the sound and meaning of utterances are composed of these in elementary ways. Thus "bottle" in (ii) maps onto a sound; the pronoun "it" maps onto the flaw in (i), the fly in (ii), and the book in (iii), where flaws are the kinds of things-in-the-world that can be in arguments the way flies are in bottles, a book is something that John can buy, read, and shred while Mary still has it. And so on, for every word we choose to look at carefully.

There is a sense in which sounds, flaws, books, . . . "are clearly not logically absurd entities, nor . . . particularly odd entities" (Ludlow). Sounds, for example, are perfectly robust things. We have no problem assigning sounds to the expressions of (1), or none of the above, in a vast range of normal cases. If we are satisfied with this result, we can avoid the hard problems of experimental phonetics.

Proceeding, we can quantify over flaws and average guys; also sounds and books. When we use a nominal phrase X, we can usually say, intelligibly, that the X's are the things we are talking about: flaws, books, the darkening sky, Joe Sixpack, the bank that was burned to the ground after rejecting the loan application, the threat of global warming, a biochemical process that is tragic (life), the purpose of America (the title of a book by one of the founders of realist international relations theory), . . . Again, if satisfied with this result, we can avoid the hard problems of the study of meaning, reference, and language use generally.

If we adopt this course, are we putting forth "substantive metaphysical claims," claims about the constituents of the world? Surely not in terms of the enterprise of natural science; on that there is no dispute.

Let's turn to ethnoscience, which studies "folk science," seeking to determine how Jones and others conceive of the elements of which the world is constituted and their functioning (see reply to Egan). Suppose that ethnoscientific inquiry attributes to Jones internal non-linguistic concepts that are one-one associated with nominal phrases of his I-language, and concludes that these are the constituents of Jones's commonsense understanding of the world. That is a metaphysical claim, on the order of the claim that an insect has a mechanism of path integration, or that Jones's I-language has a phonological system that enables him to distinguish /r/ from /l/ (unlike his Japanese friend). The relation of nominal expressions to these concepts is Ludlow's relation R^0 , a syntactic relation that he dismisses.

Suppose ethnoscientific inquiry concludes that Jones also believes that the mind-independent world is constituted of entities that correspond to mind-internal words (or maybe their images under R^0). I presume Ludlow would be skeptical about these conclusions. At least, I would. I doubt that people think that among the constituents of the world are entities that are simultaneously abstract and concrete (like books and banks), or that have the amalgam of properties we discover when we explore the meanings of even the simplest words (“river,” “person,” “city,” etc.). I suspect that serious ethnoscience will find that intuitive commonsense understanding of the world looks quite unlike that.

Nonetheless, at some level, we can accept the idea that when we use the words “book” or “river,” we are talking about books or rivers, and can happily say that there are books and rivers; similarly for virtually every nominal expression, and, with elaborations, for other expressions. We can accept all this at the level at which we abandon curiosity about language and the mind, about human action and its roots and properties. Though sounds are perfectly robust and simple, much more so than books and rivers and flaws, scientists concerned with the sound aspects of language have not been satisfied with such accounts, any more than Kayne was satisfied with the hints provided by the French Academy or traditional grammars and dictionaries. They have sought to discover the internal entities PHON(E) and to determine how they relate to the kinds of mind-external entities that are studied in the sciences, investigating a relation between internal and external events mediated by sensorimotor systems; a phonological analogue to R^0 has never been considered. They have not been satisfied with the informal notion of thing in the world that is fine for ordinary life but useless if we hope to understand the world – in this case, the robust notion *sound of E*, a fine thing in the world, in informal usage.

Have articulatory and acoustic phonetics followed the right course, or should they just have stopped with the observation that E has a sound, as we all agree? That depends on one’s interests. It is possible to disguise lack of interest with technical terminology. In this case, we could invent the term “phonetic value,” taking the phonetic value of the expression E to be its sound. Similarly, we can say that the “semantic value of ‘water’ just is water” (Ludlow).

These moves should, I think, be understood as registering lack of interest in the problems. That may be entirely reasonable; no one seeks to study everything. But we should not mislead ourselves into believing that by invoking sounds, flaws, John Doe, attention, escaping, . . ., or rivers, water, cities, books, trees, . . ., and taking them to be related to pronouns and other words by an invented technical notion called “reference,” we have even begun to investigate, let alone to have solved, the problem of how people use language to refer to things in the world, or any other kind of language–world relation. That’s taken for granted in the study of sound, and should be in the study of meaning and referring as well, I think.

Consider a Martian scientist M who wants to study humans the way humans

study insects. Suppose M is interested in questions of the kind just raised, and is informed by his human subjects that there is no real problem: we can account for everything by invoking the metaphysical thesis that among the things of the world are sounds, flaws, books, . . . , saying that words like “it” *refer* to these things, and so on. But M wants to know more: What are sounds, flaws, books, . . . , and what is the relation “refer” that has these curious properties? Suppose we assure M again that there’s no problem: we all understand very well what sounds, flaws, books, . . . are and how to use “it” to refer to them. True enough, but no help to M, who wants to comprehend what it is that we understand, how we achieved this state of mind, and how our linguistic states relate to the outside world (and from another point of view, the inner world of language-external thought). M will, in short, adopt the stance that humans adopt when studying insects, or the sound side of language; and should adopt if they are interested in meaning and referring as well. M is just us when we are engaged in naturalistic inquiry; that is, seeking to understand something about the world (including the special case of how we engage in self-conscious naturalistic inquiry or unreflective exercises of commonsense understanding, or any amalgam of these enterprises).

Consider Davidson’s analysis of the event structure of propositions. I share Ludlow’s conclusion that these ideas have been extremely productive, a major contribution to our understanding of language (see Pietroski’s essay). But the work remains syntactic in the broad sense, and falls short of significant metaphysical claims, until we are given some further explanation of what events are and of the notion “existence of events,” an explanation that responds to the concerns of the Martian M: that is, the kind of explanation we rightly demand when we are told of “the existence of sounds.” What is needed is a language-independent account of the postulated “events,” either in the domain of ethnoscience or other branches of naturalistic inquiry. We speak freely of events in ordinary life, but that’s no help to M. What counts as an event in the relevant sense is not at all obvious, surely no more so than what counts as a book, or a river, or a nameable thing.

In the same connection, consider Ludlow’s notion *I-substance* (which he introduces with reservations). I-substances are “the stuff that we are intuitively talking about when we use language (the intuitive referent of ‘water’).” That’s fine, as long as we do not conclude without evidence that our commonsense conception of the way the world is constituted and functions actually employs I-substances, pre-empting the conclusions of ethnoscience and contrary to what even superficial reflection reveals, I believe. The I-substance is what it appears we are “talking about when we use language” only if we misunderstand our use of language, ignoring the warnings of the later Wittgenstein, Gilbert Ryle, and others. The problem of “matching up right” doesn’t arise, if we avoid these pitfalls.

There is nothing original about these observations. They are at least as old as Aristotle. He was not satisfied simply to say that a house is a robust entity (just

like a sound or a flaw), but wanted to know what it is. We can “define a house,” he concluded, “as stones, bricks and timbers,” in terms of material constitution, or as “a receptacle to shelter chattels and living beings,” in terms of function and design; and we should combine both parts of the definition, integrating matter and form, since the “essence of a house” involves the “purpose and end” of the material constitution (*Metaphysics*, bk VIII, ch. 3, 1043^a; *De anima*, bk I, ch. 1, 403^b). That is a step towards understanding, though only very partial. Houses have far more intricate properties, as does every “object,” as we discover when we go beyond casual inspection (Chomsky 1975a, 2000, among others). The most elaborate dictionaries, monolingual or pedagogical, never give a hint of these properties, quite rightly; even if they had been noticed, spelling them out would only confuse the user, whose knowledge of these facts comes from “the original hand of nature,” in Hume’s apt phrase.

We move a step beyond when we reach the seventeenth century and its radical revision of Aristotelian science. One aspect was to shift the burden of explanation of topics such as these from the postulated structure of the world to the structure of the mind. What we can know is determined by the “modes of conception in the understanding,” in Cudworth’s phrase. That led to illuminating work in the theory of meaning by neoplatonists, empiricists from Hobbes to Hume, and others, influencing Kant, with later resonance until today (see reply to Gopnik for some references). The Martian scientist would be well advised to pursue a similar course, adopting the ethological perspective that seeks to discover the organism’s *Umwelt*, its particular mode of interpreting the world.

For non-human animals, so it is alleged, internal computational systems are “representational” in something like the sense of technical referential semantics. Introducing a series of experimental papers, Gallistel (1990) argues that representations play a key role in animal behavior and cognition; here “representation” is understood as isomorphism, a one–one relation between mind/brain processes and “an aspect of the environment to which these processes adapt the animal’s behavior,” as when an ant “represents” the corpse of a conspecific by odor. For humans, this notion is completely inappropriate, as we see by inspecting even the simplest words, criteria of individuation, the basic mechanisms of intended referential dependence, and other elementary properties of language and its use, even the concept “nameable thing” (as noted in comments that Ludlow quotes).

Pursuing this course, we can learn a lot about humans, and about how the expressions of their I-languages relate to language-external entities – sounds and things in informal usage. But we do not arrive at “substantive metaphysical theses,” except about the nature of the human mind and the “modes of conception in the understanding” that enter into constructing the *Umwelt* in which we live and act. Non-human representational systems might yield directly a kind of mind-independent “metaphysical thesis” if Gallistel is correct, but not human language, which does not seem to be a “representational system” in anything like this sense, in either its sound or meaning aspects.

The decision to be satisfied with informal talk about sounds and meanings and things in the case of (1), or “house,” or any other aspect of language not only cuts off inquiry into important questions, but leaves us with misleading conclusions. Take “flaw” and “fly.” They function in a similar way in (1i, ii), but not in other expressions. Compare (2)–(4):

- (2) (i) there is a fly in the bottle
- (ii) there is a flaw in the argument
- (3) (i) there is believed to be a fly in the bottle
- (ii) there is believed to be a flaw in the argument
- (4) (i) there is a fly believed to be in the bottle
- (ii) there is a flaw believed to be in the argument

4(ii) is deviant, unlike the others. Such constructions as (4) have existential import in some manner beyond (2), (3). The observation generalizes, and falls within an explanatory framework with a variety of consequences, and interesting open problems (Chomsky 2001). Many such aspects of the meaning and structure of expressions, and their use in talking about the world, are likely to be missed if we are satisfied to say that we understand the expressions of our language and need inquire no further, just as the discoveries of articulatory and acoustic phonetics would have been lost if investigators had been satisfied with the observation that expressions have sounds, perfectly robust entities.

For these reasons, I do not entirely agree with Ludlow’s critique of model-theoretic semantics. The basic problem seems to me to go beyond the difficulty of identifying the intended model. It has to do with the models themselves. What are the individuals in the model? Are they flaws, books, houses, water, rivers, nameable things, sounds, . . . ? If so, we still have to face the problems of how language is used to refer. If we posit semantic values in the manner suggested, we face the same problems. Either we are positing some mind-internal domain of obscure “individuals” that are related to expressions by something like Ludlow’s R^0 , and we face the problems of justification that arise when we posit other kinds of syntactic objects (phonology, phrase structure, etc.); or we are choosing to content ourselves with informal talk that would not answer the questions of the Martian scientist or ourselves, as scientists, though as subjects of inquiry we understand this talk very well, just as bees understand the waggle dance; no help to von Frisch. If questioning stops at this point, the enterprise is a form of syntax – which is, of course, in no way to impugn its contribution to the understanding of I-language and its importance as a preliminary to investigation of the use of language to refer and in other ways, just as phonology is an essential prerequisite to the study of how people produce and interpret the sounds of language.

On efforts to translate expressions such as (1) to “logical form,” I share Ludlow’s skepticism. If the notion “substantive metaphysical thesis” is taken

seriously, I suspect that any such endeavor will soon find itself engaged in something like a Carnapian *Aufbau* project (adding metaphysical claims that Carnap would hardly have welcomed). If the translation project is understood to be in principle part of the empirical sciences, then the operations proposed will at least have to be suitably general and independently justified, like postulated rules of phonology. That may well be possible in such cases as “bad singer”; for (1), it looks hopelessly implausible for reasons Ludlow mentions, and others: e.g., if “flaw” is a modifier like “bad,” then we have to be told how the pronoun “it” can be referentially linked to it in (1i), though never to “bad” in the same way. Problems proliferate rapidly, case by case.

Let’s turn finally to Ludlow’s relations R^i , and the “interpretive puzzle” he poses. R^0 relates a linguistic expression to an entity in an internal symbolic system (call it “S-Mentalese”). A phonetic analogue would relate the expression to an entity in P-Mentalese. S-Mentalese and R^0 have to meet the usual criteria of justification for syntactic theory. Whatever the prospects, I think Ludlow is right to conclude that model-theoretic semantics does not “deliver language/world connections,” though not only for the reasons he mentions. And I think the same is true of approaches that invoke “semantic values” without telling us what these are in some terms that would be helpful to the Martian investigator M: that is, to us, if we hope to understand language and thought. If so, then these approaches do not “explore some form of language/world relations” (Ludlow), though they may purport to do so, misleadingly.

Ludlow’s R^1 is a “direct” relation (expression, “the world”), R^2 a relation (speaker, expression, context, “things in the world”). One could add others, e.g., the “direct” relation $R^3 =$ (expression, things-in-the-world) and the relation $R^4 =$ (speaker, expression, context, the world). R^4 and (under idealization) R^1 seem to me the appropriate choices for the investigation of sound and meaning.

Ludlow quotes my comment that R^2 is “entirely innocent,” because of course “people use expressions to refer to things in particular circumstances.” Probably I was not sufficiently clear in the correspondence he cites. The remark quoted is without metaphysical intent; it is intended simply to repeat the truism that ordinary usage is fine as it is, but shouldn’t be misrepresented through distorting doctrinal prisms and used to draw illegitimate conclusions about the world or our beliefs about the world.

How should we proceed in studying expression–world relations? On the sound side, it is assumed without comment that the approach should be something like R^4 , though actual research introduces radical idealizations, even beyond R^1 , to restrict the terrain to something that can be seriously investigated. Though analogies cannot be pursued too far, it seems to me that a somewhat similar approach is appropriate on the meaning-referring side of language. From lexical semantics to more complex expressions, we can investigate the syntactic interface level SEM and the principles that relate it to the rest of FL. We can then ask

how the information in SEM is used in expressing thought, talking about the world, and other kinds of human action, perhaps mediated by a conceptual-intentional system of the mind/brain (see Horwich and reply). It is an open question whether something will emerge that is similar to the referential semantics devised for formal symbolic systems, and that might have some counterpart in animal cognition, making use of a notion similar to the theoretical term “reference.” I know of no convincing reason to believe so, and would not be surprised if the quest turns out to be as illusory as the analogue is taken to be (without comment) on the sound side. But far too little is understood to say anything with much confidence.

Reply to Horwich

Paul Horwich’s proposals about the faculty of language FL focus on issues of central importance, and also considerable obscurity. I will try to review some of the considerations that seem to me to bear on them.

Horwich notes “a striking convergence of opinion between Chomsky and Wittgenstein regarding the dubious legitimacy of philosophical theories of meaning.” That’s true, and of course no coincidence. When I began to think about these questions seriously as a student, I was much influenced by recently published work of the later Wittgenstein. My skepticism extended to what Horwich calls “referentialism,” though not entirely for the reasons he gives. A more serious problem, I think, is that the best question of use of language – what is the “aspect of the external world” to which a linguistic item is claimed to “refer” – is left unexamined in any serious way, and it is doubtful that this is the right picture to explore at all. See reply to Ludlow. Horwich is right to abandon it, effectively, I think.

Under these and related influences, I assumed from my earliest writings in the mid-1950s a kind of “use theory of meaning,” not in Wittgenstein’s terms but perhaps not inconsistent with them: the internalized language (I-language) generates syntactic objects, each of which “provides a basis for a description of how, in fact, language is used and understood” (Chomsky 1955: 75; quite a substantial basis, it is argued). The generated objects provide information accessed and put to use by performance systems, including the sensorimotor systems and the conceptual-intentional systems (P/A and C/I in Horwich’s notation). In this regard, the theories of sound and meaning are similar. For such reasons, I have an initial sympathy for Horwich’s preference for an approach to meaning in terms of “regularity of use.” But I do not see how to implement the kind of approach he outlines without returning to the one he rejects.

Horwich’s primary concerns are (I) and (II):