

Appears in a special issue of *Croatian Journal of Philosophy* on the philosophy of linguistics.
Volume 6, no. 18, 2006, pp. 401-429.

TESTIMONY AND ILLUSION

Alex Barber

I. Introduction: the threat of sentence eliminativism

Could the existence of sentences be mere illusion? If speakers and hearers were deluded in exactly the same way about the nature and existence of the sentences supposedly being uttered during their conversations, communication between them would be unaffected by their mistake. Being mutual, this mistake would cancel itself out. My aim here is to rule out this sceptical possibility.¹

Section II introduces the topic of sentence eliminativism. Most of what I say, here and throughout the paper, is intended to generalize to include other standard linguistic entities, to borrow Georges Rey's useful label for sentences, nouns, noun phrases, morphemes,

¹ The possibility is actuality according to Georges Rey, on whose statement of the sceptical view I will be focusing. His is the most explicit as well as the most philosophically developed (Rey 2006, Rey 2005, and Rey in preparation). He happily acknowledges debts to the arguments of others, e.g. Saussure 1916, Sapir 1949, Fodor, Bever et al. 1972, Jackendoff 1983, and Chomsky 2000. But these others tend towards idealism rather than eliminativism, or else conflate or waver between the two. In an ongoing exchange (Rey 2005, Devitt 2006: 184-8, Rey 2006, and Devitt [this volume]), Michael Devitt has criticized Rey's first paper for alleged weaknesses both in its case against his own affirmative position and in its sceptical alternative. My objection to Rey is distinct from Devitt's and I will not address that dialectic directly except to argue in section V that Rey's global rejection of relational views in this issue is premature – which, if I am right, would be unsurprising given that his remarks are meant only to herald a longer work in progress.

phonemes, etc. Section III presents a number of affirmative accounts of sentential reality. They differ over what sentences are but share a commitment to their existence. Each has trouble dealing with an objection brought by Rey, or else with ramifications of that objection.

Sections IV and V contain the substance of the paper. In IV I present and criticize Rey's own negative account. In brief, Rey claims that speaker and hearer are locked in a kind of *folie à deux*: sentences do not exist but this fails to disrupt or hinder their conversation because both suffer the same illusion that sentences do exist. Indeed, communication is *more* efficient for their non-existence. Against this I ask whether the transmission of knowledge through verbal or written testimony would be possible if communication involved a *folie à deux*. The problem is an apparent tension between Rey's view and a widely endorsed constraint on knowledge that follows from what epistemologists call the *no-false-lemmas principle*. According to the principle (in its only-if direction), knowledge, including testimonial knowledge, cannot be premised on false assumptions. Of course, Rey may simply have stumbled upon a plausible counterexample to the principle, one from a hitherto unnoticed class of counterexamples that all involve communication through illusion. After all, a seemingly analogous scenario, in which two people send messages to one another by revolving what they think is a bent stick but which is in fact a straight stick delved in water, is perfectly coherent. I argue that, when suitably refined, the no-false-lemmas principle is in fact compatible with the testimonial transfer of knowledge using a "bent" stick but that its compatibility with Rey's position is undermined by a fundamental difference between sticks and linguistic expressions.

A by-product of this discussion of communication through illusion is a different affirmative account of sentential reality. I develop and defend it in V. If I am right, the syntax realized in a speech act, like a speech act's meaning, is an illocutionary property. That is, whether a particular sentence is realized turns on whether the performer of the speech act intends the act to be judged by his or her audience as the realization of that sentence – much

as a speech act's meaning turns on whether the performer intends it to be judged as having that meaning. This position is immune to Rey's criticism of other affirmative stances. Several independent criticisms of the proposal – that intention is too high a level of cognition, for example, or that the identity of the sentence realized depends at least in part on the sound produced and not entirely on the utterer's intention – are less forceful than they first seem.

II. What are sentences and how could there not be any?

Before we can say whether there are any sentences, we need a statement of what sentences are. It is simple enough to give *the form* of such a statement: sentences are abstract entities with realization conditions of some kind – or, in equivalent terminology, they are types, tokened under certain conditions. It would seem, for example, that a sentence was realized when I typed out the opening eighty-one characters of this paragraph, but most sentences have not been and never will be realized. The precise nature of the realization conditions for sentences, left open by this purely formal statement, is controversial. A substantial statement will specify them. Eliminativism about sentences, in the sense I am concerned with here, is the view that sentences are never realized, i.e. the realization conditions, whatever they should turn out to be, are never met.

According to what I will call *basic realism*, sentence realization consists in the occurrence of an acoustic or graphic event of some kind. This simple version of basic realism is ripe for qualification in at least three respects. I mention these only so as to set them aside (at least for now, though the second will re-emerge in section III). The first respect turns on the apparent fact that the same sentence can be realized in different media, including speech, signing, standard writing, Braille, semaphore, and, more controversially, thought. The second complication is that the same sound pattern can realize sentences that are distinct by virtue of having distinct meanings, either because of lexical or structural ambiguities within a language or because speakers of distinct languages could attach distinct meanings to the same acoustic

pattern, e.g. ‘Empedokles liebt’ and ‘Empedocles leaped’ (Davidson 1968: 135). This last example points to a third complication: the existence of a hierarchy of levels of realization, from weak to strong. A sentence may belong merely to a possible language; it may belong to one from a proper subset of possible languages, namely the humanly possible ones; it may belong to an actual human language; and it may have been actually spoken or written. I am concerned here only with what constitutes ascendancy to this fourth level of reality.

Suppose that the basic realist view, refined to accommodate the complications, is correct. And suppose, too, as is plausible, that the physical world does not contain acoustic patterns with the complexity commonly attributed to sentences by linguists. (Acoustic patterns are complex, just not in *that* way.) Eliminativism about sentences, in the sense given above, follows more or less immediately. I elaborate on this simple argument further in the next section, and consider how successful two other affirmative accounts of sentence realization are in defusing the sceptical threat.

III. Affirmative accounts of sentence realization and their shortcomings

Here is Rey’s brief statement of basic realism (or ‘physical tokenism’ as he calls it):

It would seem to be a commonplace that people, when they talk, produce tokens of such things as words, sentences, morphemes, phonemes and phones – I’ll call tokens of all such types “Standard Linguistic Entities (SLEs)”. Part and parcel of this commonplace would be the presumption (“physical tokenism” ...) that these entities can be identified with some sorts of *acoustic* phenomena, e.g., wave patterns in space and time. (Rey 2005: 234)

Michael Devitt and Kim Sterelny, unlike Rey, endorse this ‘commonplace’. In a passage cited by Rey they write:

Tokens are datable, placeable parts of the physical world. The obvious examples of word tokens are inscriptions on a page or sounds in the air. Types, on the other hand, are kinds of tokens. Inscription types and sound types are identifiable by their overt physical characteristics and so we might call them “physical” types.²

And, of course, the view that sentences and other linguistic expressions are realised as acoustic and graphic events is widespread among philosophers of language, not to mention the folk and many linguists.³


So what is wrong with this view? Like most contemporary species of eliminativism with respect to the objects of commonsense ontology, Rey’s originates from science, in this case linguistics. Linguists represent sentences, or inflectional phrases (IPs), as structured entities rather than simple lists of words. For example, sentence (1), normally written out as in (1a), might be more fully represented as in (1b).⁴

² Devitt & Sterelny 1999: 71. Devitt 2006: 184, n. 24, points out that they did not assert a type identity since that would leave them unable to deal with ambiguity. They may therefore be more properly classified as qualified realists (see below). Depending what form it takes, qualified realism can be either a supplemented version of, or a competitor to, basic realism. The classification question does not matter in the long run. Devitt also repeats much the same claim, again stressing multiple realizability: ‘[A] sentence might look like anything: it can be physically realized as a sequence of acoustic vibrations, as gestures in a sign language, as marks on paper, as a sequence of flags, as electric pulses of various kinds, and so on. ...[T]here is nothing incoherent in the idea that sentences could be realized in a neural medium as well.’ (Devitt 2006: 9.4)

³ William Lycan, for example, writes in a popular philosophy of language textbook that ‘[s]ome strings of marks or noises are meaningful sentences’ (Lycan 1996: 4).

⁴ Copied verbatim from Rey 2005. The sentence is the highest IP, a complex entity with the constituent structure indicated. Sentences are often said to be representations themselves, with several levels. The tree represents the

(1) (a) *John seems to Bill to want to help himself*



Nothing stands in the way of a sentence *qua* abstract entity having this structure. But, on the basic realist account of realization, is anything with this structure ever realized? Take, for the sake of focus, the presence of PRO as a constituent of (1), leaving aside other considerations such as PRO's being co-indexed to both *John* and *himself*. PRO is in no obvious sense a component of the acoustic pattern that results when someone gives voice to (1) (or writes the graphic pattern in (1a) for that matter). Indeed, it is hard to see how acoustic events, as such, could non-arbitrarily be assigned constituent structure at all, let alone constituent structure that includes PRO among its elements. Acoustic events are sequential changes in frequency and wavelength in the vibrations of air and eardrum. It would seem to follow that sentences

sentence at, at most, one level of representation. I will ignore the claim about different levels of representation since it adds to the stakes but does not affect the essential nature of Rey's objection.

with the properties linguists attribute to them are never realized and, to that extent, do not exist. Rey summarizes the problem by considering what happens when he speaks:

Now *what actual thing in the world possesses this structure [i.e. (1b) above]*? Well, what is that (purported) utterance? A token of an English sentence. According to [the commonplace voiced in the previous quotation from Rey], this is something in space/time. But does anything I actually produced in *space and time* have the above structure? I think not. (Rey 2005: 245)

Sentences are structurally elaborate; the acoustic patterns associated with speech acts are not, or at least, not in the right way; so the acoustic patterns associated with speech acts do not as such realize sentences.

This reasoning assumes that a realizer must have the same degree of structural elaboration as the realized. The principle seems plausible, at least in the present context. Consider a mouthing of (2) in the context of a normal conversation.

(2) /ˈvɪzɪtɪŋ ˈrɛlətɪvz kæn bi: ˈtaɪəsəm/ ⁵

If it realizes any sentence at all, this mouthing realizes either two distinct sentences – the one about being visited as well as the one about being a visitor – or it realizes just one. If it realizes both then, for all that basic realism tells us, the realization conditions of two distinct sentences are identical and any structural difference between them seems otiose. So the

⁵ I am making the simplifying but false assumption that the phonetic alphabet represents acoustic patterns. Rey is eliminativist about the posits of phonology as well as about sentences (Rey 2005: 246-7), for similar reasons, but that is a distraction in the present context.

practice of assigning structure to sentences, and individuating them thereby, seems to presuppose that a particular mouthing of (2) realizes just one of these sentences. Which one is realized must reflect, somehow, the difference in structure between the two sentences. Exactly how this difference is reflected is still an open question. The argument of the preceding paragraph shows merely that the discriminating factor cannot be a feature internal to the event's acoustic pattern.

Recollection is overdue of an obvious fact: assigning elaborate structure to sentences is supposed to render explicit something buried deep within the mind of the speaker, not readily discernible in vibrations in the air. Given this, it is hardly surprising that acoustic events outside the head of the speaker do not *as such* realize sentences as linguists characterize them. One radical option would be to hold that the acoustic pattern is entirely epiphenomenal with respect to whether a sentence is realized. Whether a sentence with the properties linguists are interested in is realized would in that case turn entirely on what happens in the mind/brain of language users. I consider this *idealist* option towards the end of the present section, after looking at an intermediate position.

The broad idea behind *qualified realism* is that sentence realization depends not only on the acoustic pattern produced but also on the mental states of actual or potential speakers and/or audiences. There are various forms this proposal might take. On a dispositional version, a speech act realizes an IP in all its complexity because its acoustic pattern has certain dispositional properties. For example, a particular mouthing of (1c), below, might be said to realise (1), replete with PRO, etc. (as explicit in (1b)), because certain human speakers and hearers tend to judge tokens of that event's acoustic type to have this structure.

(1) (c) /dʒɒn si:mz tə bil tə wɒnt tə help hɪm'self/

Such a dispositional property of the acoustic event would depend on the relevant parts of the relevant humans' brains filtering and supplementing any information processing carried out by the bare auditory system. The resulting judgements would be at least partially unconscious, a fact that complicates but need not seriously damage the proposal. Generalizing from the example, a dispositionalist would say that tokens of an acoustic event type realize a sentence to the extent that there is a standing tendency on the part of certain humans to treat those tokens as realizations of that sentence, either in speech or in perception.

The dispositionalist version faces a difficulty: some sound types seem to be associated with a plurality of sentences but just one disposition. (2) is a case in point. So the dispositional property associated with the acoustic pattern of a particular acoustic event does not always determine which sentence is realized. To fix this, we would need to associate the acoustic event type with more than one disposition. So, perhaps we are disposed to assign the relatives-are-visiting structure (call it structure A) to tokens of (2) on some occasions and the relatives-being-visited structure (call it structure B) on others. With that distinction in place one could say that the sentence realized by a given token of the acoustic type depends on which disposition the token is grounded in. This attempt to patch up the dispositionalist position trips up when we ask for a principle for categorizing tokens of (2) into two groups. What makes an acoustic event belong to the A-category rather than the B-category would, presumably, be its origin in one mental state rather than some other. The account would need to draw on this fact. Once it does, however, standing dispositions fall away as redundant. We need only say that a sentence with a particular structure is realized in an acoustic event to the extent that it is causally related in some way to a mental state that represents the event as possessed of that structure. This causal version of qualified realism is more promising, and when I talk without qualification of qualified realism, this is what I will have in mind. Will it do?

Call the mental state from which a given acoustic event supposedly inherits its syntactic structure the *imbuing* mental state. Mental states, and thus imbuing mental states, must have content. The acceptability of qualified realism therefore requires some account of how a mental state could represent an acoustic event of type (1c) as having the structure depicted in (1b). But many theories of mental content are challenged by, for example, the silent presence of PRO. Consider a simple causal-covariance theory.⁶ According to causal-covariance theories, the content of a simple mental representation turns on the existence of some form of reliable (causal/nomological) co-variation between tokenings of the representation and tokenings of the entity or property it represents. For example a token of representation C represents an instance of a cow if and only if tokens of C are reliably caused by and only by instances of cows. This account runs into trouble when it is sentence constituents that are being represented. One worry is an apparent circularity. The thought behind qualified realism is that an acoustic event is the realization of a word – the name *John*, say – to the extent that it is caused by a mental representation that expresses the concept *the name John*. But now we are being told that a representation expresses that concept to the extent that it is of a type that is reliably caused by instances of the name.⁷ A more serious difficulty concerns PRO and other syntactic constituents with a null phonology. By hypothesis the imbuing mental state that determines whether (1) is realized in a given acoustic

⁶ There are sophisticated mutations of the simple version but none of the improvements helps here. A classic statement of the simple view is Dretske 1981. This is criticised and purportedly improved on in Fodor 1987. Teleosemantic approaches, e.g. Millikan 1984, Dretske 1988: ch. 3, and Papineau 1993 all draw essentially on a causal connection between the representer and the represented, albeit in different ways. For a survey of some pros and cons of different theories of content, see Barber 2005, chs. 4-7. The discussion that follows takes its lead from Rey's claim (Rey 2005: 243-5) that externalist theories of content are incompatible with what I am calling qualified realism.

⁷ This is distinct from the circularity worry in Rey 2006: 2.4.2, more on which in section V.

event contains **PRO**, a representation of PRO.⁸ Something with no physical presence, however, cannot stand in a causal co-variance relationship with a mental representation, or with anything for that matter.

Perhaps the solution is to give up on causal-covariance theories. Conceptual role⁹ theories also look to the causal relations representations enter into for an account of their content. But the other relata in the causal relations include other representations and not (or not only) external physical objects. The causal inefficacy of non-physical entities is therefore no obstacle to their being represented. Perhaps **PRO** represents PRO because the first is deployed in some characteristic way within the subject's cognitive economy, not because there is some standing causal relation between the two.

To evaluate the proposal, consider a bare-bones rendition of the conceptual role view according to which the mental representation COW represent cowhood because it occurs in various strings within the belief box of the subject whose representation it is:

COWS EAT GRASS

COWS LIVE ON FARMS

COWS HAVE FOUR STOMACHS

COWS... etc.

⁸ The usual convention is to put representers in upper case lettering and representeds in lower case. But 'PRO' is already in upper case, and necessarily so since 'big PRO' contrasts with 'little pro', a contrast that will come into play later. For representations of PRO I therefore use boldface.

⁹ Or inferential role, causal role, functional role, or procedural profile. There is no systematic terminology, and there are different implementations of the general approach. A classic and optimistic survey of the stance can be found in Block 1986; Fodor & Lepore 1992 contains trenchant criticism. As with causal co-variance theories I shall not be concerned here with stock objections and replies.

EAT, GRASS, LIVE, and so on receive their content in turn from the relation they stand in to other representations, including to COW itself since the theory is holistic. Something falls under the concept expressed by COW if and only if it satisfies all of COW's *descriptors*, i.e. it eats grass, lives on a farm, has four stomachs, etc. The equivalent list for **PRO**, which accounts for its representing PRO, might include:

PRO IS A CASELESS PRONOMINAL

PRO STANDS FOR THE MISSING SUBJECT OF A NON-FINITE CLAUSE

PRO IS PHONOLOGICALLY NULL

Let us say for simplicity that this list of **PRO**'s descriptors is complete. It leans pathetically on what someone sympathetic to mainstream generative syntax might say. That these strings are found anywhere in the head of an ordinary speaker is highly implausible. But we can leave these details aside. The point is just that internal causal role (the precise nature of which depends on how the rules of a grammar are implemented) rather than causal co-variance is what matters to content. Talk of descriptors extracted from explicitly tokened strings is merely a way of making the proposal vivid.

Armed with this theory of how **PRO** has the content it does, can we salvage qualified realism? Suppose the subject hears an acoustic event, e , in which someone mouths (1c). She forms a judgement by tokening Mentalese string (J).¹⁰

(J) **PRO IS A CONSTITUENT OF THE SENTENCE REALIZED IN ACOUSTIC
EVENT E**

¹⁰ I am assuming that (J) does not itself create another descriptor for **PRO**. That is, I am assuming that it is partial rather than total causal role that determines content. This is to take sides arbitrarily in an important intra-theoretical debate for conceptual role theorists. It does not affect the wider issue.

If we regard that judgement as an imbuing mental state, its being made constitutes *e*'s containing PRO as a constituent – or, if the same is true for other constituents, *e*'s realizing the full IP in (1b). (J) is therefore true self-fulfillingly and sentences do exist in the sense of being sometimes realized. Rey's argument is defeated.

This is too quick. The three descriptors give **PRO** its content, but they do not require us to say that the subject's tokening of (J) is self-fulfillingly true. This further claim, which I will later call into question, is generated only if the conceptual role account of content is allied with qualified realism in the form illustrated using (J). It is open to us to divorce the two, use the conceptual role theory to assign content to **PRO**, but say that the subject's tokening of (J) is false rather than self-fulfillingly true. This is no less a coherent position for all that has been said so far. Indeed, it is something like the position adopted by Rey. He claims that although human subjects discern PRO and the rest of (1b) in *e*, they are mistaken. I will leave qualified realism dangling, taking it up again only in section V. In the next section I argue that Rey's scepticism is untenable, adding urgency to the subsequent evaluation of qualified realism. I end up keeping some elements of the view and rejecting others.

Before considering Rey's sceptical position I will look briefly at one final affirmative position, the idealist view mentioned earlier. This has affinities with both Rey's view and with qualified realism, but has failings they lack. It is therefore important to distinguish it from both.¹¹ In common with the negative position Rey himself advocates, it holds that communication involves a felicitous co-ordination between the producer of an utterance and the utterance's audience. Both parties assume that utterances have a certain acoustic structure, with PRO being among the components of that acoustic structure if the utterance is of (1). Both producer and audience are mistaken in this, but their mutual error cancels itself out. The

¹¹ The view is like that of McX in Quine 1948 or Parmenides in Plato's *Sophist*.

speaker forms the intention (at an unconscious level) to perform a sound containing PRO and the rest, and with luck the hearer thinks she perceives PRO and the rest as part of the resulting acoustic pattern. The acoustic pattern does not in fact contain PRO, etc., but neither is any the worse off since they both interpret the speech act in the same way, treating it as an utterance of (1b) rather than of some PRO-less doppelgänger.

This thought does not determine an answer to the question of whether sentences are ever realized. Rey denies they are. Idealists say that, despite never being realized in acoustic events, sentences are realized in the minds of language users. Ray Jackendoff adopts this position in claiming that discrete vowels and words (along with shapes, musical structures, and, presumably, syntactic structures) are elements in a ‘projected world’, not a real world ‘out there’, “‘out-thereness” being as much a mentally supplied attribute as, say, squareness’ (Jackendoff 1983: 26). The projected world is the world ‘as unconsciously organized by the mind’ (Jackendoff 1983: 29). Since the mind is the brain and theoretical linguistics is one of the cognitive sciences, it is still an objective matter that a given sentence has a particular structure, e.g. one containing PRO.

The most convincing reasons for resisting idealism in this context are rooted in classical objections to parallel views about other categories of object.¹² Many are tempted to say that Valhalla, even if we cannot reach it via the rainbow bridge, exists in our minds. Yet

¹² Rey adds to these classical objections an objection that is specific to the case of idealism about sentences: if a sentence is in the mind of its producer, how can the audience be caused to hear it? And if it is in the mind of the hearer, how can a speaker produce it? This strikes me as unfair to Jackendoff. After all, to the extent that mental events are physical, mental events have causal powers. So causation is not a problem. And if we allow that there could be regular co-incidence in respect of which sentences occur in the speaker’s and the hearer’s projected worlds despite their non-realization in the acoustic stream, much as, on Rey’s own view, there could be communication despite speaker’s and hearer’s *folie*, we might reasonably describe this coincidence as the hearer perceiving the same sentence that the speaker produces.

what goes for Valhalla must surely go for its five-hundred and forty doors. And if Valhalla has doors, they must be spatially situated in its walls; but ideas do not have spatial extension. A different worry is that all one's beliefs about Valhalla and other non-existents would come out true, since one could hardly be mistaken about something whose sole existence consists in its being thought about. Indeed, all one's beliefs about anything are at risk of being automatically true, since any mad belief about Big Ben, say, might easily be thought of as a true belief if interpreted as being about a mental object.

IV. Rey's negative stance and the use of illusion in communication

Rey is encouraged towards his negative stance by faults with affirmative accounts of sentence ontology. One option not yet considered is that sentences do not in fact contain PRO or have complex structures, and linguists are wrong to claim they do. This seems a desperate strategy. It would undermine what have long been accepted as plausible principles of explanation.¹³ Nor is it any help to stipulate that sentence are physical types and conclude that IPs are not sentences. That merely raises the question in a different form: are there any IPs?

According to Rey, then, sentences, with their complex structures, are not realized in acoustic patterns considered alone; nor do acoustic events realize sentences in some way that is dependent on the mind of the speaker or perceiver of the event; nor are sentences realized but only in the mind and not in the external world. Rather, sentences are never realized, we only think they are. They belong to the category of what, after Brentano 1874, he calls *intentional inexistents*, defined as follows (Rey 2006: 1.2).


¹³ (1b), for example, but not (1a) or (1c), reveals why *himself* takes *John* not *Bill* for its antecedent. Fodor 1995, section 8.2 ('Are empty categories real?') is a helpful discussion of the empirical case for empty categories.

y is an intentional inexistent for a representational system S iff_{df} there is a sufficiently stable representation in S that has the content $[y]$ and y doesn't exist.

(The square “Brentano brackets” are a way of signalling that we seem sometimes to need to neuter words of their referential commitments in order to use them to express the content of a representation.) Other intentional inexistents are said to include Zeus and Kanizsa triangles.¹⁴ Like them, neither PRO nor any sentence containing PRO exists. But the assumption that they do exist, because common to all conversational participants, renders their non-existence harmless, Rey claims. Communication can still take place. Here is the *folie à deux* view in Rey's own words:

To a first approximation, instructions issue from speakers' phonological systems to produce certain SLEs, and these instructions cause various motions in their articulatory systems, which in turn produce various wave-forms in the air. These wave-forms turn out, however, not to reliably correspond to the SLEs specified in the instructions. All that seems to be true is that when they impinge on the auditory system of an appropriate hearer, this hearer's phonological system will be able to make an extremely good guess about *the intentional content* of the speaker's instructions, not about any actual SLEs, which, *ex hypothesi*, never actually got uttered. Indeed, this sort of guessing in general is so good, and the resulting perceptual illusion so vivid, that it goes largely unnoticed, and speakers and hearers alike take

¹⁴ Rey 2005: 248. I say a little more about the Kanizsa triangle illusion (below) in note 20. This particular alleged intentional inexistent is (or isn't) in the centre of the image, pointing upwards.



themselves to be producing and hearing the SLEs themselves. It is in this way that it's a kind of *folie à deux* (or *folie à n*, for the n speakers of a common language): the speaker has the illusion of producing an SLE that the hearer has the illusion of hearing, with however the happy result that the hearer is usually able to determine precisely what the speaker intended to utter. Indeed, were SLE tokens actually to exist, it would be something of an accident. Their existence is completely inessential to the success of normal communication and to the needs of linguistic theory. (Rey 2005: 240)

I aim to reject this claim that speaker and audience are comfortably demented together, and replace it with an alternative according to which sentences do, after all, exist. The route to that place will be a winding one, but we can start, at least, with a simple objection. Communication allows, among other things, the transfer of knowledge. Yet knowledge through testimony seems to be ruled out by Rey's view if we grant the no-false-lemmas principle as a requirement on knowledge.

No-false-lemmas principle: Belief is knowledge only if it is not based on falsehoods.¹⁵

¹⁵ This is only the necessity direction of the principle. The sufficiency direction is controversial but irrelevant in the present context. The principle was introduced to deal with Gettier's examples (Gettier 1963), but similar examples that respected the principle soon emerged, suggesting that it is not sufficient for knowledge that a belief be based on no false lemmas (and true and justified). See Lycan 2006 for more on the principle, first propounded in Harman 1965: 92. Most epistemologists are committed to the principle's necessity, either explicitly, or implicitly as a consequence of their explicit position.

Assuming Rey is not prepared to embrace full-blooded scepticism about testimony, he must either explain why his more limited scepticism cannot be combined with this principle to yield the more full-blooded scepticism, or argue that the principle itself should be rejected.

It is by now a canonical move in the philosophy of linguistics to dismiss critical appeals to knowledge's supposedly essential properties as the illicit use of *a priori* intuitions to undermine empirical discovery. Thomas Nagel, John Searle, and others, for example, objected to the attribution, by linguists to ordinary speakers, of unconscious knowledge of complex grammatical rules. Knowledge, they said, must be accessible to consciousness, at least in principle. The stock response to this involves pointing out that the Cartesian notion of knowledge, relative to which this accessibility constraint may indeed be legitimate, is simply irrelevant to cognitive explanation. The existence of knowledge states of the kind linguists are interested in need not be appreciated, actually or potentially, by the subject. The same dialectical pattern occurs for other allegedly essential features of knowledge, for example its inferential promiscuity, its internal justifiability, or the subject's possession of the constituent concepts.¹⁶ Appeal to the no-false-lemmas principle looks suspiciously like one more instance of this form of reasoning, inviting the reaction that the no-false-lemmas principle is but an *a priori* intuition grounded in folk or philosophical notions, not a serious obstacle to scientific advancement.

This would be a mistake. The objection to Rey's argument does, it is true, draw on a claim about an essential property of knowledge. It was, moreover, arrived at by reflecting *a priori* on folk intuitions and tweaking accordingly. But the issue here is not whether the relation between speakers on the one hand and the principles of an explanatorily adequate grammar on the other is the knowledge relation. Rather, it is whether Rey's claim about the

¹⁶ See Nagel 1969 and Searle 1990. Chomsky 2000: 93-8 responds to both. The introduction to Barber 2003a details other instances of this pattern of dispute.

ontology of linguistic objects is compatible with attributions of knowledge at a far higher level – for example, knowledge that Julius Caesar first invaded England in 55 BC, or any other judgement adduced through testimony.

Am I shunting the standard dialectic into a siding too hastily? The representational states Rey is talking about are situated at a very low level of cognition. So perhaps their representing falsely falls outside the scope of the principle, which concerns only ‘belief’. This reaction fails to take account of *why* the no-false-lemmas principle is so plausible. If we infer one belief from another, then the inferred-to belief inherits the inferred-from belief’s epistemic value so long as the inference is legitimate. If the inferred-from belief is false there is nothing worth inheriting. Being inferable from it confers no epistemic value. The same could be said about inferring from a misrepresentation of any kind, not merely from a false belief. If a representational state is operative in the generation of a further representational state, and the generation is of a kind that preserves accuracy, then the second representational state inherits the accuracy of the first. But if the first is a *mis*representation, its accuracy is not worth inheriting.

So the threat the principle poses to Rey’s position needs, at the very least, to be taken seriously. But might not testimony through illusion be, precisely, a setting in which the principle fails? To help us determine whether PRO’s realization is as illusory as Rey claims, it will be useful to have before us an example of communication involving a more blatant illusion: the bent-stick illusion.

R and J are star-crossed lovers. Kept apart by their feuding families and attended at all times by watchful minders, their estates are separated by a tall wall. This wall crosses a pond that straddles the border, hovering a few centimetres above the water’s surface (figure 1).




Figure 1. The scene J interprets as a proposal by R to avoid meeting again before tomorrow.

R devises a code whereby, even when both are under watch, R and J can communicate by delving sticks into the pond. The code is passed to J in a sealed envelope delivered by a sympathetic monk. R's note announces that a clockwise-rotating bent stick means they should not rendezvous again before tomorrow. Neither R nor J is aware that the bent stick illusion is an illusion. They think a straight stick bends as it is delved in water. R rotates a partially submerged straight stick and J interprets the act as R intends. Communication succeeds despite the shared false belief. This scenario appears to undermine the no-false-lemmas principle. But if testimonial knowledge could arise from false lemmas, nothing stands in the way of supposing that testimonial knowledge could arise from the illusion that sentences exist. Rey could, then, be commended for having discovered a surprising class of counterexamples to a popular and otherwise plausible constraint on knowledge.

Against this, the no-false-lemmas principle stands in need of modification for independent reasons, and the modified version better accommodates the bent stick example just described. The independent reasons have to do with examples such as that of K. A, B, C,

and D all work in K's office. K has good evidence that A, who doesn't own a Ford, owns a Ford, as well as good evidence that B, C, and D, who do own Fords, all own Fords. K assumes that someone in her office owns a Ford. It does not seem to matter to us in our assessment of K as a knower that one of her grounding beliefs is false. Another example: K lives in a culture that brands left-handed people as witches and warlocks, both literally in that pentangles are seared onto their foreheads as children, and figuratively in that they are judged to worship the devil, sacrifice cats, and so on. K sees P with a pentangle on his forehead, infers that he is a warlock, and from this concludes that he is left-handed, as indeed he is. Although K has inferred from a false belief, it seems fair to say that she knows that P is left-handed.¹⁷

Those sympathetic to the no-false-lemmas principle respond to such cases with a modification:

The no-essential-false-lemmas principle: Belief is knowledge only if it is not based essentially on falsehoods.¹⁸

This steers round the problem cases. It also steers around our stick case. It is clearly inessential to R and J's ability to communicate that J suffers from the bent-stick illusion. If she knew about the illusion and so was aware that the stick was straight, she would carry on interpreting as before. She would just assume that R was ignorant in an irrelevant way.

Rey could, at this point, conjure up a variant of the stick example, one in which J's delusion that the stick is bent *is* essential for testimonial success. Suppose that R's advertised

¹⁷ The first example is due to Lehrer 1965; see also Rozeboom 1967. Both are repeated in Lycan 2006. The second example was suggested to me by Derek Matravers.

¹⁸ See e.g. Harman 1973: 47, Peacocke 1986, p. 129 (final line), and Lycan 2006. Peacocke uses the principle in this form to argue against the counterfactual account of knowledge in Nozick 1981.

code for no further rendezvous that day is, as before, a rotating bent stick, but in addition his advertised code for meeting that same evening is a rotating straight stick. R places a straight stick in the water. It looks bent, but whether it *is* bent suddenly matters. With blind faith in the illusion, J would regard the straight stick in the water as a bent stick when decoding. If R, too, is naïve about the effect of the water's surface, they will both fix on the same interpretation when R rotates a straight and partially submerged stick. So J's (and R's) false assumption *is* now essential to communicative success and hence to knowledge transfer.

I will attempt to show that the modified no-false-lemmas principle is not undermined even by this new example. This will have the (for me) unwelcome side effect of suggesting that Rey's own position, too, is consistent with the principle. But there is, I will go on to claim, a relevant disanalogy between sticks and linguistic expressions. This disanalogy undermines Rey's position at the same time as it suggests an alternative account of what it is for a sentence to be realized. But first, why is the example in the previous paragraph not a counterexample to the modified no-false-lemmas principle?

We cannot say, without qualification, that J's false assumption that the stick is bent is essential to her coming to an accurate interpretation of R's act. It is essential only given her ignorance as to whether R suffers the illusion. If R and J have had long and sophisticated conversations about refraction prior to their forced isolation from one another, J might assume that his using a straight but partially submerged stick invoked those innocent times, and go along with the humour by interpreting the stick as a genuinely bent stick, knowing full well that it is not really. Alternatively, if she knows that he is aware of the illusion but that he lacks a sense of humour and is given to boorishly emphasising the straightness of apparently bent submerged straight sticks, she might interpret it as straight. If, on the other hand, she knows he is ignorant about the illusion, she will again interpret it as if it were bent. If none of these was clearly the case and she was in no position to guess, she would probably not know what to make of his act and no knowledge transfer would take place. There are other imaginable

cases. Traversing this diversity is a true generalization: where knowledge transfer does take place, no matter what permutations of ignorance and awareness exist on the part of each about either the illusion or about the other's awareness of the illusion, J must be able to discern R's intention concerning how J should treat the stick for purposes of interpretation using the explicit code. That is, for communication to occur, she must treat the stick as bent for interpretation purposes just so long as this is how R intends her to treat it for interpretation purposes. Call this the key communicative discernment. (I bother to label it because it will be important in the next section.) In the second stick example, J makes this key communicative discernment and knowledge transfer occurs. But she does not draw essentially on any false assumption. She sees what *appears* to be a bent stick, and assumes that R intended her to interpret it as a bent stick. Any judgement as to the stick's actually *being* bent is inessential to her coming to make the discernment. Even if she in fact makes that false judgement (because she takes the appearance at face value), the false judgement is epiphenomenal relative to her coming to know what R is proposing. In sum: if the modified no-false-lemmas principle is read as ruling out only misrepresentations *whose being taken at face value is essential*, the second example is no threat to it. J makes the judgement that the stick appears bent, but this is not a false essential lemma because it is not false. She also makes the judgement that this appearance is veridical, but this is not a false essential lemma because it is not essential to her reasoning.

At this point Rey could claim that his own account of communication is compatible with the modified no-false-lemmas principle for the same reason. He could say that linguistic communication trades systematically on false representations of the presence of PRO, etc., in acoustic events, just as J's and R's system of communication trades systematically on the visual system's generation of false representations. In neither case is it essential that the misrepresentation be taken at face value and adopted as a premise in the (unconscious)

reasoning that results in testimonial knowledge. Nonetheless, the representation is a misrepresentation: in reality, there are no sentences just as, in reality, there are no bent sticks.

This would be to miss a considerable disanalogy between the stick case and the language case. The appearance of a bend in a submerged stick is a misrepresentation of the shape of the stick only because there is something it would be for the appearance to be correct. There is no such contrast between appearance and reality in the case of language. What is falsely represented as being the case? How would the world have to change in order for the illusion that PRO is often an *acoustic* constituent to become veridical? Would the PRO in (1b) be acoustically realized if it manifested itself as a glottal stop? Or as /prəʊ/? Or /bɪg prəʊ/ to contrast with /lɪtəl prəʊ/? Any of these would be entirely arbitrary. It seems far more plausible to say that there is nothing it would be for PRO to be acoustically realized. In that case, nothing is being represented to be the case that is not the case. So there is no illusion.

Perhaps Rey would concede the major premise – that unless there is something it would be for an appearance to be veridical we cannot talk of it as appearance in the first place, even false appearance – but deny the minor premise. For at one point he says we can...

...*imagine* building a kind of Rube-Goldberg tree structure, replete with little ornamental morphemes suspended from each terminal node, and with wires and pulleys that permitted some movements or connections here and prevented others there. But it is an interesting fact that the noises we produce when we intend to utter sentences are nothing like this, and don't need to be.¹⁹

¹⁹ Rey 2005: 246. Rube Goldberg was a cartoonist who depicted wild and unnecessarily complex machines for everyday tasks, the US equivalent of, for example, Britain's Heath Robinson.

I disagree. What could be true of a sound that would make it the case that PRO rather than, say, little pro, occupied the relevant terminal node in (1b)? It is not merely that the discriminating features of PRO ‘don’t need to be’ given concrete form in an acoustic realization of a Rube-Goldberg machine for (1b). Rather, and *pace* Rey, they cannot even be *imagined* in that guise. Being PRO has to do in part with the non-phonological characteristics of PRO, such as **PRO**’s first and second descriptors listed earlier. What gives PRO the properties that distinguish it from pro, or from the divide between the third and fourth movements of Beethoven’s Fifth Symphony for that matter, are assumptions on the part of speaker and hearer, assumptions that unconsciously lead the latter to discern PRO but not these other things as a constituent, and that unconsciously lead the former to expect that this is what the latter will discern. No wires or pulleys, or their acoustic equivalents, could represent movement or co-indexing and so on intrinsically. Rey might insist on the point by noting that we can envisage a linguistic community whose members communicate with verbal read-outs of tree diagrams like (1b), being sure to say ‘big PRO’ rather than just ‘PRO’ at the relevant point. But we are tempted to say that this would constitute an acoustic realization of (1b) rather than of (1b*) – where (1b*) is like (1b) save for the replacement of big PRO by little pro – only because a convention exists among linguists of labelling these two theoretical posits this way round. There is nothing intrinsic to the sound event – the reading aloud of the tree diagram – that would make it the realization of one rather than the other structure.

None of this shows that Rey’s position is floored by the no-false-lemmas principle. But our consideration of the bent stick illusion does suggest that he must confront the no-false-lemmas principle alone, for two reasons. First, the no-false-lemmas principle is not undermined by other easily imaginable cases involving testimony through illusion. He cannot, therefore, simply reject the principle by citing these other cases in his support. But second, nor can he explain why his position is compatible with the principle by looking to how these other easily imaginable cases are compatible with it and demanding parity. Compatibility

depends on an appearance-reality distinction that bears scrutiny in the stick case but not in the language case.²⁰ I conclude that an affirmative account, in which knowledge is not based on falsehood, would be preferable. I develop and defend one in the next section.

V. Illocutionary syntax?

Of the affirmative positions considered in section III, one, qualified realism, was left dangling. The thought there was that acoustic events inherit their syntactic properties from imbuing mental states, which represent the events as having that syntax. These states would thus be self-fulfillingly accurate in what they represent. I will defend the inheritance idea, but only after dissociating it from the apparent corollary that imbuing states are self-fulfillingly accurate.

That a belief about the syntactic properties of an acoustic event self-fulfillingly bestows the acoustic event with precisely those syntactic properties is ruled out if the following is a condition on all beliefs: no belief could be correct solely by being held. That is, for any p , someone's believing that p cannot in itself constitute their believing it correctly. But such minimal objectivity *is* a plausible condition on belief. Why so?

The difference between the states of believing that p and believing that q turns on differences in what their correctness requires of the world. For the first the world must be p -like and for the second it must be q -like. If two beliefs were correct solely by virtue of being held, there would be nothing to distinguish them. So at most one belief state is self-fulfillingly

²⁰ And perhaps not in the Kanizsa triangle case either (see note 14). The case is complicated by vagueness in what suffices for the realization of a triangle, but at least we can say, and we can say why, it is an objective fact that there is no linear division between darker and lighter zones between each pair of "pacman corners".

true.²¹ In the present context, this would yield the unacceptable result that the attribution of a (1b)-syntax to an acoustic event and the attribution of, say, a (1b*)-syntax would always come to the same thing. But in any case, a self-fulfillingly true belief is objectionable for its lack of objective utility. For any p , if the condition for the belief that p being correct is just that the belief is held, holding the belief correctly could serve no non-subjective purpose. But beliefs about the structure of sentences clearly do have some non-subjective purpose, since they allow us to find out facts about the world and about the minds of others.

The argument just given was directed against the proposal of self-fulfilling beliefs as the imbuing states. A similar line could be taken against the proposal of self-succeeding intentions, e.g. the intention to perform an acoustic event with the structure of (1b) rather than (1b*). The equivalent constraint to the one that rules out self-fulfilling beliefs is that an act cannot be successful solely because it is performed with a specific intention. In particular, the acoustic stream coming out of one's mouth cannot have a particular syntax solely because the utterer intends it to do so, any more than a painting will fetch €3,500 solely because the painter painted it with that intention. Something extra must happen. In the painting case, someone needs to pay. The reasoning behind this constraint is similar. What makes it the case that an act was performed with an intention to bring about i rather than with an intention to bring about j is that the act is successful just in case it results in i , not just in case it results in j . Self-succeeding intentions would be indistinguishable from one another. Besides that, they would impose no requirements on the world and so it is hard to see how it could be of use to anyone to act on one.

I nevertheless want to retain the thought that imbuing states are responsible for the syntax of acoustic events. I have been arguing only against the thought that imbuing states

²¹ I am talking about states that are self-fulfillingly true *constitutively*, not beliefs that are self-fulfillingly true because, for example, they cause you to behave in a certain way – e.g. to leap the crevasse – when you could only do so because you believed you could.

could be self-fulfillingly true or self-fulfillingly successful. That rules out only certain simplistic accounts of the content of the imbuing state, e.g. (for belief) that the utterance just heard has a particular structure, or (for intention) that it have that structure.

A clue to the nature and content of the imbuing states can be found in the discussion of the bent stick example. In order for testimonial transfer of knowledge to occur, J has to make what I earlier called the key communicative discernment. She must discern R's intention concerning how she will regard the stick for purposes of interpretation using their code. The equivalent in syntax to this requirement is that, in order for a hearer to understand a speaker, she must recognize which sentence the speaker intended the hearer to treat as the object of interpretation. This is an epistemic claim, but there is nothing to prevent us from supposing that a constitutive fact underpins it, namely, that the identity of the sentence uttered turns on the intention of the speaker to be recognized by his or her audience as having produced that sentence. (The equivalent shift from epistemic to metaphysical claim is not available in the stick-in-a-pond scenario because a stick's being genuinely bent or straight is independent of R's and J's mental states. This is the difference between communication through sticks and communication through natural language that, I suggested earlier, Rey has failed to take on board.)

This proposal amounts to treating the syntax of an acoustic event as an illocutionary property. An illocutionary act (in the sense in which I am using the term²²) is an act like any other in that it is performed with an intention. It is also like other acts in that it is successful just in case it results in the satisfaction of the intention that gives rise to it. What sets illocutionary acts apart is that their success also consists in the recognition, by some intended

²² Though I am using the term 'illocutionary' stipulatively here, my use is broadly in keeping with that of Searle 1969: 47.

audience, of the intention with which it was performed. In other words, an act is illocutionary if and only if it belongs to some kind *k* that satisfies the following schema:

The illocutionary schema: A person *ks* if and only if they act with the intention of being judged to have *ked*.

Archetypal illocutionary acts include promises, assertions, and questions. Having an illocutionary *property* consists in being a token of some specific illocutionary action type. Being a promise, being an assertion, and being a question are all examples of illocutionary properties of acoustic events or graphic events as it may be. My claim is that being the utterance of a particular sentence, e.g. the sentence with structure (1b), is also an illocutionary property of an acoustic event.

The advantage of treating syntactic properties of an acoustic event as illocutionary is that they need no more be physically manifest than, say, its meaning properties are. A sentence can therefore have ‘hidden’ structure replete with PRO, trace elements, and so forth. Attributing hidden structure need not be thought of as any kind of error on the part of the conversants, even a harmless one of the kind Rey ascribes. In a successful exchange, S believes H will, as S intends, treat the acoustic event he is about to produce as containing PRO (and interpret accordingly); H, on the other hand, will believe that S intended her to regard the acoustic event as containing PRO. If syntax is illocutionary, both beliefs are true, so there is no conflict with the modified no-false-lemmas principle. For the same reason, there is no *folie à deux*.

There are several reasons why treating syntactic properties as illocutionary may seem absurd. For one thing, the illocutionary schema looks to be a mark of the meaningful, but to produce a piece of syntax is not as such to do anything meaningful. For another, it is traditional to think of producing a sentence as producing a characteristic sound of some kind.

While some modification of this idea may be tolerable and necessary in view of the empirically motivated hypothesis of phonologically null constituents and other kinds of hidden structure, the present proposal seems to divorce sentences from sound types altogether. And finally, syntax is standardly thought of as the product of a modular system, while intentions are not. I will consider these three reservations in turn, before looking to see if any of the criticisms of qualified realism in Rey's contribution to the present issue (or of 'relational theories' as he calls them) apply to my version of it. In the process, the nature of the proposal should itself become clearer.

There are several attractions to treating the illocutionary schema above as the mark of a meaningful act. It arguably captures an interesting characteristic of all and only those acts we would pre-theoretically call meaningful, including literally asserting, implicating, asking a question, giving an order, and perhaps even producing an artwork. Moreover, it gels with a plausible account of communicative success and failure: understanding is what occurs when an illocutionary intention is recognized, misunderstanding is what occurs when it is not. Finally, the dependency of meaning on the mind of the utterer is explained without releasing the utterer from all constraints on what she can or cannot mean. She cannot, like Humpty, intend to be understood one way when she produces a sentence that she has good reason to believe will be interpreted quite differently.

Producing a piece of syntax does not in itself amount to doing something meaningful in the pre-theoretical sense. So I am committed to denying that being illocutionary is the mark of the meaningful, even if it is a necessary characteristic. Any upset caused by this concession is ameliorated by the smooth way in which this account of sentence realization fits into a hierarchy of illocutionary intentions. I have in mind the hierarchy lying behind Davidson's notion of 'first meaning', his surrogate for literal meaning.²³ Consider a famous example

²³ Davidson 1986. An exegesis can be found in Barber 2004, on which this paragraph leans.

(Grice 1975: 33). Professor A is writing a testimonial for a philosophy job candidate. His letter reads: *Dear Sir, Mr X's command of English is excellent, and his attendance at tutorials has been regular. Yours, etc.* Professor A's utterance involves nested intentions. Ultimately, members of the appointment committee are intended to spot that, in writing these words, Professor A means that Mr X is a poor candidate. But they are intended to do so only after interpreting him as asserting that Mr X's command of English is excellent. The nested sequence of intentions is thus:

- (i) to utter (write) the sentence 'Mr X's command of English is excellent'

in order that:

- (ii) the committee interpret his utterance as a statement that Mr X's command of English is excellent,

in order that:

- (iii) the committee interpret his utterance as implicating that Mr X is a poor candidate

in order that:

- (iv) Mr X not be offered the position unless they are quite desperate.

On the normal view, while the second and third of these intended outcomes are illocutionary in a way that is characteristic of meaningful acts, the first, like the fourth, is not. The first meaning of an utterance is read off from the first intended outcome that is 'self-referring', as

Davidson puts it, or illocutionary as I would put it. So the first meaning of Professor A’s act is that Mr X’s command of English is excellent, since Professor A intends his utterance to be interpreted by members of the committee as a statement that Mr X’s command of English is excellent, and Professor A will be successful in this only if the committee recognizes the intention itself. Success in what is intended requires, constitutively, recognition of that intention by an intended audience.

I hold that this hierarchy needs to be broken down further. Professor A has the intention:

(ia) to concatenate keystrokes thus:

M^r^spacebar^X^apostrophe^s^spacebar^c^o^m^m^a^n^d^spacebar^o^f^spacebar^E
 ^n^g^l^i^s^h^spacebar^i^s^spacebar^e^x^c^e^l^l^e^nt^fullstop [or to mouth /'mɪstə
 'ɛksɪz kə'mɑ:nd ɒv 'ɪŋlɪʃ ɪz 'ɛksələnt/ if it is a verbal reference]

in order that:

(ib) the committee take him to have produced the sentence ‘Mr X’s command of English is excellent’ *with* all its hidden syntactic complexity,

and on as before to (iv) via (ii) and (iii). (ia) is not illocutionary. It is possible to produce an acoustic event without having any intention to be recognized as having done so. If I am right, however, (ib) is illocutionary, for all that it is anterior to the ‘first meaning’ given by (ii). Professor A’s intention to bring about (ib) will be satisfied just in case this intention is

recognized by the committee.²⁴ Producing a sentence involves having intentions that are in this sense ‘headed towards’ meaning. So it is perhaps not surprising that producing sentences has features in common with producing meaningful acts, specifically in respect of being illocutionary.

The second of the three reservations I listed was that, if syntactic properties are illocutionary, any connection between syntax and acoustics looks to have been severed entirely. It is conceivable that a person could intend to be recognized as producing (1b), and *ipso facto* produce (1b) according to the present view, even if the only sound they mouth is (1d) as opposed to (1c).

(1) (d) /dʒɒn wɒntz bɪl tə si:m tə help hɪm'self/

This is, at the very least, counterintuitive.

It is also a mistaken accusation, for reasons touched on already. One attraction of treating meaningful acts as illocutionary is that, though what an act means depends on the performer’s intentions, the performer is constrained in what she or he can mean when she utters a particular sentence. She can only form the intention to mean such and such if she has some expectation that this is how she will be interpreted by her intended audience. She cannot form that expectation by an act of will. She needs some evidence, or what she takes to be evidence, perhaps in the form of acquaintance with relevantly similar acts having been interpreted along relevantly similar lines in the past. What goes for semantics goes for syntax. Forming the intention to mouth (1d) (as opposed to (1c)) as a means of bringing it about that your audience judges you to have uttered a sentence with structure (1b) would require you to

²⁴ Of course, one can rehearse in private just as one can say meaningful things in private. But there are standard ways of dealing with the latter cases – see Barber 2003b – that would apply equally to the former.

have highly unusual beliefs about the phonological form of the constituents of (1b), and it is not open to speakers to generate such unusual beliefs at will. We are constrained in what we can believe about phonology by a mixture of the evidence of our senses and innately specified knowledge. The result: we cannot produce a particular sentence using any sound we wish. But this fact is an empirical consequence of my account of what it is to produce a sentence rather than a matter of definition.²⁵

The final reservation I wish to anticipate (before looking at Rey's) is that talk of intentions regarding syntactic structure is ludicrously high-minded given the unreflective and speedy nature of syntactic processing. This nature is more consistent with sentence production and identification being the work of a dedicated module than with its being the product of conscious, rational deliberation involving the propositional attitudes. To the extent that the illocutionary schema is prompted by reflection on R and J's predicament, we appear to have yet another reason not to seek to understand natural language by using thought experiments applicable, at best, to non-linguistic forms of communication.

Any debate over the existence of intentions, here, might be expected to track the more developed debate alluded to earlier over whether ordinary speakers and hearers genuinely possess the knowledge of language that linguists are in the habit of ascribing to them. Is this knowledge propositional? Is it inferentially promiscuous? Is it conscious? These are delicate questions. If this 'knowledge' lacks too many of the features we typically attribute to knowledge states, carrying on calling it knowledge starts to sound hollow and the explanatory potential of knowledge attributions is diminished. Chomsky has sought to sidestep debates over the essential properties of knowledge by introducing a term of art, 'cognizes', in place of 'knows', at least as a temporary measure. I suggest we could do the same by replacing

²⁵ I am ignoring legalistic reasons one might have for insisting on minimal acoustic conditions since they are irrelevant to the explanation of how language can be used in communication. For more on this issue, see Barber 2003b: 389-93.

‘intends’ with ‘conates’. I don’t claim that either move is satisfactory, only that the debate has the potential to pan out in similar ways. Besides, we need *some* way of describing astonishingly fast intentional or quasi-intentional actions. Learning to touch type or play the piano involves turning simple, conscious, and undeniably intentional acts into something resembling the ‘intentions’ that, I claim, enable the realization of syntax. Whether we call them intentions is not really the issue.

Rey 2006 sketches various objections to what he calls relational views, his label (roughly) for what I am calling qualified realism. There are many kinds of relational view (‘dispositional, aetiological, statistical’, p. 562), and Rey’s criticisms are, doubtlessly, effective against some. I will confine myself to asking how far his variability objection (pp. 562-3), his circularity objection (pp. 563-4), and his redundancy objection (pp. 564) disturb the illocutionary view I am proposing.

The variability objection is that type identities across acoustic events can only be purchased by imposing suspect restrictions on acoustic criteria, often normative ones that have no place in scientific explanation. Consider what counts as an acceptable execution of the word ‘Wednesday’. There is massive variation in the contextual factors that lead people to judge an acoustic event as the realization of this or any word. Must it be /wɛdnzdeɪ/ (with two /d/s)? That would be prescriptive. The problem is not that insisting on this requirement is illiberal. It is equally a prescriptive judgement, and so equally problematic in the context of scientific enquiry, to be permissive in allowing, in addition, /wɛnzdeɪ/ and /wɛnzdi/. This and other ways of systematizing the variability are likely, as Rey puts it, to end up characterizing ‘an idiolect with a gun-boat – and an obsessive, *a priori* philosopher at the helm’.²⁶

²⁶ Rey 2006: 563; the allusion is to an aphorism in Weinreich 1945, that a language is a dialect with an army and a navy.

Since I am not tying conditions on realization to anything other than the intention, I cannot be accused of tidying up acoustic criteria with suspect normative constraints. If anything I could be charged with the opposite: dissociating the realization conditions from acoustics entirely. I have already addressed this opposite charge.

The second of Rey's criticisms of relational theories hits closer to home. The problem, he suggests, is one of circularity.

Leading theories in phonology seem driven to characterize SLEs [standard linguistic entities] in terms of either a hearer's "categorical" *perception of a noise as an SLE*, or of a speaker's *intention to produce one* (see, e.g., Liberman & Mattingly 1985 and Bromberger & Halle 2000: 24-5). But how are these perceptions or intentions to be understood on a relational analysis? As an intention to produce a sound that will be produced by just such an intention? (Rey 2006: 563)

On my view (though not explicitly on the views of the two sets of authors cited), the imbuing intention is the intention to be recognized as having produced that particular SLE. Rey's objection can, I think, be read as parallel to Grice's against Searle. Searle 1969: 47 treats meaning much as I am treating syntax, namely, as an illocutionary property. He argues that this improves on Grice's account of the content of meaning generating intentions. Presumably thinking of the reoccurrence of the relevant substitution instance of '*k*' on the right-hand side of the biconditional in the illocutionary schema, Grice replies:

[Searle's and other's] position hardly seems satisfactory when we see that it involves attribution to speakers of an intention which is specified in terms of the very notion of meaning which is being analyzed (or in terms of a dangerously close relation of that notion). Circularity seems to be blatantly abroad. (Grice 1989: 352)

The proper answer to this? Don't be fooled by the bi-conditionality of the illocutionary schema into thinking that it is meant to provide a conceptual analysis of the act type substituted for 'k'. Satisfying the schema is necessary and sufficient for an act type to be illocutionary, and being illocutionary is, modulo the remarks made in answer to the previous objection, the mark of meaningfulness. This latter is a substantive claim that, if true at all, is not vacuously true or circular. But many different act types satisfy the schema, so satisfying the schema does not as such discriminate one type of meaningful act from another. For example, asking after the nearest bank and calling last orders are both illocutionary acts, so something else must serve to discriminate the one from the other. Similarly, if uttering a sentence is, as I maintain, an illocutionary act, something other than satisfying the schema must distinguish uttering (1b) from uttering some other sentence. The trick is to say what. In the syntactic case, which is my concern here, looking for purely acoustic criteria will not do since acoustic criteria cut both too coarsely and too finely. (Too coarsely because of, for example, structural ambiguities; too finely because of variability in enunciation.) Treating sentence realization as an illocutionary matter bypasses this lack of correspondence. When I mouth (1c), its being a realization of (1b) rather than (1b**) – identical to (1b) save for the presence of *won* in place of *want* – must turn on non-acoustic features of the event. I claim that the difference has to do with which of the two words the speaker intends the audience to judge is a syntactic constituent of the sentence uttered. The two intentions are discernible from one another because a representation of the word *want* will of necessity have different descriptors than a representation of the word *won*. After all, the two words express distinct concepts, belong to distinct syntactic categories, and are acoustically distinct in some contexts even if they are acoustically indistinguishable in the context of (1b)/(1b**).

Rey's final charge against relational views is that even if they are achievable they are redundant:

Perhaps there is some ingenious combination of dispositional and aetiological approach that could circumvent these two difficulties and provide a relational analyses. But, given these difficulties, one might wonder why it's so important to provide one. (Rey 2006: 564)

After quoting Chomsky's view that extra-mental standard linguistic entities (e.g. 'phonetic values') would be 'spinning wheels' (Chomsky 2000: 129), he continues:

Why does he think it would be wheel spinning? Because if one looks at the theoretical work SLEs are supposed to perform in a theory, their definition in terms of acoustic phenomena – indeed, their very existence – is entirely needless and extraneous to linguistic theory. (Rey 2006: 564)

But is Rey's own position any less guilty of this same charge of being 'entirely needless and extraneous to linguistic theory'? Theories of content aim to state, in a very general way, the conditions under which internal representations are accurate. Chomsky's view is that attempts to supply a theory of content for internal representations are otiose.²⁷ Here is not the place to

²⁷ That is, Chomsky's quarrel is really with the project of giving a theory of objective content for internal representations (2000: 158-60; see Rey 2003a; 2003b for a critical examination of his position). Talk of the inaccuracy of a representation makes sense relative to imposed interpretations, but objective inaccuracy is as irrelevant to linguistics as objective accuracy: 'Accessed by performance systems, the internal representations of language enter into interpretation, thought, and action, but there is no reason to seek any other relation to the world, as might be suggested by a well-known philosophical tradition and inappropriate analogies from informal usage. Misperception raises no difficulties for this approach; it is a matter of how people assign interpretations to interactions they observe – to the reactions of a frog or person in an experiment, a photoreceptor that is

evaluate that claim. I wish only to make the *tu quoque* observation that the redundancy charge does not discriminate in favour of accounts of sentence realization that say internal representations are inaccurate and against accounts that say internal representations are accurate. If establishing that internal representations are accurate is redundant relative to the aims of linguistic theory, so is establishing that they are inaccurate.

VI. Summary and conclusion

Sentences are realized by acoustic (or graphic, etc.) events, but the acoustic properties of these events do not provide criteria for this realization because many features of sentences, though relevant to which sentence has been realized, are not acoustically manifest. In the positive part of this paper I defended the thought that what makes an acoustic event the realization of a particular sentence turns on the intention of the event's producer. This is distinct from other versions of the claim that the mind of the producer or perceiver determines the identity of the sentence being realized. For example, it is distinct from the suggestion that intentions to produce a particular sentence are self-fulfillingly successful, or that beliefs that particular sound events realize a particular sentence are self-fulfillingly true.

It is also distinct from Rey's sceptical suggestion that sentences are never in fact realized. I raised a worry about this latter proposal, namely, that it violated a widely accepted principle of epistemology, the no-false-lemmas principle. In evaluating this worry I invoked a scenario involving communication through what is, without question, an illusion, namely, the bent stick illusion. On first acquaintance this scenario, too, seems to be inconsistent with the principle, to the principle's cost. If that were so, Rey could fairly bat the principle back to the epistemologists for re-evaluation. But on reflection, delved-stick communication turns out to

“deceived”, etc. – a fair topic for internalist inquiry into the psychology of the person who is deciding what to call a “misperception”.’ (Chomsky 2000: 160.)

be consistent with the principle. This consistency requires an appearance/reality distinction that is unavailable in the linguistic case. So Rey cannot look to the stick illusion either to support his case against the principle or to back any assertion that his own position is similarly consistent with the principle.²⁸

²⁸ Presentation of an earlier draft of this paper to a workshop in Dubrovnik in September 2005 has led to substantial changes. I am grateful to everyone there for helping me see that using the adverbial theory to do away with intentional inexistents would introduce more difficulties that it could possibly remove, and to Barry Smith in particular for pointing out that my objection to Georges Rey rested on the no-false-lemmas principle. My thanks, also, to colleagues at the Open University Philosophy Department for useful discussion at a presentation in February 2006.

References

- Barber, Alex, ed. (2003a). 'Introduction'. *Epistemology of Language*. Oxford: Oxford University Press.
- Barber, Alex (2003b). 'Truth conditions and their recognition.' In A. Barber, ed., *Epistemology of Language*. Oxford: Oxford University Press.
- Barber, Alex (2004). 'Idiolects.' In E. N. Zalta, ed., *Stanford Encyclopedia of Philosophy*. [<http://plato.stanford.edu/archives/win2004/entries/idiolects/>](http://plato.stanford.edu/archives/win2004/entries/idiolects/).
- Barber, Alex (2005). *Language and Thought*. Milton Keynes: Open University.
- Block, Ned (1986). 'Advertisement for a Semantics for Psychology.' In P. A. French, ed., *Midwest Studies in Philosophy*. Minneapolis: University of Minnesota Press. 10: 615-78.
- Brentano, F. (1874). *Psychologie vom empirischen Standpunkte*. Leipzig: Duncker & Humblot; 2nd edn, ed. O. Kraus, Leipzig: Meiner, 1924; repr. Hamburg: Meiner, 1973; trans. A.C. Rancurello, D.B. Terrell and L.L. McAlister and ed. L.L. McAlister, as *Psychology from an Empirical Standpoint*, London, Routledge, 1969; 2nd edn, 1995.
- Bromberger, Sylvain and Morris Halle (2000). 'The ontology of phonology (revised).' In N. Burton-Roberts, P. Carr and G. Docherty, eds., *Phonological knowledge: conceptual and empirical issues* 19-37.
- Chomsky, Noam (2000). *New Horizons in the Study of Language and Mind*. Cambridge, UK: Cambridge University Press.
- Davidson, Donald (1968). 'On Saying That.' *Synthese* 19: 130-46.
- Davidson, Donald (1986). 'A Nice Derangement of Epitaphs.' In *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Cambridge, Mass.: Blackwell.
- Devitt, Michael (2006). *Ignorance of Language*. Oxford: Oxford University Press.
- Devitt, Michael and Kim Sterelny (1999). *Language and reality, 2nd edition*. Oxford: Blackwell.

- Dretske, Fred (1981). *Knowledge and the Flow of Information*. Oxford: Blackwell.
- Dretske, Fred (1988). *Explaining Behavior: Reasons in a World of Causes*. Cambridge, Mass.: MIT Press.
- Fodor, J. A. (1987). *Psychosemantics*. Cambridge, Mass.: MIT Press.
- Fodor, J. A., T. Bever and M. Garrett (1972). *The psychology of language*. New York: McGraw-Hill.
- Fodor, J. A. and E. Lepore (1992). *Holism: A shopper's guide*. Oxford: Blackwell.
- Fodor, Janet Dean (1995). 'Comprehending Sentence Structure.' In??, ed., *Invitation to Cognitive Science, Vol. ??* Cambridge, Mass.: MIT Press.
- Gettier, Edmund L. (1963). 'Is Justified True Belief Knowledge?' *Analysis* 23: 121-23.
- Grice, H. P. (1975). 'Logic and conversation.' In P. Cole and J. Morgan, eds., *Syntax and Semantics*. Volume 3. London: Academic Press. Reprinted in H. P. Grice, *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press, 1989.
- Grice, H. P. (1989). *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.
- Harman, Gilbert (1965). 'The Inference to the Best Explanation.' *Philosophical Review* 74: 88-95.
- Harman, Gilbert (1973). *Thought*. Princeton: Nj Princeton Univ Pr.
- Jackendoff, R. (1983). *Semantics and Cognition*. Cambridge, Mass.: MIT Press.
- Lehrer, Keith (1965). 'Knowledge, Truth and Evidence.' *Analysis* 25: 168-75.
- Lieberman, Alvin M. and Ignatius G. Mattingly (1985). 'The motor theory of perception revised.' *Cognition* 21(1): 1-36.
- Lycan, William G. (1996). *The Philosophy of Language*. London: Routledge.
- Lycan, William G. (2006). 'On the Gettier Problem Problem.' In S. Hetherington, ed., *Epistemology Futures*. Oxford: Oxford University Press.

- Millikan, Ruth Garret (1984). *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, Mass.: MIT Press.
- Nagel, Thomas (1969). 'Linguistics and Epistemology.' In S. Hook, ed., *Language and Philosophy*. New York: New York University Press; reprinted in Gilbert Harman, Noam Chomsky: Critical Essays, Garden City: New York, 1974.
- Nozick, Robert (1981). *Philosophical Explanations*. Cambridge, Mass.: Harvard University Press.
- Papineau, David (1993). *Philosophical Naturalism*. Cambridge: Blackwell.
- Peacocke, Christopher (1986). *Thoughts: An Essay on Content*. Oxford: Blackwell.
- Quine, Willard V. (1948). 'On What There Is.' *Review of Metaphysics* 2: 21-38.
- Rey, Georges (2003a). 'Chomsky, Intentionality and a CRTT.' In L. Antony and N. Hornstein, eds., *Chomsky and His Critics*.
- Rey, Georges (2003b). 'Intentional Content and a Chomskian Linguistics.' In A. Barber, ed., *Epistemology of Language*. Oxford: Oxford University Press.
- Rey, Georges (2005). 'The intentional inexistence of language - but not cars.' In R. J. Stainton, ed., *Contemporary Debates in Cognitive Science*. Oxford: Blackwell.
References are to page proofs.
- Rey, Georges (2006). 'Conventions, intuitions and linguistic inexistentents: A Reply to Devitt.' *Croatian Journal of Philosophy*, this issue.
- Rey, Georges (in preparation). *Representing nothing: phones, feels, and other intentional inexistents*. Oxford: Oxford University Press.
- Rozeboom, William W. (1967). 'Why I Know So Much More Than You Do.' *American Philosophical Quarterly* 4: 281-90.
- Sapir, Edward (1949). 'The psychological reality of the phoneme.' In D. G. Mandelbaum, ed., *Selected Writings of Edward Sapir in Language, Culture, and Personality*. Berkeley: University of California Press.
- Saussure, F. (1916). *Cours de linguistique générale*. Paris: Charles Bally.

Searle, John R. (1969). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.

Searle, John R. (1990). 'Consciousness, Explanatory Inversion, and Cognitive Science.' *Behavioral and Brain Sciences* 13: 585-96.

Weinreich, Max (1945). 'Der yivo un di problemen fun undzer tsayt.' *Yivo Bleter* 25(1): 3-18.