

[forthcoming in *Linguistic Intuitions, Evidence, and Expertise*, ed. Samuel Schindler (Oxford University Press)]

Linguistic Intuitions: Error Signals and the Voice of Competence

Steven Gross

Johns Hopkins University

Abstract

Linguistic intuitions are a central source of evidence across a variety of linguistic domains. They have also long been a source of controversy. This chapter aims to illuminate the etiology and evidential status of at least some linguistic intuitions by relating them to error signals of the sort posited by accounts of on-line monitoring of speech production and comprehension. The suggestion is framed as a novel reply to Michael Devitt's claim that linguistic intuitions are theory-laden "central systems" responses, rather than endorsed outputs of a modularized language faculty (the "Voice of Competence"). Along the way, it is argued that linguistic intuitions may not constitute a natural kind with a common etiology; and that, for a range of cases, the process by which intuitions used in linguistics are generated amounts to little more than comprehension.

Keywords: linguistic intuitions, monitoring, error signals, grammaticality, acceptability

1. Introduction

A substantial portion of the evidential base of linguistics consists in linguistic intuitions—speakers' typically non-reflective judgments concerning features of linguistic and language-like stimuli. These judgments may be elicited, for example, in answer to such questions as:

Is the following sentence natural and immediately comprehensible, in no way bizarre or outlandish (cf. Chomsky's (1965) gloss on "acceptability")?

She likes chocolate anymore.

Just going by how it sounds, /ptlosh/ is not a possible word in English, but /losh/ is. Please rate the following candidates on a scale from 1 (definitely not possible) to 5 (definitely possible):

/fant/, /zgant/, ...

Do the bolded terms in this sentence co-refer (Gordon & Hendrick 1997)?

*John's roommates met **him** at the restaurant.*

Which phrase are you most likely to use with a friend when leaving (Labov 1996)?

(a) *goodbye* (b) *bye* (c) *bye-bye* (d) *see you* (e) *so long*

Because of their evidential centrality, linguistic intuitions have been the focus of much methodological reflection. There are well-known worries concerning both how they are collected and how they are used: for example, linguists often use themselves as subjects, risking confirmation bias; they may gather too few intuitions to enable statistical analysis; and they may rely on intuitions too much, failing to seek converging (or disconfirming) evidence of other sorts. There are also well-known replies to these worries: for example, intuitions are often now gathered from a statistically well-powered number of naive subjects in a controlled setting; the comparisons that such work has enabled with linguists' own intuitions have tended to validate linguists' intuitions; and there is an ever-growing exploration of other sources of evidence. Much more can be said on these matters. (For reviews with further discussion and references, see, e.g., Schütze 1996; Schütze 2011; Sprouse & Schütze forthcoming.) I mention these familiar debates to set them aside and to distinguish them from this chapter's main question. All parties agree that linguistic intuitions *can* be and often are a good source of evidence. Why are they? What about their etiology enables them to be a good source of evidence?

This chapter suggests that we assign a role to error signals generated by monitoring mechanisms. It will not establish that this is so, but instead aims to render it a plausible, empirically-motivated hypothesis and to consider some of its philosophical consequences. There exists a sizeable body of psycholinguistic research on language-related monitoring. But its potential relevance to the etiology, and thus evidential status, of linguistic intuitions has not been much explored.¹

It is not intended that the proposal extend to *all* linguistic intuitions. Methodological discussions of linguistic intuitions often focus on acceptability judgments as evidence in syntax, but judgments concerning other features—in the examples above, pronounceability, co-reference, and likelihood of use, but not only those—play a significant evidentiary role as well. It is far from obvious that the same account can be given for each. Thus, after exploring the possible role of error signals in generating some judgments of unacceptability, I suggest that linguistic intuitions in fact do not form a natural kind with a shared etiology, discussing in particular the role of utterance comprehension. It's also no part of my proposal that, in those cases where error signals do play some role, there are no other significant causal factors or sources of warrant.

The etiology of linguistic intuitions is of interest for several reasons, beyond the intrinsic interest of better understanding any instance of the mind-brain's goings-on. For one, progress on this specific case contributes to our understanding of intuitive judgment more generally, a topic of significance both for psychologists and philosophers (DePaul & Ramsey 1998). For another, there is the aforementioned question of linguistic intuitions' evidential status. While their status as good evidence may not *require* a deeper knowledge of their etiology (Culbertson & Gross 2009), such knowledge can certainly help clarify and further secure it. Finally, a better understanding of linguistic intuitions' etiology enables us to more fully answer a challenge raised by Michael Devitt (2006a; 2006b) to mentalist conceptions of linguistics—conceptions on which linguistics is a branch of psychology investigating mental mechanisms and processes implicated specifically in language acquisition and linguistic behavior. Indeed, it is this challenge—and its bearing on broader questions in the philosophy of linguistics—that motivates and frames the

¹ Important exceptions include Sprouse (2018) and, especially, Matthews (ms.).

present study. Accordingly, I begin by providing some background on Devitt's views and the discussion it has elicited.

2. Devitt on the “Voice of Competence” and his modest alternative

Why can linguistic intuitions serve as evidence in linguistics? Devitt (2006a; 2006b) contrasts two answers.² According to the “Voice of Competence” view that Devitt rejects, linguistic intuitions are the product of a modularized language faculty that alone delivers the relevant information, or content, to mechanisms responsible for judgment. Judgments with such an etiology, on this view, can provide fruitful evidence for linguistic theorizing because they directly reflect constraints built into mechanisms specifically implicated in language acquisition and linguistic behavior and thus provide speakers privileged access to the linguistic facts. It is this view that Devitt ascribes to proponents of a mentalist conception of linguistics.³

According to Devitt's own “modest” view, while linguistic competence may supply access to the phenomena linguistic intuitions are about, it does not supply the intuitions' content. Rather, intuitions are arrived at via ordinary empirical investigation using the mechanisms responsible for judgment more generally (“central systems”). Linguistic intuitions, thus produced, can provide evidence for linguistic theorizing because experienced language-users, immersed in language, make fairly reliable judgments about many linguistic matters, just as those immersed in their local flora may be fairly reliable about aspects of it. Devitt calls his view “modest” because it needn't advert to any mental states or processes beyond those to which any account of judgment is committed. Importantly, according to Devitt, linguistic intuitions, because they are empirical judgments, are theory-laden as all such judgments are.

Devitt argues that his view provides a *better* answer to the question ‘Why are linguistic intuitions a good source of evidence?’. Among his main arguments is that, not only do we lack a positive account of how a module embodying grammatical constraints might generate intuitions suited to play the evidential role the mentalist requires, but it is hard to see how one might go: we lack so much as “the beginnings of a positive answer” (Devitt 2016b: 118). (Indeed, linguists themselves sometimes lament our relative ignorance of aspects of linguistic intuitions' etiology—cf. Schütze 1996; Goldrick 2011.) It is this challenge that the present chapter aims to address. (We return below to some other considerations Devitt raises; still others are addressed in Maynes & Gross 2013.)

It might illuminate why Devitt thinks there's a problem in the first place if one notes that he raises this challenge specifically for conceptions of linguistic modules according to which grammatical constraints are embodied in computational operations rather than explicitly represented. If grammatical constraints were explicitly represented, then—Devitt suggests—linguistic intuitions might be derived within the language module in a quasi-deductive fashion.

² Devitt has developed and defended his views in a large number of subsequent papers, which can be found on his webpage. See also his reply in this volume.

³ Context makes clear that Devitt uses the word ‘information’ here, not in the information-theoretic sense, but to indicate representational content. Henceforth, I use ‘content’ to avoid confusion. Devitt uses ‘modular’ in the Fodorean sense (Fodor 1983). Mentalism about linguistics does not require accepting all aspects of Fodorian modularity (cf., e.g., Collins 2004 for differences between Fodor and Chomsky on linguistic competence and modularity); and there are now a variety of conceptions of modularity on the market (e.g., Carruthers 2006). I will attempt to bracket these matters.

(Devitt assumes that the relevant intuitions are judgments of *grammaticality*. But in current practice, judgments of grammaticality are typically not sources of evidence, but rather reflective judgments made by *theorists* to *explain* judgments of acceptability and other sources of evidence—cf., e.g., Myers 2009. We return to this shortly.) Devitt’s challenge is raised in reply to those who reject the explicit representation of grammatical constraints, arguably the vast majority of researchers in the field. It asks how else such intuitions could arise in a way that affords the speaker privileged access to the linguistic facts, suggesting there may not be any other way. (Devitt rejects the Voice of Competence view as well for conceptions on which grammatical constraints *are* explicitly represented, albeit on other grounds.)

But more is at stake than just the source and epistemic status of linguistic intuitions. Devitt’s argument for his “modest” view is part of a larger argument against mentalist conceptions of linguistics. Recall that, according to mentalist conceptions of linguistics, linguistics is a branch of psychology, investigating mechanisms specifically implicated in language acquisition and linguistic behavior. According to Devitt, linguistics is not, or ought not to be, so conceived. Rather, its object is, or should be, linguistic reality: the facts about language, or specific languages, which exist, independently of any specific speaker, as conventions among populations (as opposed to as Platonic *abstracta a la* Katz 1981). Devitt thus endorses an E-language, rather than an I-language, conception of what linguistics is, or ought to be, about (Chomsky 1986). His view of linguistic intuitions fits into his larger argument as follows: *If* the Voice of Competence view best explained why linguistic intuitions can be evidence, that would supply a consideration in favor of the mentalist conception. But, argues Devitt, it doesn’t best explain it; so, it doesn’t supply such a consideration. Answering Devitt’s challenge thus speaks to this element of his abduction in favor of his anti-mentalist conception of linguistics.

3. Clarifying the options and locating the current proposal

In fact, matters are more complicated than deciding between the Voice of Competence view and Devitt’s “modest” alternative. These two views do not exhaust the possibilities, and indeed, in previous work, I have argued against both options. Briefly reviewing these arguments will help clarify the claims of the current chapter.

Against Devitt’s view, Culbertson & Gross (2009) argue that one doesn’t find the divergence in linguistic intuitions it predicts. Devitt maintains that linguistic intuitions are theory-laden and so can diverge across speakers with different relevant background beliefs, including different commitments concerning linguistic theories. Indeed, Devitt—far from worrying about confirmation bias—argues that linguists should prefer their own linguistic intuitions to those of native speakers who are naïve regarding linguistics; for, the better (more reliable) linguistic intuitions will be those of speakers with better theories. But we found a high degree of consistency among subjects with varying degrees of expertise in linguistics, ranging from subjects with none at all up to practicing syntacticians. This suggests that linguistic intuitions—at least of the sort we elicited—may be fairly stable across changes in relevant background beliefs and experience and thus are not theory-laden in a way or to a degree that matters to linguistic inquiry. They may rather reflect their pre-judgmental etiology to a particularly robust degree.⁴

⁴ Devitt (2010) replies and Gross & Culbertson (2011) responds. ‘Reliable’ is used here not in the psychologist’s sense of being consistently produced in similar circumstances, but in the philosopher’s sense of tending to be accurate (as with a reliable thermometer)—what

On the other hand, Maynes & Gross (2013) argue *inter alia* against the Voice of Competence view—or at least reject the idea that mentalists should see themselves as committed to it. Recall that Devitt builds into his characterization of the Voice of Competence view that the language faculty itself outputs the *content* of the intuition (henceforth, the “content requirement”). But there is nothing about mentalism that requires this. Consider the judgment that some string is unacceptable. Mentalists needn’t commit themselves to the view that the language faculty itself outputs a state with the content *That string is unacceptable*. It can suffice that the parser fails to assign a structural description to the string and that the *absence* of a parse can in turn play a causal role in the process that leads the speaker to judge that the string is unacceptable.

The inclusion of the content requirement stems from Devitt’s emphasis on speakers’ privileged access to linguistic facts. For, if the language module supplies the content of linguistic intuitions, that might explain the source of this privilege. Recall, however, that judgments of grammaticality (as opposed to, for example, acceptability) are not, or are no longer, typical of the meta-linguistic judgments linguists rely on as evidence. Mentalists, in relying on the kinds of linguistic intuitions they in practice do, thus need not assume that speakers have privileged access to whether strings are grammatical. (Perhaps speakers have defeasible privileged access regarding acceptability.) Mentalists need only maintain that linguists’ theorizing involves an abduction from linguistic intuitions—and any other available considerations—to claims about a language faculty. (Thus might they, for example, elicit acceptability judgments under varying conditions and with varying stimuli, intending to control for alternative explanations.) This does not require that speakers have privileged access regarding the ground or causal source of their judgments—in particular, privileged access to why they judge a sentence (un)acceptable. Indeed, sentences can be unacceptable for any number of reasons. To take a classic example: multiply center-embedded sentences may be judged unacceptable owing to memory limitations, instead of a grammatical violation.

If linguistic intuitions are not theory-laden in the way Devitt expects, and if mentalists may reject the content requirement, then the positions Devitt discusses are not exhaustive. Thus, with Devitt’s alternative rejected in Culbertson & Gross (2009), Maynes & Gross (2013) defend a mentalist conception of linguistic intuitions *sans* the content requirement. This conception rejects as well the idea that a mentalist conception of linguistic intuitions’ evidential status requires that speakers possess privileged knowledge regarding grammaticality, while allowing that the special role grammaticality constraints can play in the generation of linguistic intuitions may enable those intuitions to serve as evidence for those constraints, in a manner relatively unaffected by changes in relevant belief and expertise.

Against this background, we can clarify the aims of the present chapter. The suggestion bruited above that a failure to parse can cause a judgment of unacceptability is a rather bare etiological claim, even if “sufficient unto the day” in the context of Maynes & Gross’ (2013) response to Devitt. In what follows, I buttress this reply by developing further suggestions concerning the etiology of linguistic intuitions. In particular, I suggest that error signals generated by monitoring mechanisms may play a role in some cases. I also suggest, more briefly,

psychologists would call validity. Note that, although relative expertise in linguistics did not matter in our experiment, one group—those with no formal exposure to the mind-brain sciences at all—was an outlier. Culbertson & Gross (2009) hypothesize a deficiency in task knowledge.

that in other cases the intuition's etiology may amount to little more beyond that of comprehension itself.

Interestingly, these further suggestions provide some grounds for entertaining a stronger thesis than I previously defended. For, although mentalism *per se* needn't build-in the content requirement, the error signal story, as we shall see, may allow the content requirement to be satisfied, at least by some intuitions—similarly, in some cases, for the comprehension account. The Voice of Competence view—or something like it (see below)—may thus be true after all, at least in those cases! But it's important to note that this is indeed a further claim: one can parry Devitt's etiological challenge without endorsing satisfaction of the content requirement.

Devitt might reply that it is *essential* to the Voice of Competence view, as he conceives it, that speakers' have privileged access to whether strings are grammatical. If so, satisfaction of the content requirement doesn't suffice for the Voice of Competence view, even if it provides for a view that is otherwise like it. Likewise, it's possible Devitt sees his etiological challenge as presupposing a mentalist endorsement of speaker's privileged access concerning grammaticality. If so, our reply to the etiological challenge *sans* this presupposition is a reply to a variant of Devitt's challenge, one suggested by it and worth addressing. But, in considering just what content error signals may have, we will also mention the even more speculative possibility that some error signals have content more specifically about grammaticality. The error signal story may thus even provide resources for someone attracted to a Voice of Competence view with some such privileged-access component built-in. (Some parallel indications are marked but not developed for comprehension cases.) Again, this would be a further claim, one that would go beyond maintaining that error signals have an etiological role and also beyond adding that they enable satisfaction of the content requirement.⁵

4. Candidate monitoring mechanisms

There's no consensus regarding the mechanisms involved in monitoring language use (for a brief survey, see Nozari & Novick 2017, to which this section's summary is indebted). But that some such mechanisms are involved in the prevention, detection, and correction of linguistic errors is a widespread view; and the correct details do not matter for my main point. Nonetheless, it's worth indicating some of the more specific extant ideas, both for the sake of concreteness and to underscore that my suggestion isn't *ad hoc*, but rather adverts to on-going, independently-motivated theorizing. (In that sense, my suggestion is thus also modest.) In addition, though the mere existence of error signals generated by monitoring mechanisms might suffice for my reply to Devitt's challenge, the details do matter for some more specific questions flagged below.

⁵ Henceforth, I use the 'Voice of Competence' label for the view that linguistic intuitions are the product of a modularized language faculty that alone delivers the relevant content to mechanisms responsible for judgment; and I take up the question of privileged access as a possible further requirement rather than building it in. It's of course less important which positions are allowed the label, so long as the positions themselves are clear. Note that Rey's (this volume; forthcoming) mentalist defense of the Voice of Competence view involves dropping the content requirement. He thus rejects Devitt's *characterization* of the view, whereas Maynes & Gross (2013) deploy Devitt's characterization (after all, it's his term) in rejecting Devitt's *ascription* of the view to the mentalist.

Initial ideas in this area were developed in theorizing about monitoring for speech *production* errors, and so we start with some examples of these. Production monitoring might seem not directly relevant to our topic, since linguistic intuitions are elicited in response to presented stimuli. But monitoring mechanisms have been posited in comprehension as well, and, according to some, monitoring mechanisms in production and comprehension are intimately related (Pickering and Garrod 2013).

Maynes & Gross (2013) cite Levelt's (1983; 1993) perceptual loop theory, according to which we monitor our production via comprehension. The basic idea is simply that we listen to what we ourselves say. Evidence for this includes that blocking auditory feedback with ambient white noise negatively affects our ability to catch production errors (Oomen et al. 2005). But it seems that this is not the only, perhaps not even the central, mechanism: among aphasic patients there is a double dissociation between comprehension and self-speech error detection (reviewed in Nozari et al. 2011).

An alternative view hypothesizes that an efference copy of the motor command is sent to a forward model that generates, for checking, an expectation concerning future states of motor and/or sensory systems (Tian & Poeppel 2010). This temporally more plausible approach applies to linguistic production a widely-held view of motor control more generally (Wolpert 1997; Shadmehr 2010). It's less clear, however, how classic versions of such views work or are motivated for higher-level, more abstract linguistic features such as syntactic structure, since such features are "up-stream" from motor commands.

Conflict models (Nozari et al. 2011) do not require efference copies. On such views, what is monitored is the comparative activation level of candidate linguistic representations (lexical items, phonemes, etc.) with a conflict signal produced when the difference in activation is insufficient for there to be a clear winner. In production, this might occur, for example, when /b/ and /c/ both get activated, to a sufficiently close degree, when you're trying to say 'cat'. In comprehension, an ambiguous signal, for example, might likewise lead to competing candidate representations with insufficiently differentiated activation levels.

Finally, there are many models of comprehension that incorporate prediction (Kutas et al. 2011). For example, an incremental parser, governed by various grammatical constraints, may generate expectations regarding syntactic features of lexical items to come. A mechanism might then monitor discrepancies between what is expected and the in-coming signal. (Forward models are examples of monitoring mechanisms incorporating prediction on the production side. See Pickering & Garrod 2013 for an attempt to integrate prediction in production and comprehension.)

Research in this area is active and on-going. But we needn't place bets. We can consider the consequences should some such model pan out.

5. Error signals and linguistic intuitions

Whatever the model, suppose problems generate an error signal. Consider, for instance, a failure to parse owing to ungrammaticality. Perhaps, on a conflict model, no structural description is activated at a level sufficiently greater than the rest to "win" the competition among candidates; or, with a predictive parser, perhaps the signal's completion fails to meet the expectations generated by previous material—perhaps even after attempted reanalysis. If monitoring mechanisms in such situations generate an error signal, then we are not limited to suggesting, as Gross & Maynes (2013) did, that a failure to parse yields an *absence* of output. The parser may

issue an error signal.⁶ The error signal, if it can in turn play a role in generating a judgment of unacceptability, enables us to elaborate our reply to Devitt's etiological challenge.

We can bolster the suggestion by noting various features such signals can have that mesh with features commonly associated with judgments of unacceptability. First, there's negative valence. It's an *error* signal after all. Likewise, subjects may express their negative judgment of a string by saying that it sounds *bad*. (Cf. Pietroski's (2008) use of 'yucky'—or 'icky,' another technical term I recently encountered in a linguist's talk.) Second, error signals can have motivational force—corrective in production, corrective or aversive when interpreting others. (Maynes & Gross 2013 note the possible connection to social cognition and in-group/out-group identification.) Linguistic intuitions are typically divorced from actual use, concerning rather what one could or would say or understand; but they still may be associated with an off-line or dispositional motivational force. Third, error signals might suggest the violation of a norm—indeed, perhaps they are a source of linguistic prescriptivism. Linguistic intuitions likewise may be associated with a sense of wrongness. Fourth, error signals may be graded—that is, they can come in various strengths. Defeated probabilistic predictions, for example, can be associated with varying degrees of surprisal. Similarly, the gradedness of linguistic intuitions has long been noted, whether in linguists' use of varying marks (*, **, ?, *?) to record their own judgments or in the graded results of formally collected judgments ranked on a scale.

Finally, each of these features may have associated phenomenology: a *felt sense* of badness, motivation, and/or norm-violation, of some particular strength. Often they may not, or not to a noticeable degree—otherwise the signals generated by the prevalent disfluency in ordinary speech might get in the way of conversational flow. But perhaps they may if attention is appropriately deployed—as it may well be when one is a subject performing an acceptability judgment task. Conscious or not, the signal may play a causal role in the generation of a linguistic intuition. Likewise, either way they may be relevant to the judgment's epistemic status. But noting the possibility of error signals rising to consciousness lends some measure of introspective plausibility to the proposal. Moreover, states capable of consciousness arguably can play a particular sort of epistemic role that in-principle unconscious states cannot: they can serve as justifications for the epistemic agent (as opposed to serving at best as warrants unavailable to the agent)—in the case at issue, justifications for the person forming meta-linguistic judgments.⁷

6. Error signals as the Voice of Competence?

The parallels between error signals and judgments of unacceptability do not *establish* that error signals play a role in the generation of linguistic intuitions, but they help render it a plausible hypothesis worth exploring. Even harder to establish is the more speculative claim that such signals might satisfy the content requirement. Nevertheless, the possibility is worth serious consideration, and we can provide some motivation.

⁶ Indeed, a possibility—not the only possibility (here is a place the details do matter)—is that the absence of a parse, given the activation of the parser by the string, *itself* plays this very role. The monitoring system may be so constructed so as to construe the absence (given the cueing up of the parser by a language-like stimulus) as a signaling of error. (This would not allow, however, for gradedness.) Alternatively, the absence may *partially* constitute the error signal—or the error signal may be a completely distinct state.

⁷ The nature and types of justification and their relation to conscious access is too large and controversial a topic to develop here. For some discussion and pointers, see Pappas (2017).

The error signal is a *signal*. It functions to deliver information to monitoring systems concerning what is occurring in language-related mechanisms, so as to initiate repair, reinterpretation, or some other corrective measure or appropriate response (perhaps including asking for clarification—cf. section 8 below). Moreover, the states it functions to deliver information about are representational states, concerning, for example, the structure of the presented string. Given the signal's functional role—with these sorts of causes and with these sorts of downstream effects—it is a natural thought that the signal itself might have representational content. Whether that thought is correct depends on the criteria for a mental state's bearing content, a highly contentious matter. But a leading view is that assigning a mental state content on the basis of what causes it and how it is consumed is warranted when doing so yields explanatory illumination (e.g., Shea 2007; 2012). Arguably, this condition is met by error signals: our understanding of what a monitoring system does involves our seeing the system as capable of being informed of a problem, and it is the error signal that does the informing. Note, moreover, that it seems in-principle possible for a monitoring system to itself be in error (see fn. 10), so that a signal can *misrepresent* the state of the system being monitored. The possibility of misrepresentation is often considered a necessary condition of intentionality (Dretske 1986).

Just what would be the content of an error signal? There are various natural candidates. But all seem to enable the satisfaction of at least some version of the content requirement. Consider again a failed parse and a subsequent judgment of unacceptability. The candidate gloss that most obviously would make the case is *This string is unacceptable*—or perhaps such close relatives as *This string has unsurmountable problems* or *Something is wrong with this string* (or variants that refer to the utterance). If the error signal's content is thus glossed, then clearly, at least for such intuitions, one could defend the Voice of Competence view, even with its content requirement. In these cases, speakers' judgments can reflect constraints built into a modularized language faculty that outputs the content contained in the judgment.

That's the clearest, best case, so far as satisfying the content requirement goes. But a consideration of alternatives only lends further support. Our opening gloss makes reference to the string. But perhaps the state is more purely interoceptive, indicating only how things are with the subject. If so, a better gloss might be: *I have—or this mechanism has—encountered unsurmountable problems*. It's not obvious that the error signal itself makes no reference to the string as the locus of the problem; but, even if this is so, the content requirement may be satisfied. For it remains the case that linguistic processes yield a representation of the stimulus (minimally, a representation as of these phonemes and lexemes in this order). Plausibly this representation, *together* with the error signal, can supply the content of the unacceptability judgment.

But perhaps the error signal indicates, not that the problem has become unsurmountable, but only that problems persist: *There is (still) a problem*. The further information that there are no options left or that further efforts are not worth the cost might require further states—for example, perhaps a monitoring state that indexes effort expended. Yet it would remain the case that error signals not only play an etiological role in the generation of the unacceptability judgment, but also supply content to monitoring states that (perhaps in concert with representation of the string itself) supply the content of the judgment. The content requirement would thus, again, still be satisfied. Similar remarks apply to the suggestion that error signals have imperatival content (cf. Klein's (2007) imperatival view of the content of pain states)—perhaps glossed as *Appropriately respond to this problem!*. (Suppose they *only* have such

content. If this is just *additional* content, no issue is raised.) The persistence of such states can play a role in generating states that indicate that there is nothing more worth doing.

An intriguing—still more speculative—possibility is that the signal carries more specific information about the source or nature of the problem. Recall that we motivated the assignment of content to error signals by reflection on their functional role: to inform a monitoring system of a processing problem so as to respond appropriately. But, just as a fire department needs information concerning where an alarm is coming from if it is to do anything about it, so may the monitoring system require more specific information about the signal's source or the kind of problem. Suppose the signal itself carries that information. In our example, given the source of the problem (a failure to parse owing to embodied grammatical constraints), one might argue that the signal has a content that reflects this specificity: not just that the string is unacceptable, but that it is *ungrammatical* (as opposed to *unpronounceable*, or *pragmatically unacceptable*, etc.). Indeed, perhaps in principle an error signal could indicate the nature of the violation even *more* specifically—e.g., *subjacency violation*. Some such suggestion might be particularly tempting if the monitoring mechanisms that consume such signals are domain-general, since a domain-specific monitoring system might not need to sort among signals' varying causes. Whether the monitoring systems implicated in language use are domain-general or domain-specific (or some of each) is an unresolved empirical question, one where the details matter (Nozari & Novick 2017—but see Dillon et al. 2012 for some ERP evidence that error signals can encode their cause).⁸

Suppose error signals do have more specific content. It would not follow that judgments caused by or based on them also have this more specific content. It's one thing for a state to have some content and another for the subject to conceptualize it as such or even to be able to. Compare the representation of color features. The visual system might represent some object as having certain color features. But the color categories available to conceptual systems may be much coarser than those available to vision. (This could be so even if color perception is categorical, though see Witzel (forthcoming) for arguments that it isn't.) Similarly, even if error signals represent more specifically the source or kind of error, it's a further question whether the judgments subjects make on their basis do so as well. Thus, in principle an error signal with the content *That string is ungrammatical* could cause and be the basis for a judgment with the more generic content *That string is unacceptable*. If so, should we say the content requirement is satisfied? The signal and judgment do not have the *same* content. But the signal's content warrants the judgment's in a way that seems analogous to how perceptual contents more generally can warrant perceptually-based judgments with different but closely related contents (cf. Peacocke 2004 on canonical conceptualizations). I suggest that this should suffice for a content requirement worth preserving.

The possibilities entertained so far are all consistent with denying speakers privileged access to grammaticality facts. Even if the signal's content is supplied by linguistic competence (whatever level of specificity that content may have), it doesn't follow that the speaker knows that it is. We generally have unreliable introspective access to the causal source of our judgments, and content not conceptualized as such is not available to judgment. Thus, the

⁸ In support of more specific content, one might also advert to subjects' often being able to provide some indication of where the problem lies and to suggest fixes. It's unclear, however, whether this on its own favors more specific error-signal content over a Devitt-like specification in judgment, in light of a less specific error signal.

“voice” of competence may be “heard”—in the sense of its causing and supplying a basis for a judgment—without the subject knowing that it is *competence* that’s “speaking.” Theorists would still need to engage in the hard work of inference to the best explanation—as, for example, is the case in trying to sort out whether the judgment data concerning binding phenomena reflect syntactic or pragmatic constraints (e.g., Chomsky 1981 and Reinhardt 1983).

But suppose we contemplate the possibility that error signals about ungrammaticality can cause and support linguistic intuitions with corresponding content—meta-linguistic judgments with the content *That is ungrammatical*. This could help provide grounds—at least regarding some intuitions—for a Voice of Competence view that also requires privileged access to grammaticality facts. (Privileged access might require reliability as well. Perhaps cases such as sentences with multiple center-embeddings could be deemed outliers.) Thus might one try to defend a limited application of the full view Devitt opposes. The application, however, would be quite limited indeed if grammaticality judgments are in fact not often invoked as evidence, as opposed to in explanation. Still, perhaps subjects may in some cases feel and be able to express that a problem seems grammatical or in some sense structural. Linguists, with their theoretical expertise and fuller array of concepts, may even achieve further specificity: perhaps a particular case may *feel* like a subadjacency violation.⁹ That said, methodological caution suggests that we distinguish more solid evidence from what may be merely theoretical hunches. But we needn’t dismiss the *possibility* of such linguistic intuitions out of hand, even if they are not given much weight or play in practice. (We mark below some other avenues of possible support.)

The question of specificity bears on a challenge Devitt raises to the Voice of Competence view. He asks: If our competence “speaks” to us, why does it not say more? Why do linguists not find themselves with a broader array of speakers’ intuitions to draw on—specified in the linguist’s language of c-command, heads, etc.? Given our remarks above, we can divide the question in two: Why doesn’t the language module deliver more specific content (if it doesn’t)? And, supposing it does, why is it that speakers typically can’t non-reflectively conceptualize that content in judgment? It’s not clear that the mentalist’s inability to answer at present would be particularly problematic. But the questions are interesting nonetheless. One way to take them is as design questions, so that we might ask what purpose would be served by having things otherwise than they are, and whether things overall would be in some sense better if things were that way. We posed the question of domain-generality vs. domain-specificity of monitoring mechanisms in such a functional way. It’s important to bear in mind, however, first, that speakers may not achieve optimality regarding the generation of linguistic intuitions (there’s certainly no reason to expect our capacities to be optimal for enabling successful linguistics!); and, second, that our capacity for yielding intuitive meta-linguistic judgments may be a by-product—for example, of monitoring systems that may operate in large part unconsciously. (Questions of function briefly recur in section 8.)

7. Other linguistic intuitions, other sources

⁹ It is not enough that one possess the relevant concepts, though that is necessary and is a further—and later—achievement than acquiring language itself (Hakes 1980). If the concepts are to contribute content to intuitive judgments, one must also be able to deploy them unreflectively in conceptualizing one’s experience in response to presented strings. Acquiring syntactic concepts in a linguistics course may not suffice for this.

We have focused on intuitive judgments of unacceptability (remarking as well on the possibility of intuitive judgments of ungrammaticality). But what of other linguistic intuitions? The various judgments that have been called ‘linguistic intuitions’—for example, those with which this chapter starts—may not form a natural kind. Different kinds of linguistic intuition may have different etiologies and require different accounts of why they are evidence. In principle, Devitt’s view could be right about some, wrong about others; and where it’s wrong, a Voice of Competence view incorporating the content requirement may likewise be right for some, but not for others. Let’s consider some cases.

What of judgements that a string *is* acceptable? There is an obvious asymmetry here. In such cases, there is presumably no error signal to play an etiological role, so the content of the intuition would not seem to be the content of some output of the language faculty. A possibility is that, while error signals lead to judgments of unacceptability, their *absence* leads to judgments of acceptability. As for satisfaction of the content requirement: well, the suggestion was restricted to *some* linguistic intuitions—perhaps judgments of acceptability are not among them. But an alternative to invoking merely the absence of error signals would instead extend the account to encompass positive, “non-error” signals. Perhaps monitoring mechanisms, when functioning properly, should be construed as always vigilant and thus always in effect receiving a signal. The absence of an error signal would then be itself a signal of proper functioning—perhaps, further, a state with content to the effect that: *This string is acceptable*. If so, the content requirement could be satisfied after all, and the Voice of Competence view extended to judgments of acceptability.¹⁰

There is, however, another possibility to consider. Perhaps what’s causally and epistemically most significant for judgments of acceptability is neither the absence of error signals nor the presence of a positive ‘no error’ signal, but rather the speaker’s having comprehended what was said. As we saw, Chomsky incorporates comprehensibility into his gloss on acceptability. Indeed, having noted this possibility, one might challenge the need for error signals in accounting for *unacceptability* judgments as well: perhaps it suffices that one *not* comprehend the string. But it wouldn’t follow that error signals do not play a role. Even if one factor would suffice, both might be present. (The same source could cause both the signal and (ultimately) the failure to comprehend.) In any event, adverting to comprehension cannot serve as a complete account. For there are strings that are readily comprehended but also readily judged unacceptable, such as ‘She seems sleeping.’ Chomsky’s gloss doesn’t require *only* comprehensibility.¹¹

¹⁰ One might worry about cases where an *unacceptable* sentence is judged acceptable (at least at first or unreflectively)—as with plural attraction (‘The key to the cabinets are on the table’) and, more persistently, comparative illusions (‘More people have been to France than I have’). Why doesn’t the error signal yield an unacceptability judgment here? If it doesn’t, this needn’t be an objection: error signals might not always succeed in generating appropriate linguistic intuitions; and our judgments of (un)acceptability needn’t be infallible. But, in any event, the question’s presupposition that there *is* an error signal in such cases may be mistaken. Perhaps the question should be: how do such strings get pass the parser without generating an error signal? Just as visual illusions provide insight into the fine structure of visual processing, such cases can illuminate the quirks of linguistic processing (Wagers et al. 2009; Wellwood et al. 2018).

¹¹ Another, more contentious reply to the suggested alignment of comprehension and acceptability would invoke alleged cases of uncomprehended strings judged acceptable, such as

That said, we can certainly assign comprehension or the lack thereof a significant role in the generation and grounds of (un)acceptability judgments, without challenging the error signal suggestion. The error signal account needn't exhaust the etiology for linguistic intuitions generally or for any specific kind of linguistic intuition. And allowing a role for (in)comprehension is not in tension with our aim of elaborating on Maynes & Gross' (2013) reply to Devitt's etiological challenge. Indeed, it meshes with it, since the fact of (in)comprehension also yields data for the linguist's inference to best explanation, data that can be relatively robust to variation in background theory. Note that, contrary to Devitt's complaint, we *do* have some idea in this case how one gets from embodied grammatical constraints to the linguistic intuition. Of course, major gaps exist in our knowledge of utterance understanding (that's what much of linguistics is about), but there's no further special gap introduced by linguistic intuitions of this sort.

Adverting to comprehension in such cases, however, responds to the etiological challenge without satisfying the Voice of Competence's content requirement. Devitt (2006b: 118), in raising the challenge, maintains that what's delivered to central systems is the "message"—i.e., bracketing delicate semantic/pragmatic issues, something like the content of what is said—not information that leads to arriving at it (recall that Devitt has in mind information about grammaticality). I would suggest that what's delivered is rather something like that the speaker said that P in uttering S (*mutatis mutandis* for other speech acts).¹² Either way, the content of the intuition itself is indeed not delivered. But yet we can see how the comprehension story could allow a simple transition from the *fact* of (in)comprehension to a judgment of (un)acceptability—*modulo* the contribution of an error signal.

Comprehension plays a significant role in other linguistic intuitions as well—for example, judgments of co-reference and truth-value judgments. But here it's not just a matter of *whether* one comprehends, but also of *what* one takes the content to be. For example, to answer whether the bolded terms in 'John's roommates met **him** at the restaurant' co-refer, one in effect reports whether one understood it to be John whom the roommates met at the restaurant. Of course, answering the question so formulated requires some meta-linguistic awareness (and more specifically in this case possession of the concept of co-reference), as does any meta-linguistic judgment. To that extent, the judgment goes beyond mere comprehension of what it said, though

lines from *Jabberwocky* or particularly inscrutable bits of philosophy (Rey this volume; forthcoming). To my knowledge, judgments concerning such cases haven't been investigated in a controlled experimental setting (though see Pallier et al. 2011 for neurolinguistic investigation). But, in addition, a potential counter-reply is that they are *sufficiently* understood, insofar as subjects construct a metalinguistic or deferential concept to handle the problematic open-class terms (*borogoves, whatever they are*)—cf. Higginbotham (1989). Incidentally, another possible use of such cases (if they are granted) might be to bolster the possibility of useable intuitive judgments of grammaticality (cf. Rey this volume; forthcoming).

¹² That the string and not just its content is delivered can be so even if the string is then more easily forgotten (Sachs 1967). Note that including the string in what's delivered as output to central systems raises the question of how the string is represented at this stage. To the degree that structural information is preserved, this again might provide resources for those who would defend the possibility of privileged access to more specific linguistic features. Whether—and if so, to what extent and in what ways—syntactic features are *perceptually experienced* (as phonemes, morphemes, order, etc. are) is a related, delicate question.

not much.¹³ (Parallel remarks apply to other judgments concerning what “readings” a subject gets.) For truth-value judgments it likewise matters, not just that one comprehends the sentence, but also what content one assigns it. In this case, the judgment goes beyond one’s comprehending the sentence in a further way: one must assess that content against some scenario. Having a capacity to do so in a certain range of cases, however, is arguably in part constitutive, or highly diagnostic, of one’s capacity for comprehending the sentence. (Of course, one must also exploit one’s grasp of the concept of truth to form the appropriate meta-linguistic judgment.)

Moreover, in these cases where what matters is what one takes the content to be, not only do we have some understanding of the intuition’s etiology (to the extent that we understand the etiology of comprehension), but also—insofar as the intuition amounts to little beyond comprehension—the content requirement is fulfilled and arguably the Voice of Competence is vindicated.¹⁴ Note in particular that Devitt’s characterization of linguistic competence’s role in generating *grammaticality* judgments—as supplying material for a central system response—seems inapt in these cases where the intuition just is, or goes just a bit beyond, comprehending what was said.

Do error signals play a role with these linguistic intuitions? In positive cases—where the terms do co-refer (where one gets the reading) or where what is said is judged true—what one understands the sentence to say seems much more significant, even if the absence of an error-signal (or the presence of positive non-error signals) is concomitant. But the matter is less clear with at least some negative cases. Suppose the terms do not co-refer, or one judges the sentence false. Does the subject’s understanding of what was said supply the resources necessary to arrive at these judgments without recourse to error signals? One way to suggest a role for error signals would implicate them in processes of belief formation and maintenance in general. Perhaps such processes involve a monitoring system for one’s epistemic states, where clashes between an entertained thought and other beliefs can yield an error signal (perhaps sometimes a conscious, pre-judgmental feeling of wrongness) that leads to the thought’s rejection—and perhaps such processes are implicated in our subject’s negative verdicts. The suggestion would require independent motivation (why think error signals are required, as opposed to operations that directly adjust beliefs in response to reasons?).¹⁵ In any event, these would not be error signals generated by specifically linguistic monitoring mechanisms.

¹³ One could eliminate the meta-linguistic element in this case by telling the subject that John’s roommates met him at the restaurant (using those words) and asking whether it was John whom they met. This task would not involve linguistic intuitions if one defines them to be meta-linguistic judgments. The difference does not seem to amount to much in this case. Note further that if comprehension involves delivering to central systems a content that relates what is said to the string uttered (as suggested above), then, depending on the details of what’s represented, this metalinguistic element may be included in comprehension, even if not in the expressed judgment.

¹⁴ More fully defending this would require unbracketing the delicate semantics/pragmatics issues. But see Sperber & Wilson (2002).

¹⁵ David Pereplyotchik rightly points out that an analogous charge could be raised against positing error signals in language monitoring. The question is simply which models are empirically justified.

But it might also be suggested that error signals generated from monitoring mechanisms specifically for language could play a role, at least in some cases. When one considers whether the bolded terms in ‘**Sally** gave **her** the book’ co-refer, one might not only consider how one naturally understands the sentence, but also attempt to find a reading where they do co-refer—all the more so if one’s asked whether they *can* co-refer. In such a case, one’s failed attempts might involve error signals (whether from a violated grammatical constraint, a mismatch between the parse and the semantic supposition, or a pragmatic violation) that ground one’s judgment. The possibility of a language-specific error signal is less obvious for truth-value judgments. But candidates could be truth-value judgments concerning the special case of strings that are “analytically” false in the sense that they can be known to be false just by an exercise of one’s linguistic competence (Pietroski 2002).

I have been suggesting that comprehension may play a much more significant role with some linguistic intuitions than do error signals. But there are also cases where it’s clear that comprehension (in the sense of understanding) plays no role whatsoever. Consider phonological intuitions, such as judgments as to whether a stimulus is a pronounceable, or a possible word, in one’s language. Comprehension plays no role in one’s assessment of /fant/ and /zgant/. The error signal account, on the other hand, neatly extends to such cases (though, since now the source of the signal is different, perhaps so is the signal’s content, if the content is specific). Again, an account that focuses on comprehension is incomplete.

Finally, consider sociolinguistic intuitions such as judgments of leave-taking (the last of our opening examples). The processes by which these are formed may be altogether different. One might try arguing that there’s a role for error signals or comprehension. Perhaps one attempts to simulate leave-taking in imagination and registers a feeling of wrongness with some candidates (though the feeling’s source may not be specifically linguistic); perhaps one “understands” some of the candidates as differing in their linguistically-encoded level of formality or politeness. But another possibility is that one forms an empirical judgment based on one’s memory of past usage. Such cases might fit Devitt’s model well.

Linguistic intuitions thus may differ in their etiologies, and a single intuition may have several language-specific bases. We have suggested error signals play an etiological role, but not that an account that adverts only to them is complete. Comprehension plays a significant role with some linguistic intuitions, and perhaps other bases matter as well—for example, indices and feelings of effort and fluency, briefly mentioned in Section 6 (cf. Luka 2005).

8. Why not intuitions elsewhere?

In considering the etiology of other linguistic intuitions, we mentioned the possibility of monitoring mechanisms for other aspects of cognition and behavior; and earlier we noted that forward models are common in accounts of motor control. The possibility that monitoring systems are ubiquitous—more specifically, monitoring that involves error signals—connects to another challenge Devitt raises for the Voice of Competence view. In some sense, he says, we have “embodied rules” for swimming, typing, and other skills. If “embodied rules” of language yield intuitions that provide fruitful evidence in linguistic theorizing, why don’t “embodied rules” for swimming etc. do the same? Adapting this question to our present discussion, we can ask: if such monitoring is ubiquitous, why don’t error signals concerning swimming etc. provide the basis for fruitful theorizing in those domains? I conclude by providing a reply.

First, judgments analogous to linguistic intuitions do play a central evidentiary role in various domains—most obviously in perception science. So, to that extent, linguistics is not a

special case. Indeed, even acceptability judgments more specifically (or something very much like them) are exploited in some domains, such as music cognition (Patel et al. 2008; Featherstone et al. 2013) and moral psychology (Cushman 2015).

But what about the sort of motor skills Devitt emphasizes? There is no *a priori* reason why intuitions could not supply fruitful evidence in these domains—and some instances can be found. For example, Ward & Williams (2003) had soccer players of different skill levels view video clips and then answer questions concerning to which players teammates could pass and which was most likely. We may grant, however, that judgments about exercises of a motor skill or their products are less commonly exploited as evidence concerning those skills. Why might this be so?

There are a variety of possible (non-exclusive) explanations. It could simply be that other methods have proven sufficiently fruitful to have occupied these fields so far. But there may also be differences that would explain why we should not expect intuitions ever to play a central role. Most relevantly to our topic, the usefulness of intuitions in linguistics, as opposed to swimming, may in part reflect the crucial role of a dedicated, more-or-less modular mechanism with its proprietary quirks, recursivity, and interface constraints. That is, intuitions may provide fruitful evidence for theorizing about (aspects of) language, but not about swimming, precisely because something like the Voice of Competence view is correct in this domain.

Other differences may also be relevant. For example, many skills arguably lack an analogue of comprehension, in two senses. First, for most skills, to exteroceptively perceive some exercise of it is arguably not itself to exercise the skill: watching someone swim is not itself swimming, perhaps contrary to some mirror-neuron theorists.¹⁶ Second, linguistic comprehension involves understanding what someone said. But, for many motor skills, their exercise does not involve understanding the actions of an agent, in particular understanding their (communicative) intentions. This is relevant to intuitions in that, as we've noted, for some linguistic intuitions, comprehending what was said just is, or is almost, the forming of the intuition itself. Finally, comprehension in this second sense is connected to the particular demands for coordination found in communication, which is among the primary functions that language subserves. To be sure, there is synchronized swimming. But arguably the specific forms coordination takes in language use (asking for clarification of an ambiguous or just hard-to-hear utterance, engaging in lexical negotiation) require or at least promote a capacity for meta-linguistic awareness that, perhaps as a by-product, renders language-users especially able to provide useful intuitions in this domain. Language-related education—most obviously, years of concentrated training in reading and writing—may likewise play a role (cf. Schütze 1996). This is not to deny that the exercise of other skills can involve coordination with others and the comprehension of intentions. But when they do, as with soccer, we've seen that judgment data may be useful after all. More generally, exercising different kinds of motor skill may make varying demands on cognitively-driven conscious control; Devitt's examples may simply fall on the less demanding end (cf. Montero 2016; Gregory et al. 2016).

Here and above, this exercise in empirically-motivated speculation deploys a fair number of 'maybe's and 'perhaps's. But, whether the suggestions pan out or not, they at least provide

¹⁶ Interoception in production may be in some sense an aspect of the exercise—and, again, in principle could serve as a source of evidence. But, as noted above, judgments formed on the basis of interoception in production are not analogous to typical linguistic intuitions, which are responses to externally-given language-like stimuli.

“the beginnings of a positive answer” to the questions we have posed, however transformed those questions may be from Devitt’s.

References

- Carruthers, Peter (2006). *The Architecture of the Mind*. Oxford: Oxford University Press.
- Chomsky, Noam (1965). *Aspects of the Theory of Syntax*. Cambridge: MIT Press.
- Chomsky, Noam (1981). *Lectures on Government and Binding: The Pisa Lectures*. Dordrecht: Foris Publications.
- Chomsky, Noam (1986). *Knowledge of Language*. New York: Praeger.
- Collins, John (2004). ‘Faculty Disputes’, *Mind & Language* 19: 503-33.
- Culbertson, Jennifer, and Steven Gross (2009). ‘Are Linguists Better Subjects?’, *British Journal for the Philosophy of Science* 60: 721-36.
- Cushman, Fiery (2015). ‘Punishment in Humans: From Intuitions to Institutions’, *Philosophy Compass* 10: 117-33.
- DePaul, Michael, and William Ramsey, eds. (1998). *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*. Lanham, MD: Rowman and Littlefield.
- Devitt, Michael (2006a). ‘Intuitions in Linguistics’, *British Journal for Philosophy of Science* 57: 481-513.
- Devitt, Michael (2006b). *Ignorance of Language*. Oxford: Oxford University Press.
- Devitt, Michael (2010a). ‘Linguistic Intuitions Revisited’, *British Journal for the Philosophy of Science* 61: 833-65.
- Devitt, Michael. this volume.
- Dillon, Brian, Andrew Nevins, Allison Austin, and Colin Phillips (2012). ‘Syntactic and Semantic Predictors of Tense in Hindi: An ERP Investigation’, *Language and Cognitive Processes* 27: 313-44.
- Dretske, Fred (1986). ‘Misrepresentation’, in R. Bogdan (ed.), *Belief: Form, Content, and Function*. Oxford University Press, pp. 17-36.
- Featherstone, Cara, Catriona Morrison, Mitch Waterman, and Lucy MacGregor (2013). ‘Syntax, semantics or neither? A case for resolution in the interpretation of N500 and P600 responses to harmonic incongruities’, *PLoS ONE* 8: e76600. doi:10.1371/journal.pone.0076600
- Fodor, Jerry (1983). *The Modularity of Mind*. Cambridge: MIT Press.
- Goldrick, Matthew (2011). ‘Utilizing Psychological Realism to Advance Phonological Theory’, in J. Goldsmith, J. Riggle, and A. Yu (eds.), *Handbook of Phonological Theory* (2nd ed.). Oxford: Blackwell, pp. 631-60.
- Gordon, Peter, and Randall Hendrick (1997). ‘Intuitive Knowledge of Linguistic Co-Reference’, *Cognition* 62: 325-70.
- Gregory, Emma, Michael McCloskey, Zoe Ovans, and Barbara Landau (2016). ‘Declarative memory and skill-related knowledge: Evidence from a case study of amnesia and implications for theories of memory’, *Cognitive Neuropsychology* 33: 220-40.
- Gross, Steven, Jennifer Culbertson (2011). ‘Revisited Linguistic Intuitions’, *British Journal for the Philosophy of Science* 62: 639-56.
- Hakes, David (1980). *The Development of Metalinguistic Abilities in Children*. New York: Springer-Verlag.
- Higginbotham, James (1989). Knowledge of Reference. In A. George (ed.), *Reflections on*

- Chomsky. Oxford: Blackwell.
- Katz, Jerrold (1981). *Language and other abstract objects*. New York: Rowman & Littlefield.
- Klein, Colin (2007). 'An Imperative Theory of Pain', *The Journal of Philosophy* CIV: 517-532.
- Kutas, Marta, Katherine DeLong, and Nathaniel Smith (2011). 'A Look around at What Lies Ahead: Prediction and Predictability in Language Processing', in M. Bar (ed.), *Predictions in the Brain*. Oxford: Oxford University Press, pp. 190-207.
- Labov, William (1996). 'When Intuitions Fail', *Chicago Linguistics Society: Papers from the Parasession on Theory and Data in Linguistics* 32: 76-106.
- Levelt, Willem (1983). 'Monitoring and self-repair in speech', *Cognition* 14: 41-104.
- Levelt, Willem (1993). *Speaking: From Intention to Articulation*. Cambridge: MIT Press.
- Luka, Barbara (2005). 'A Cognitively Plausible Model of Linguistic Intuitions', in S. S. Mufwene, E. Francis, and R. Wheeler (eds.), *Polymorphous Linguistics: Jim McCawley's Legacy*. Cambridge, MA: MIT Press, pp. 479-502.
- Matthews, Robert. ms. 'Linguistic Intuition: An exercise in linguistic competence'.
- Maynes, Jeffrey, and Steven Gross (2013). 'Linguistic intuitions', *Philosophy Compass* 8: 714-30.
- Montero, Barbara (2016). *Thought in Action*. Oxford: Oxford University Press.
- Myers, James (2009). 'Syntactic judgment experiments', *Language and Linguistics Compass* 3: 406-23.
- Nozari, Nazbanou, Gary Dell, and Myrna Schwartz (2011). 'Is comprehension the basis for error detection? A conflict-based theory of error detection in speech production', *Cognitive Psychology* 63: 1-33.
- Nozari, Nazbanou, and Jared Novick (2017). 'Monitoring and control in language production', *Current Directions in Psychological Science* 26: 403-10.
- Oomen, Claudy, Albert Postma, and Herman Kolk (2005). 'Speech monitoring in aphasia. Error detection and repair behavior in a patient with Broca's aphasia', in R. J. Hartsuiker, R. Bastiaanse, A. Postma, and F. Wijnen (eds.), *Phonological encoding and monitoring in normal and pathological speech*. New York, NY: Psychology Press, pp. 209-225.
- Pallier, Christophe, Anne-Dominique Devauchelle, and Stanislas Dehaene (2011). 'Cortical representation of the constituent structure of sentences', *PNAS* 108: 2252-7.
- Pappas, George (2017). 'Internalist vs. Externalist Conceptions of Epistemic Justification', in E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition), URL = <https://plato.stanford.edu/archives/fall2017/entries/justep-intext/>.
- Patel, Aniruddh, John Iversen, Marlies Wassenaar, and Peter Hagoort (2008). 'Musical syntactic processing in agrammatic Broca's aphasia', *Aphasiology* 22: 776-89.
- Peacocke, Christopher (2004). *The Realm of Reason*. Oxford: Oxford University Press.
- Pickering, Martin, and Simon Garrod (2013). 'An integrated theory of language production and comprehension', *Brains & Behavioral Sciences* 36: 329-47.
- Pietroski, Paul (2002). 'Small verbs, complex events: Analyticity without synonymy', in L. Antony and N. Hornstein (eds.), *Chomsky and His Critics*. Oxford: Blackwell: 179-214.
- Pietroski, Paul (2008). 'Think of the Children', *Australasian Journal of Philosophy* 86: 657-69.
- Reinhart, Tanya (1983). *Anaphora and Semantic Interpretation*. Chicago: University of Chicago Press.
- Rey, Georges. this volume. 'A Defense of the Voice of Competence'.
- Rey, Georges. forthcoming. *Representation of Language: Foundational Issues in a Chomskyan*

- Linguistics*. Oxford: Oxford University Press.
- Sachs, Jaqueline (1967). 'Recognition Memory for Syntactic and Semantic Aspects of Connected Discourse', *Perception and Psychophysics* 2: 437-42.
- Schütze, Carson (1996). *The Empirical Base of Linguistics*. Chicago: University of Chicago Press.
- Schütze, Carson (2011). 'Linguistic Evidence and Grammatical Theory', *Wiley Interdisciplinary Reviews: Cognitive Science* 2: 206-21.
- Shadmehr, Reza, Maurice Smith, and John Krakauer (2010). 'Error correction, sensory prediction, and adaptation in motor control', *Annual Review of Neuroscience* 33: 89-108.
- Shea, Nicholas (2007). 'Consumers need information: supplementing teleosemantics with an input condition', *Philosophy and Phenomenological Research* 75: 404-35.
- Shea, Nicholas (2012). 'Reward prediction error signals are meta-representational', *Nous* 48: 314-41.
- Sperber, Dan, and Deirdre Wilson (2002). 'Pragmatics, modularity and mindreading', *Mind & Language* 17: 2-23.
- Sprouse, Jon (2018). 'Acceptability judgments and grammaticality, prospects and challenges', in N. Hornstein, C. Yang, and P. Patel-Grosz (eds.), *Syntactic Structures after 60 Years: The Impact of the Chomskyan Revolution in Linguistics*. Berlin: Mouton de Gruyter, pp. 195-224.
- Sprouse, Jon, and Carson Schütze. forthcoming. 'Grammar and the use of data', in G. Popova, B. Aarts, and J. Bowie (eds.), *Oxford Handbook of English Grammar*. Oxford: Oxford University Press.
- Tian, Xing, and David Poeppel (2010). 'Mental imagery of speech and movement implicates the dynamics of internal forward models', *Frontiers in Psychology*. doi: 10.3389/fpsyg.2010.00166
- Wagers, Matthew, Ellen Lau, and Colin Phillips (2009). 'Agreement attraction in comprehension: Representations and processes', *Journal of Memory and Language* 6: 206-37.
- Ward, Paul, and A. Mark Williams (2003). 'Perceptual and cognitive skill development in soccer: The multidimensional nature of expert performance', *Journal of Sport & Exercise Psychology* 25: 93-111.
- Wellwood, Alexis, Roumyana Pancheva, Valentine Hacquard, and Colin Phillips (2018). 'The anatomy of a comparative illusion', *Journal of Semantics* 35: 543-83.
- Witzel, Christoph. forthcoming. 'Misconceptions about colour categories', *Review of Philosophy and Psychology*.
- Wolpert, Daniel (1997). 'Computational approaches to motor control', *Trends in Cognitive Sciences* 1: 209-16.

Acknowledgments

This chapter stems from a brief aside in Maynes & Gross (2013), the remains of a paragraph excised for reasons of space—my thanks to Jeff Maynes. Material related to this paper was presented at the Johns Hopkins Cognitive Science Brown Bag; the Norwegian Summer Institute on Language and Mind; the Southern Society for Philosophy and Psychology; the Society for Philosophy and Psychology; and the Aarhus Workshop on Linguistic Intuitions, Evidence, and

Expertise. For help of various sorts, I thank Sam Schindler, Megan Hyska, Fritz Newmeyer, Sam Featherston, Colin Phillips, Jon Sprouse, Carlos Santana, Lila Gleitman, Mike McCloskey, David Pereplyotchik, Tal Linzen, Nick Allott, and OUP's anonymous referee (my apologies for omissions). Special thanks to Georges Rey and Bonnie Nozari. Extra special thanks to Michael Devitt for his support and encouragement. The first person to whom I floated some of these ideas, back in 2012, was Akira Omaki, whose enthusiasm and kindness are much missed.

Requested biographical paragraph:

Steven Gross is Professor of Philosophy at Johns Hopkins University, with secondary appointments in Cognitive Science and in Psychological and Brain Sciences. He has published on a variety of topics in philosophy of language, philosophy of mind, and the foundations of the mind-brain sciences. His most recent publications have focused on perceptual consciousness and on cognitive penetration. Current projects include “anti-Bayesian” updating in vision and whether linguistic meaning is perceived or computed post-perceptually.