Routledge
Taylor & Francis Group

Inquiry

An Interdisciplinary Journal of Philosophy

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/sinq20

Encapsulation, inference and utterance interpretation

Nicholas Allott

To cite this article: Nicholas Allott (13 Oct 2023): Encapsulation, inference and utterance interpretation, Inquiry, DOI: 10.1080/0020174X.2023.2267084

To link to this article: https://doi.org/10.1080/0020174X.2023.2267084

INQUIRY
https://doi.org/10.1080/0020174X.2023.2267084

Routledge
Taylor & Francis Group

Encapsulation, inference and utterance interpretation

Nicholas Allott

Department of Literature, Area Studies and European Languages, University of Oslo, Oslo, Norway

ABSTRACT

While utterance interpretation is standardly understood as context-sensitive inference, there are prima facie reasons to doubt it is fully unencapsulated. First, utterance interpretation is normally fast and automatic, while it has been argued that unencapsulated processes are slow and reflective. Secondly, certain illusions appear to show that the processing of utterances is unrevisable in light of further information. I argue that these reasons are not conclusive, and utterance interpretation relies on indefinite tracts of background knowledge and contextual information.

Examples show that interpretation requires the use of information that is not in the input and is neither purely linguistic nor purely stereotypical knowledge. I show that this applies to arriving at the proposition expressed (e.g. the proposition asserted) as well as implicatures.

'Grammatical' and 'semantic' illusions show less than one might think about encapsulation of utterance interpretation. Some do not depend on reaching any stable interpretation. The inversion illusion is due in part to background knowledge. While the Moses illusion is revisable, such revision may be performed by a distinct capacity, so is not clear evidence for cognitive penetrability of the normal utterance interpretation process.

ARTICLE HISTORY Received 2 August 2021; Accepted 26 September 2023

KEYWORDS Inference; pragmatics; modularity; utterance interpretation; verbal communication; illusions

1. Introduction

People standardly communicate by uttering phrases or sentences with certain intonation patterns, accompanied by facial expressions, eye contact and often a variety of gestures. If all goes well the addressee understands what is meant by the utterance, usually very rapidly and

CONTACT Nicholas Allott n.e.allott@ilos.uio.no Department of Literature, Area Studies and European Languages, University of Oslo, Postboks 1003 Blindern, Oslo 0315, Norway

without conscious effort. Linguistic pragmatics investigates how such utterance interpretation works[1] and why it sometimes goes wrong. Given that what the speaker intends to convey by her utterance typically, perhaps always, outstrips what is encoded by the linguistic material uttered, a central task of pragmatics is to explain how that gap is bridged.

Most people who work in pragmatics are committed (with varying degrees of explicitness) to an inferential model;[2] more specifically, to the claim that utterance interpretation is a form of non-demonstrative inference, the hearer attempting to find the best explanation of the utterance. It is generally taken for granted that in this inference any information can be brought to bear, although there is no consensus on how this is achieved. It is controversial whether pragmatic inference is carried out by general central cognition (Sperber and Wilson 1986, 66–67; see also Chiappe and Kukla 1996; Sperber and Wilson 1996), perhaps aided by a capacity for shared attention (Harris 1996); by a 'theory of mind' module applied to utterances (Bloom 2002); or by a dedicated utterance interpretation module (Carston 1997b; Mazzarella and Noveck 2021; Sperber 2000; Sperber and Wilson 2002; Vaccargiu 2023). But participants in this debate either assert or take for granted that '[p]ragmatic processing can call on any background information imaginable, from any source (perception, memory, earlier discourse), so it is informationally UNencapsulated' (Carston 1997b, 5, emphasis in original).[3]

There are, however, some *prima facie* reasons to doubt that the inferential processes the hearer brings to bear are quite as unconstrained, or 'unencapsulated,' as standard intelligent inferences tend to be. One is the speed of utterance interpretation. It has been argued that unencapsulated processes are slow, since they need both to search memory and consider the relevance to the current task of all available information (Fodor 1983, 70).[4] But utterance interpretation is typically very fast, occurring online while the utterance is still proceeding and terminating soon after, with disambiguation, reference assignment to indexicals, and

[1] I focus here mainly on spoken verbal communication. Pragmatics also investigates both non-verbal communication and the interpretation of written texts.

[2] See section 3 below for brief discussion of alternative, anti-inferential views.

[3] It follows that the term 'module' as used by these authors does not mean an informationally encapsulated processing unit (Fodor's notion: see below). Rather, it means (approximately) a dedicated domain-specific processing unit (Carston 1997b, 20; Sperber and Wilson 2002, 9; Mercier and Sperber 2017, 74–75; Mazzarella and Noveck 2021, e199–e200, e202).

[4] This is related to Fodor's claim that unencapsulated processes cannot be studied. For discussion and arguments against this view as applied to pragmatics, see Sperber and Wilson 1996; Allott 2019; Allott and Wilson 2021.

even comprehension of figurative speech typically occurring within about a second.[5]

Another reason to doubt that utterance interpretation is fully unencapsulated is that certain illusions and related phenomena appear to show that the processing of some verbal utterances is unrevisable in light of further information. That provides *prima facie* evidence that the speaker's background beliefs cannot affect the processing of at least those utterances, and casts doubt on the claim that utterance interpretation in general is unencapsulated (Carston 1997b, 21).

I want to argue that these reasons are not conclusive, and that interpretation of utterances does indeed rely on indefinite tracts of background knowledge and contextual information as well as on the parsed linguistic structure (where parsing may be encapsulated to some extent). This paper offers the first fully articulated discussion of what the standard types of example of pragmatic context-sensitivity establish about (lack of) encapsulation.

Intuitively, while some episodes of utterance interpretation involve more or less laborious conscious reflection about what the speaker meant, most seem effortless and spontaneous: 'inferential comprehension is in general an intuitive, unreflective process which takes place below the level of consciousness.' (Sperber and Wilson 2002, 9). It is rather plausible that more reflective, consciously controlled interpretation is unencapsulated, with no particular bounds on the addressee's search for reasons why the speaker spoke as she did. This paper defends the stronger claim that even spontaneous, intuitive (i.e. unreflective) utterance interpretation is unencapsulated.

I will suggest (in section 2) that there are (at least) three kinds of process – or processing mechanism – which can produce representations of external stimuli: informationally encapsulated ones which are spontaneous (fast and unreflective);[6] unencapsulated, reflective ones; and spontaneous unencapsulated ones. This paper argues that utterance interpretation is normally of the third kind: unencapsulated, unreflective and fast. Unreflective inference has been relatively neglected in

[5]See e.g. Swinney 1979 on lexical disambiguation; Tanenhaus et al. 1995 on structural disambiguation; Arnold 2010, 187 on reference resolution; Noveck 2018, 167–169 on metaphor comprehension. Swinney says his results suggest a 'process which acts to select a single meaning from those originally [primed … which is] apparently completed by the time that three syllables of additional information have been processed (approximately 750–1000 ms)' (p. 657). Subsequent work has found similar timescales.

[6]For this use of 'spontaneous' see Sperber and Wilson 1986, 67–69.

philosophical discussions.[7] But in pragmatics we cannot ignore it. While utterance interpretation is mostly fast and automatic,[8] its output is – and needs to be – sensitive to practically any information. A second contribution made by this paper, then, is to draw attention to unreflective but unencapsulated processing by showing that it is at the centre of one very familiar everyday activity.

In section 3 I will explain why people working in pragmatics take utterance interpretation to be unencapsulated: mainly because of examples showing that interpretation is sensitive to information that is neither in the input nor proprietary to language. A further important point demonstrated by such examples is that in many cases some of the information used is too specific to be part of a 'frame' or 'template' that might be retrieved automatically when perceiving a certain object or hearing a particular word.

I also address here some potential alternative views. One of these concedes that working out implicatures and perhaps the speaker's attitude to the proposition expressed are unencapsulated, but claims that an encapsulated process generates the representation of asserted content or the proposition expressed.[9] I will argue for the stronger view that even arriving at this 'direct' (non-implicated) utterance content is unencapsulated, again on the basis of examples showing sensitivity to information external to language and the input.

The unrevisability of certain illusions is one major source of evidence for encapsulation. Measuring the lines in the Müller-Lyer diagram does not dispel the illusion, for example. In section 4, I consider some well-known illusions that arise in utterance interpretation, including the Moses illusion and the comparative illusion, thus providing the first systematic discussion of what such 'grammatical' and 'semantic' illusions (as they have been called) show about pragmatic processing.[10] I argue that they do not provide compelling evidence about whether the

[7]Some recent work on implicit bias (e.g. Mandelbaum 2016) is an exception to this generalisation.

[8]Utterance interpretation is automatic in that it is stimulus-driven and involuntary (Fodor 1983, 52–53, 55), although what utterances are processed depends on where attention is focussed, which is partly under voluntary control. See Palmeri 2003 on characteristics of automatic processes and Zbrodoff and Logan 1986 on the related property of autonomy.

[9]In relevance theoretic pragmatics, the proposition expressed is (as represented in the hearer's mind) 'that propositional form which is developed by pragmatic inferences building on the incomplete logical form decoded from the linguistic form employed in the utterance (hence it is an amalgam of decoded conceptual content and of pragmatically inferred concepts)' (Carston 2002; 379).

[10]Carston 1997b, 21–22 briefly discusses the inversion illusion as a pragmatic illusion. Allott and Rubio Fernández (2002) argue that the Moses illusion is a pragmatic illusion. McCourt 2021, 102–113 argues that the Moses illusion is compatible with encapsulation of semantic processing and briefly discusses its relation to pragmatics.

normal utterance interpretation process is encapsulated. The parser is the processing capacity that assigns linguistic structures to strings of lexical items. The unrevisability of some 'grammatical' illusions suggests that the parser is encapsulated to some extent. However parsing is only one component of – or an input to – utterance interpretation proper. Conversely, I claim that the Moses illusion is a genuinely pragmatic illusion. It is revisable but, I argue, this is not clear evidence for the cognitive penetration of normal utterance interpretation processes, since revisability here might be a matter of changed input due to a shift in attention, or rational reconstruction of the inference, employing other mental resources.

I conclude with two connected suggestions. First, our repeated susceptibility to the Moses illusion is evidence that the normal utterance interpretation procedure is fairly rigid. Secondly, this helps to explain how utterance interpretation can be both unencapsulated and fast.

2. Cognitive impenetrability and reflective and unreflective inference

At least three kinds of process that have been called 'inference' can be distinguished by whether they are cognitively impenetrable and whether or not they are reflective. I first explain these terms.

2.1 Cognitive impenetrability

Encapsulated capacities are those that can use only proprietary information in processing their input (in addition to the information in the input itself) with no influence from outside the system: they are 'modular' in the terminology of Jerry Fodor (1983).[11] More specifically, cognitive impenetrability is encapsulation against 'top-down' influence, i.e. influence from 'central' or 'person-level' attitudes, memories etc. (Fodor 1983, 73–74).

A peripheral system may be cognitively impenetrable without being fully encapsulated, since it may receive input 'sideways' from another peripheral system.[12] The McGurk effect (McGurk and MacDonald 1976) appears to be an example of this kind of cross-talk between perceptual systems. Participants are simultaneously shown the lip movements of a

[11]This is a different use of the term 'module' from that used in the literature mentioned in the introduction. See footnote 4.
[12]For discussion of 'cross-modal' effects see Firestone and Scholl 2016, 4, 65.

speaker saying [ga] and played an audio recording of [ba], resulting in a perception as of [da].[13]

2.2 Intuitive and reflective processes

As noted above, most utterance interpretation strikes the hearer as effortless and spontaneous. What the speaker meant by her utterance is simply part of the experience of hearing the utterance, apparently unmediated by any need to reason about it (Fodor 1983, 55). However, it is also a common experience that in order to understand some utterances one needs to perform reasoning that demands some felt effort (whether a little or a lot) and seems to be under some degree of conscious control. We can call the first kind of experience 'intuitive' and the second 'reflective' (borrowing the terms from Carston 2010; Mercier and Sperber 2009). The point of making this distinction here is in order to state the claim that normal, intuitive episodes of utterance interpretation are unencapsulated: this is more surprising (hence interesting) than the claim that reflective utterance interpretation is unencapsulated.

A distinction along these lines is often appealed to in psychology, but its correct theoretical characterisation is controversial. In recent decades, dual-process accounts have become popular (Evans 2003; Kahneman 2011; Sloman 1996). As Evans writes, '[d]ual-process theorists generally agree that System 1 processes are rapid […] and automatic in nature: only their final product is posted in consciousness' (Evans 2003, 454). In contrast, System 2 is 'slow, consciously controlled and effortful' (Mercier and Sperber 2009, 149).

Dual-process theories have been criticised for assuming without adequate evidence that a number of distinctions line up: System 1 processes are quick, automatic, and carried out by associationist mechanisms, while System 2 processes are slow, consciously controlled and rule-based (Mercier and Sperber 2009, 149; Melnikoff and Bargh 2018).[14] This paper is agnostic about that assumption, and also about Mercier and Sperber's alternative proposal: that reflective inference processes are those which metarepresent premises and thus make the inferential link between premise and conclusion available to consciousness (Mercier and Sperber 2009).

[13]Note that the McGurk effect might not be an illusion, given that phoneme recognition is typically cross-modal, so what phoneme is present may depend on a combination of auditory and visual properties (Rey 2020, 320n24).

[14]Note, though, that as Evans says, 'some dual-process theorists prefer to emphasize the functional differences between the two systems and leave open the relation to consciousness' (2003, 454).

Within pragmatics, Carston (2010) notes that some metaphors receive slow, reflective, consciously-controlled interpretations while others are understood more spontaneously. Intuitively, this phenomenological distinction applies to processing of utterances more generally, and I adopt it here.[15] Explaining what underlies the difference between intuitive and reflective pragmatic processing goes beyond the main aims of this paper, but I make a tentative suggestion in the conclusion.

2.3 Types of processing

Fodor (1983) famously claims that peripheral input processing is modular, where cognitive impenetrability is the central feature possessed by Fodorian modules, and where peripheral input processing includes perception and linguistic parsing.[16] There is a well-known sense (which goes back at least to von Helmholtz in the nineteenth century[17]) in which perceptual processing can be called inferential: it is dedicated to producing representations of the distal objects or scenes that cause the stimuli that impinge on the senses, and the impingements on the senses (retinal image or similar) underdetermine the representations arrived at (Cf. Fodor 1983, 42, 45, 68–69). So the processing has to make certain assumptions, or to work as though it did. In effect, information that is not in the sensory input has to be brought to bear. Fodor's innovation was to point out that this does not entail that just any information can be used: that a process carries out non-demonstrative inference does not entail that it is unencapsulated.

Notwithstanding Fodor's claim, it is contentious whether peripheral processing is cognitively penetrated, and if it is, to what extent.[18] Many studies claim to have found top-down effects on visual processing. To mention just two (claimed) findings: (i) ethnic categorisation affects perception of luminance, in that we see Black faces as darker than luminance-matched White faces 'even for racially ambiguous faces that were disambiguated by labels' (Levin and Banaji 2006, 501); (ii) wearing a heavy rucksack makes slopes appear steeper (Bhalla and Proffitt 1999). However, Firestone and Scholl (2016) make a strong case that such

[15]She also suggests an account of the distinction with similarities to Mercier and Sperber 2009.
[16]I do not mean this wording to imply that parsing is non-perceptual; rather to leave the question open.
[17]von Helmholtz 1924: the relevant part is volume III, first published 1867. One can argue that encapsulated peripheral capacities are not truly inferential, and some authors hedge, e.g. 'at least since Helmholtz, it has been common to think of perception as engaged in *something like inference*' (Gross 2017, 3n1, italics added).
[18]Fodor saw modularity as a matter of degree (1983, 37).

studies do not provide compelling evidence for cognitive penetration because of conceptual and methodological flaws.[19] One area of active debate is the role of attention: are certain cases with alleged top-down effects more like looking at one thing rather than another – which does affect visual processing, but in a way that is compatible with cognitive impenetrability, given that it simply changes the input – or can the *way* that processing happens be affected, so that a given input produces a different output? (Firestone and Scholl 2016, 3, 13–15, 61–62; Nes, Sundberg, and Watzl 2023, 179–180)[20] I return to this point in section 4.

There has been a separate debate on whether syntactic parsing is encapsulated (Pickering and van Gompel 2006 is a useful survey). The view that sentence processing is modular in Fodor's sense has become unpopular in psycholinguistics, but recently, Ferreira and Nye (2017) have argued that data that have been taken as evidence against modularity are not evidence against modularity *per se*, only against a limited range of modular architectures – roughly, those similar in relevant respects to the 'Two-Stage Model' current when the debate on modularity started in the 1980s. There is no space here to go into all the technical details (see Ferreira and Nye 2017, 67ff.), but the Two-Stage model assumed that the only information available to the parser was the input string with categories attached (noun, verb etc.) plus grammatical phrase structure rules. Any effects on parsing of lexical information or prosody would thus be taken as evidence against encapsulation. But that is unreasonable, given that prosody is in the input and a great deal of lexical information is presumably in the proprietary database of the linguistic system. The model also assumed that for each input the output was exactly one structure. But there is no *a priori* reason why the output of the parser could not in some cases be two or more structures among which pragmatic processing must choose.[21] Given this possibility, effects of context in resolving structural ambiguity are not necessarily evidence that parsing is cognitively penetrated.

To approximately summarise the current state of the art: it seems that there is a good deal of peripheral processing which is pretty much

[19] On this debate, see also Pylyshyn 1999; Ogilvie and Carruthers 2016; Gross 2017.

[20] Other possible explanations for apparent counterexamples to cognitive impenetrability include top-down effects within a module and top-down effects on post-modular processing (Fodor 1983; Pylyshyn 1999; McCourt 2021), and methodological flaws including failing to control for task demands and to distinguish perception from judgement (Firestone and Scholl 2016).

[21] Another possibility is that the output of the parser might leave some structural relationships underdetermined; pragmatics would then flesh them out. E.g. Kempson and Cormack (1981) proposed underdetermined quantifier scope.

impenetrable. Illusions and other cases of unrevisability provide strong evidence.[22] As noted above, measuring the lines in the Müller-Lyer diagram does not stop them from seeming different in length. Similarly, quite a lot is now known about how the McGurk effect works, but that does not render linguists immune to it.[23] And, as briefly discussed above for syntactic parsing, much of the putative empirical evidence for cognitive penetrability of peripheral systems can be challenged by a richer characterisation of the systems' input and/or output.

A further property of many perceptual capacities and other peripheral processes such as phoneme recognition and syntactic parsing is that they are domain-specific. That is, each module is only able to take a certain kind of input. This could be whatever is coming from one particular sense, or something more specific within a sense modality, as for an edge-detection module in visual processing. Alternatively, it could be cross-modal, as with linguistic stimuli, which form a domain even though they can be auditory or visual. Fodor also argued that encapsulated modules tend to have several other properties, among them speed and automaticity (which I return to below).[24]

A second type of process or capacity is those that are reflective and not encapsulated. Fodor (1983) argued that central systems (i.e. those performing belief fixation and decision-making) are not modular – neither domain-specific nor encapsulated – on the basis of our ability to bring together apparently disparate information in reaching a conclusion. For example, to decide whether a defendant in a trial is guilty one might need to consider traffic patterns, financial facts and the psychology of the individual (for more on this point see the discussion of 'isotropy' below).

Another view is that little or no intelligence is 'general': central processing is largely or entirely performed by a suite of domain-specific capacities (Carruthers 2006; Sperber 1994; 2001). Certainly considerable progress has been made in recent decades in work on domain-specific central capacities, e.g. theory of mind, the number sense and intuitive

[22]Note, though, that Ogilvie and Carruthers 2016 argue that there are better explanations for unrevisability.

[23]Colin, Radeau, and Deltenre (2005) found limited effects of 'cognitive variables' on the McGurk effect, but these appear to be largely a matter of shifting attention rather than changing the way that a given input is processed. They therefore appear to be susceptible to (re)analysis along the lines Firestone and Scholl (2016, 14) and Gross (2017) suggest for some visual cases e.g. switching perception of the Necker cube by attending to one corner.

[24]For the relation of other properties of Fodorian modules to encapsulation see Currie and Sterelny (2000).

physics,[25] but this does not rule out the existence of systems which bridge domains.[26]

Turning to the question of whether central thought is unencapsulated: it has become notorious that it has some cognitively impenetrable aspects, as brought out vividly in a comment made by Stephen Jay Gould on the conjunction fallacy (Tversky and Kahneman 1982). Most people rank the conjunction 'Linda is a bank teller and is active in the feminist movement' as more likely than 'Linda is a bank teller', although that is logically impossible. Gould writes:

> I know that the [conjunction] is least probable, yet a little homunculus in my head continues to jump up and down, shouting at me – 'but she can't just be a bank teller; read the description'. (Gould 1991 , 469)

The standard explanation is in terms of an automatic, unreflective 'System 1' heuristic, the unrevisability of whose output suggests that it is encapsulated. (See section 4 below for more on unrevisability.) There are many similar cases in the 'heuristics and biases' literature: see Pohl 2004; Kahneman 2011. Such heuristics challenge the view that the peripheral/central and encapsulated/unencapsulated distinctions coincide.[27] On the other hand, they tend to support Fodor's claim that speed and automaticity go with encapsulation.

Fodor claimed that encapsulation *explains* how an inference process can be fast. Processing evidence takes time, as does deciding which evidence to process. Encapsulated processes avoid this (Fodor 1983, 70). But normal, unreflective utterance interpretation challenges his view (Sperber and Wilson 1986, 66ff.). It is inferential in at least the sense explained above – its inputs underdetermine the representations of distal stimuli that it constructs – and is fast and automatic, at least in its normal operation, but it is generally assumed to be cognitively penetrable. If that is right, then it belongs to a third category of inference, spontaneous but not encapsulated, demonstrating that encapsulation is not necessary for speed. In the next section I will argue that normal, intuitive utterance interpretation does indeed fall into this category.

[25] See e.g. (Dehaene 1998; Mikhail 2011; Baillargeon 2004).

[26] Induction over recent research tends to support the view that global systems are harder to investigate, as Fodor claimed: 'The more global […] a cognitive process is, the less anybody understands it' (Fodor 1983, 107). However, it seems to be domain-specificity, not encapsulation, that makes investigation tractable.

[27] Similarly, Melnikoff and Bargh 2018, 287–288 argue that they show that not all intentional processes are controllable.

3. Utterance interpretation: A (mostly) fast unencapsulated process

The thesis under investigation in this section is that normal, unreflective utterance interpretation is unencapsulated in the sense that it can (and typically does) draw on information beyond that contained in the input and a limited proprietary database. One might wonder whether that view is controversial and whether it is worth expending effort looking at evidence for it.[28]

I want to suggest five reasons why it is worth going into the evidence and arguments for the view. First, there are good *prima facie* reasons to doubt it. As noted above, both the speed of utterance interpretation and the fact that it is subject to some unrevisable illusions might be taken as evidence that utterance interpretation is encapsulated.

Secondly, pragmatics, as a branch of cognitive psychology, is in the business of trying to understand the mental processes involved in utterance interpretation. Details of particular interest are the formats of the input, the output and any internal representations, and how the representations are transformed in processing. A central task in constructing a theory of how the input is processed is therefore to find out what information can be used beyond what is present in the input, and whether only a certain limited subset of information can be used (which is compatible with encapsulation), or whether there is no limit in principle on what information can be employed.

Also of great interest is whether there are dedicated faculties for subtasks. If, as I will suggest, some sub-parts of verbal utterance interpretation (understood in a broad sense) are impenetrable, but the overall process is not, then that is evidence for separate processing units or, at the least, some functionally relevant internal structure.

The third reason is that knowing whether utterance interpretation is cognitively penetrable is relevant to ongoing debate over whether utterance meaning is perceived. What such a view amounts to is contested. For some of its proponents the claim is that utterance interpretation is as 'direct' as perception, in particular in being non-inferential (Brogaard 2020; Kissine and Klein 2013; Millikan 1984; Recanati 2002). But that a process is perceptual may not entail that it is not inferential (*contra*

[28]E.g. Stanley (2005, 132 n.1) writes, 'the interpretation of context-sensitive expressions is obviously and non-controversially guided via consciously accessible inferences from background beliefs.' *Contra* Stanley, it is not obvious that the inferences are consciously accessible – see section 4 below and Gross this volume – and while an inferential view is standard it has also been questioned, as discussed below.

Kissine and Klein 2013). As noted above, Fodor (1983) famously argued that visual processing is inferential even though it is encapsulated. Still, the properties of the utterance interpretation process should be of interest to those who argue that it is perceptual, as Drożdżowicz (2023, 212) has recently argued, and cognitive penetrability in particular has played a major role in debates about whether various other processes are perceptual (e.g. Firestone and Scholl 2016; Block 2016; Nes, Sundberg, and Watzl 2023, 176–178). While there is a range of views on how these properties relate to each other, so establishing that utterance interpretation is cognitively penetrated would not settle the question, it would be in some tension with the view that utterance meaning is perceived. Drożdżowicz (2023, 216–218) suggests that the cognitive penetration of utterance interpretation implies a process of a complexity that is hard to reconcile with perceptualism. Conversely, a supporter of perception of 'high-level' properties[29] might argue that there are many non-encapsulated perceptual processes. In an recent example of how this debate plays out in an area closely related to pragmatics, Helton (2018) argues that attribution of certain (non-communicative) intentions to others is perceptual given that it is unrevisable, where this is closely related to cognitive impenetrability. (I return to revisability in section 4 below.)

Relatedly, some people think one has a particular kind of warrant or justification in accepting the output of a perceptual process (or *believing* it or *using it as a premise* in reasoning, or similar). One argument in favour is this: if the workings of a capacity are not influenced by one's beliefs and desires, one can trust its output – or, rather, one does not have one kind of reason one might have had for *not* trusting it (Fodor 1984). If utterance interpretation is cognitively penetrated then this would not apply to it.[30]

A fourth point is that, as I will show, many of the examples that indicate that utterance interpretation in general is unencapsulated also establish this more specifically for interpretation of the proposition expressed. This clashes with a traditionalist view according to which pragmatic inference governs implicatures while propositional content ('what is said') is the concern of linguistic semantics.

Finally, and most important here, the lack of encapsulation and the examples that demonstrate it are central reasons that most pragmatic theorists take utterance interpretation to be inferential in a sense of the term that goes beyond the thin Helmholtzian sense discussed above.

[29]Helton 2016 is a useful review of this literature.

[30]This obviously does not impugn other reasons that there might be for trust: e.g. that utterance interpretation is carried out by a fairly reliable mechanism or mechanisms.

The key difference is precisely the lack of encapsulation, which implies that the possible output – the conclusion or conclusions reached – is a good deal less constrained by the input than it would be for an encapsulated process. It is also possible that the input to pragmatic processing is conceptual and propositional: namely the fact that the speaker has uttered a sentence with a certain structure, in such-and-such circumstances, accompanied by certain gestures etc.[31] If so this would contrast with the presumably non-conceptual input to low-level linguistic processing – a stream of sound – and the retinal image that is the input to visual processing.[32]

Two caveats are worth noting. First, some pragmatic theorists have proposed non-inferential accounts of utterance interpretation in recent decades. Recanati (2002; 2004) claims that recovery of the proposition expressed is carried out by non-inferential, associationist mechanisms, and Mazzone (2011; 2018) extends this claim to utterance interpretation as a whole, including arriving at implicatures. As noted above, Kissine and Klein (2013) have also argued for a non-inferential perceptual account. These theorists are aware of data of the type I discuss below which show that utterance interpretation is unencapsulated, but they resist the conclusion that utterance interpretation is inferential – at least in any sense of the term thicker than something like the Helmholtzian position outlined above (see Recanati 2002, 116). This paper does not aim to settle this debate – rather to set a precondition for it: theories of utterance interpretation need to explain how the output of utterance interpretation can be sensitive to practically any information.[33] This shifts the burden against associationist views, if we accept that sensitivity to evidence is a hallmark of inferential processes.[34]

A second caveat is that still thicker senses of 'inference' are deployed in much discussion in philosophy: e.g. Boghossian's 'reasoning that is person-level, conscious and voluntary' (2014, 2–3), properties that do not seem to apply to normal spontaneous utterance interpretation, as discussed in section 2.2 above. If one is interested only in such thicker senses of inference,[35] then normal spontaneous utterance interpretation seems to be ruled out ahead of empirical investigation. The only point I want

[31]See Sperber and Wilson (1987, 737).
[32]I do not explore this putative contrast further here due to lack of space.
[33]See Carston (2007), Allott (2008) and Mazzarella (2014; 2015) for argumentation in support of inferentialism.
[34]On evidence-sensitivity as a property of inference, see Mandelbaum (2013).
[35]This applies to Recanati (2002; 2004) who is interested in inference in a narrow sense, such that it is consciously available (although it need not be occurrently conscious).

to make here about such views is that evidence that an unreflective, automatic process may be unencapsulated supports the claim that there is a non-minimal kind of inference that is nonetheless not something conscious and voluntary.

Given all of the above considerations, it is of some interest to set out accurately the arguments and evidence for cognitive penetration of utterance interpretation and to see what counter-arguments might be raised against them. The view has been around for quite some time. Katz and Fodor (1963, 178–179) argued that 'nonlinguistic information of any kind may be involved in the understanding of [an utterance of] a sentence'[36] on the basis of examples such as:

(1a) Should we take junior back to the zoo?

b) Should we take the lion back to the zoo?

c) Should we take the bus back to the zoo?

As they say, to arrive at the correct readings we would typically need to make use of the information that lions but not children or buses are kept in cages. But more specific information would lead to different readings: e.g. if one knows that the bus in question belongs to the zoo but has been stolen, then (1c) could get a reading more similar to (1b). They say that:

> We have convinced ourselves of the truth of th[e] claim [that nonlinguistic information of any kind may be involved] by making it the basis of a party game. One person supplies a fact, however obscure, and the others try to construct a sentence which that fact disambiguates. The game is not remarkably amusing, but it is surprisingly convincing. (1963, 179n11)

Earlier comments made by Grice highlight other aspects of context that matter to interpretation:

> [I]n cases where there is doubt, say, about which of two or more things an utterer intends to convey, we tend to refer to the context (linguistic or otherwise) of the utterance and ask which of the alternatives would be relevant to other things he is saying or doing, or which intention in a particular situation would fit in with some purpose he obviously has (e.g. a man who calls for a 'pump' at a fire would not want a bicycle pump). (1957, 387)

[36]This foreshadows Fodor's later (1983) view that unencapsulated processes cannot be studied, given that Katz and Fodor argue that a complete theory of how sentences are understood is impossible since 'it would be required that the theory represent ALL the knowledge speakers have about the world' (p. 178; their caps). For counter-arguments, see Allott and Wilson (2021), 439.

Note that lack of encapsulation is not a sufficient condition for a capacity's being inferential. It is easy to imagine a non-inferential algorithm that simply combines its input with some information taken at random from the internet. It seems to be the combination of the following that matters here: the capacity aims at a representation that explains the stimulus; the input radically underdetermines the output representation; and the processing is unencapsulated.[37]

As we have already seen, the argument invokes examples that show there is no principled limit to the information that can be brought to bear on the process. Fodor (1983, 105–107; 109–111) called this property 'isotropy'. An inference process is isotropic if and only if information from any domain may be relevant to the conclusion. Fodor argues that confirmation of scientific hypotheses is isotropic, a point nicely demonstrated by the relevance of the rate of cooling of the sun to Darwin's theory of evolution (Carruthers 2003) and of rabbit designs on indigenous American Mimbres pots to the physics of supernovae (Antony 2003). Fodor further argues that scientific confirmation is a paradigm case of central cognition, and that belief fixation in general is isotropic (Fodor 1983, 104; but see Sperber and Wilson 1986, 117; Sperber and Wilson 1996).

Returning to utterance interpretation, consider the following:

(2a) A policeman arrested John yesterday; he had just stolen a wallet. (Recanati 1993, 265[38])

b) A policeman arrested John yesterday; he needed one more arrest to qualify for the end-of-year bonus.

c) A policeman arrested John yesterday; he had just taken a bribe.

In each case, the hearer (or rather, in normal cases, his pragmatics faculty) has to work out who the speaker referred to with her use of 'he'. This could be the policeman mentioned or John, or a third party. The interpretation reached can depend on an indefinite range of apparently non-linguistic facts which are not present in the input. For example, we know that policemen are supposed to arrest criminals, including thieves, and this information might result in representations in which 'he' in (2a) refers to John and in (2b) to the policeman.

[37]This is a (tentative) sufficiency claim for non-demonstrative inference in this somewhat-thicker-than-Helmholtzian sense, not a claim that these conditions are necessary for inference (in any sense). For one thing, inferential processes which work out the consequences of beliefs surely do not all aim at representations of causes of external stimuli.

[38]Cf. Winograd (1972, 33).

One might try to argue that such information is part of a 'frame' or 'schema' – a collection of stereotypical information on some topic – and one might also claim that that 'frame' is part of the knowledge one has by virtue of understanding the word 'policeman'.[39] Thus one might argue that (2a) and (2b) do not show that extra-linguistic information is involved in reaching an interpretation. Whether or not that is right, the example in (2c) suggests that reference assignment depends on factors that are too particular and episodic to be part of word meaning. The assignment of reference there could depend on the hearer's views on the honesty of John and the local police force, and surely one can know the meaning of 'policeman' without knowing whether the police in a particular city are bribeable.[40] Similarly for the specific information about the zoo bus discussed above. Crucially for my argument, such examples show not only that non-linguistic information is used in utterance interpretation, but also that specific, non-stereotypical information is used. They are therefore evidence for full-blown Fodorian isotropy rather than some more limited kind of reliance on a subset of information beyond the input.

In fact, for well-known reasons, what matters may not be what the hearer thinks about the police or the bus. Rather, what matters may be the hearer's view of what the *speaker* thinks about the topic, since hearers try to work out what the speaker intended. Or it may be that what matters in a particular case is the hearer's view of what the speaker thinks the hearer thinks, since the speaker's intentions are constrained by what she can expect the hearer to work out (and hearers seem to be guided by this). Such reasoning can in principle be iterated indefinitely.[41] It is hard to believe that we keep track of much more deeply embedded meta-representations,[42] although there is some evidence (O'Grady et al. 2015) that we are capable of doing so. Whatever the right depth of representations is here, none of this offers any comfort to someone who would like to argue that pragmatic processing is encapsulated. That is because many of the nested attitudes used in

[39]Cf. Recanati (1993, 30–31), and on lexical frames more generally Fillmore (1976) and the references listed at Sperber and Wilson (1986, 286 n). 14 and Recanati (1993, 94 n. 29). I am not attributing to any of these sources the view that all the information in a frame or schema is part of the meaning of a word.

[40]It is sometimes argued that only a certain independently-specifiable subset of information ('the narrow context') is required to fix the reference of indexicals. That applies at most to some uses of indexicals (Recanati 2004, 57).

[41]But see Lederman (2018).

[42]See Sperber and Wilson (1986), 17ff. The experimental literature on perspective-taking in conversation is also relevant here. Brown-Schmidt and Heller (2018) is a recent review.

utterance interpretation are too situation-specific to be part of either linguistic knowledge or stereotypical knowledge such as frames or schemas. For example, an interpretation of (2c) above might depend on the hearer representing (3a). Similarly, assigning the correct referent in an experimental reference task (e.g. Nadig and Sedivy 2002 ; Grigoroglou and Papafragou 2022 , 431–432) or a similar real-world situation might require the information in (3b):

(3a) The speaker doesn't know that I know that the police here are corrupt.

b) The speaker knows that I can see the green cup on the left.

Taking a different tack, it might be that some subpart of utterance interpretation is encapsulated. I'll briefly discuss how this might apply to two such parts: parsing, and arriving at the proposition expressed.

As noted above, since Fodor's work on modularity there has been a long-running debate about whether syntactic parsing is cognitively penetrated. A recent paper argues that 'the modularity hypothesis was prematurely rejected' (Ferreira and Nye 2017, 63). Even if this is right, though, not much of interest follows about whether verbal utterance interpretation is encapsulated. The possibility of a sentence being used to express different propositions is not always traceable to linguistic ambiguity in the sentence (i.e. lexical or structural ambiguity, which both by definition involve the correspondence of one surface form to two or more underlying forms). This obviously applies to assignment of referents to indexicals, but also to many other cases, including understanding what relation is indicated by a use of a possessive (e.g. 'Mary's book'[43]), whether a given use of 'and' conveys merely logical conjunction, *and then*, *and as a result* or one of several other possibilities (Carston 2002, ch. 3); embedded 'implicatures',[44] and perhaps referential and descriptive uses of definite noun phrases (Recanati 1989a; Robbins 2007, 307–308).

A stronger way of putting these points is to say that utterance interpretation is an inferential, unencapsulated process that is often fed by the output of the parser[45] (which may be encapsulated to some extent), which it treats as merely a clue – albeit an intricately structured

[43]The relation the possessive stands for may not always be a constituent of the proposition expressed by the speaker: e.g. if the phrase 'Mary's book' is intended to contribute a particular individual referent (whether individual type or token) rather than a description. In some cases, one may nonetheless need to grasp the relation in order to work out which book is the referent.

[44]See Noveck (2018, ch. 7).

[45]Given the existence of non-verbal utterances such as pointing, mimes etc., utterance interpretation does not always receive input from the parser.

one – to the speaker's communicative intentions. This is a view articulated over decades by relevance theorists, following Sperber and Wilson (see especially their 1986, 9–14; 174–178). As Sperber and Wilson point out, in addition to the sources of underdetermination already discussed here, that a particular sentence has been uttered also falls short of determining what attitude the speaker is expressing to its content: e.g. whether endorsing it or questioning it, and whether putting it forward as her own view, or as representing someone else's (op cit, 10–11; 228ff.). And of course sentences can be uttered non-communicatively (op cit, 177–178). So as Kent Bach says, 'even if what a speaker means consists precisely in the semantic content of the sentence he utters, this still has to be inferred.'[46] (2006, 24)

There is a second way of trying to carve out an encapsulated part from utterance interpretation. One might concede that working out the speaker's attitude to the proposition expressed and working out conversational implicatures of utterances are unencapsulated, but argue that an encapsulated process arrives at the asserted content or proposition expressed.[47] But this move seems already to have been blocked by many of the examples above. The differences in possible readings of (1a–c) and (2a–c) are not only or even primarily a matter of the speaker's attitude or what she implicates, but of the proposition that she expresses. The history of pragmatics since the 1980s has been in large part the discovery of ways in which the proposition expressed is underdetermined by what is linguistically encoded. (See e.g. Carston 2002; 2012; Neale 2007).[48]

Here are two of the clearest cases of linguistic underdetermination of the proposition expressed. First, if reference is not assigned to indexicals, then in many cases the interpretation reached will not be propositional. So arriving at the proposition expressed is unencapsulated to at least the extent that such reference assignment is unencapsulated.

[46]It is also possible that speakers never express sentence-type meanings, e.g. if no natural language sentences encode propositions. See the next paragraph in the text and the references there.

[47]Some semantic minimalists, while conceding that utterance interpretation is unencapsulated, have argued for an encapsulated truth-conditional semantics module (e.g. Borg 2004, 78ff.; 90ff.; see also Robbins 2007, 307ff.). It is not clear how, on such views, the semantics module relates to utterance meaning. Borg argues that we should expect the output of the semantics module to be 'divorced […] from intuitive judgements of communicated or conveyed meaning' *because* the latter are unencapsulated: they 'depend on a vast range of information located beyond the reach of the language faculty' (pp. 8–9).

[48]Some have argued that all constituents of the proposition expressed correspond to a constituent in linguistic form. Then truth-conditional interpretation would require only parsing, disambiguation and reference assignment to indexicals. Whether that is right is irrelevant to my argument here, because disambiguation and reference assignment are pragmatic processes that draw on non-linguistic and non-stereotypical information (as Neale 2007 points out).

This argument extends to lexical modulation and the expression of '*ad hoc*' concepts[49] – that is, roughly, uses or exploitations of a word to convey something other than its encoded meaning – since this can result in neologisms like the one in (4):

(4) The girl porched the newspaper. (Cf. Clark and Clark 1979)

As Wilson and Carston (2007, 240) point out, if the novel concept of 'porching' is not a constituent of the proposition expressed in this case, then there is no proposition expressed. That seems flatly wrong, since one can deny what is asserted by an utterance of (4), reason from it etc.: 'If she porched the paper and no one's moved it, then it must still be there'.[50]

There are two further arguments against the claim that arriving at the proposition expressed is encapsulated. One is that the processing that arrives at the proposition expressed sometimes seems to go via an implicature or to be very tightly connected with it. Consider this dialogue:

(5) Receptionist 1: That's Mrs Prescott to see Dr Jones.

Receptionist 2: She's busy.

Suppose receptionist 2 intended to communicate that Dr Jones is too busy to see Mrs Prescott at present. Then part of the meaning is implicated: roughly, *Dr Jones cannot see Mrs Prescott at present*. Assigning reference to 'she' may be reliant on grasping that implicature, since it relies on working out that what is asserted here is put forward as a reason why Dr Jones cannot see Mrs Prescott. (Compare an utterance in the same context of 'She's late.')

On the basis of examples like these, relevance theory claims that explicatures and implicatures are arrived at through a process of 'mutual adjustment' (Sperber and Wilson 1998, 197; Carston 2002, 143–149; 2007). One might see such cases as involving two interlocked inferences, or one inference process that delivers a conclusion about the speaker's meaning overall. Note that relevance theorists and some others would also argue that the proposition expressed here is enriched in several ways relative to the sentence's form or encoded linguistic meaning, to

[49]The terms are respectively from Cruse (1986) via Recanati (2004), and relevance theory (Carston 1997a; Sperber and Wilson 1998). They are not equivalent: some cases of lexical modulation may involve the tokening of a non-lexicalised but stored concept, while a completely novel coinage might express an *ad hoc* concept, but would not be a case of lexical modulation.

[50]These are in effect applications of the scope principle, a diagnostic test for explicated content: if a pragmatically derived aspect of meaning falls within the scope of an operator such as negation or a conditional, that is evidence that it is part of the proposition expressed, and is therefore not an implicature (Recanati 1989b, 322–325; Carston 2002, 191ff.).

something like: *Dr Jones is busy to degree X at present*, where degree X is a degree of busyness that precludes attending to a patient.

It is plausible that in such cases the hearer has to grasp the implicature to work out that the proposition expressed has been asserted, and that is so even if (as one might suppose) some non-inferential process (e.g. based on accessibility) first generates the proposition as a possible candidate to be evaluated. That is because a reason for taking this proposition to be the one expressed is that it supports the implicature.

A final argument comes from looking at experimental work on pragmatics. It is one of the fundamental assumptions of the field that changing the context can lead to different interpretations. This is commonly used in experimental design. One standard kind of task involves a target sentence which receives different interpretations in two different contexts: e.g. metaphorical and literal readings of 'Regardless of the danger, the troops marched on' (Ortony et al. 1978, 467). Such context-sensitivity is routinely exploited in work on reference assignment, disambiguation, and lexical modulation – all of which contribute to the proposition expressed – as well as on cases that one might see[51] as arriving at a representation of implicatures.[52] One only has to look at the contents of contexts used in such experiments to reliably and predictably produce different interpretations[53] to see that there is no obvious limit on the kind of information that can lead to different interpretations.

There is experimental work that explores the question of whether even early stages of pragmatic processing are influenced by context. A good deal of evidence indicates that it is. For example, Tanenhaus et al. (1995) found very early effects of visual context on structural disambiguation. In a study of pragmatic enrichments of scalar terms, Breheny, Katsos, and Williams (2006) write: '[w]e have found no evidence in our on-line studies that there is a stage of processing which is impenetrable to contextual assumptions'(459). A more general study of Quantity enrichments found that processing takes account of 'the speaker's 'conversational purpose' and the speaker's epistemic state' with 'a bias to the correct target begin[ing] to form within 200–300 ms' (Breheny, Ferguson, and Katsos 2013, 434).[54]

[51]The cautious wording here is because it is controversial whether figurative speech such as metaphor contributes to implicature or to what is 'directly' expressed (Carston 1997a; Bezuidenhout 2001; Camp 2006).

[52]For these different areas of experimental pragmatics, see Noveck (2018).

[53]Context may be presented as pictures, a video or text, *inter alia*. Here the two contexts were short texts, one about children, the other soldiers (Ortony et al. 1978, 467).

[54]For a recent overview of the literature on Quantity inferences that reaches similar conclusions, see Noveck (2018), ch.s 7 & 14.

4. Evidence from illusions?

The unrevisability of certain illusions is often taken to be a key source of evidence that some peripheral processes are encapsulated against at least some beliefs (and other attitudes), following Fodor (1983, 66, 69). It does not matter how much evidence you have or how strongly you believe that the lines in the Müller-Lyer diagram are the same length. They still look different. As Fodor says, 'it is hard to see an alternative to the view that at least some of the background information at the subject's disposal is inaccessible to at least some of his perceptual mechanisms.'[55]

There are illusions in utterance interpretation which are unrevisable. Are they evidence that utterance interpretation is encapsulated? In this section I suggest a deflationary view, for two reasons. First, I argue, certain unrevisable illusions in verbal utterance interpretation at most indicate the encapsulation of subtasks such as phoneme recognition or parsing, not utterance interpretation as a whole. Secondly, and more subtly, while a genuinely pragmatic illusion, the Moses illusion, exists and is revisable, it is not clear whether this indicates cognitive penetration of pragmatics or is due to rational reconstruction of the inference by 'higher' cognition.

What is unrevisability? A reasonable approximation is that a mental state is unrevisable if and only if it persists as long as the stimulus continues to impinge, regardless of counter-evidence possessed by the subject (Cf. Helton 2018, 245, 248).[56]

The McGurk effect (set out above) is a clear case of unrevisability of a sub-part of verbal utterance interpretation, namely phoneme recognition. Knowing about the effect and its causes does not make it go away: the speech sound continues to strike one as [da], very much as the Müller-Lyer lines continue to strike one as being different lengths. But what is unrevisable here would be merely an input to the task of working out what the speaker meant by her utterance, for reasons discussed in the previous section.

[55]For a recent statement of the view see Firestone and Scholl (2016, 52); Helton (2016, 255) argues that cognitive impenetrability is stronger than unrevisability. Ogilvie and Carruthers (2016) argue that the persistence of illusions does not show 'that high-level knowledge cannot influence or contribute to lower-level processing' (p. 726). As they say, the question is really whether there is a better explanation, case by case. In this section I consider explanations for different illusions.

[56]See also Gross (2022) on 'evidence-insensitivity'.

4.1 Grammatical illusions

There are some grammatical illusions which are unrevisable,[57] but it is not immediately clear whether what is unrevisable in these cases is an utterance interpretation or something else, such as an intuition of acceptability or meaningfulness. Among the best-known are the missing-VP illusion, illustrated in (6a), and the comparative illusion shown in (7):[58]

(6a) * The man the girl the cat scratched died.

b) ?? The man the girl the cat scratched kicked died.

c) [NP1 [NP2 [NP3 VP1] VP2] VP3]

(7) * More people have been to Berlin than I have.

Discussion of these illusions is complicated by the fact that there is no consensus on their causes. It is generally accepted that the sentences in (6a) and (7) are ungrammatical but native speakers find them acceptable (or in the case of 6a, at least better than 6b – even though 6b is grammatical, with the structure given in 6c). Also, it seems that in both cases the acceptability depends on the hearer/reader not noticing that the structure is ungrammatical.[59]

(Janet) Fodor, Nickels, and Schott (2017) have shown that grammatical multiply–centre-embedded sentences structurally identical to (6b) can be made acceptable by manipulating their prosody, as in (8a).[60] They also found that the 'encouraging' prosody that makes such sentences acceptable facilitates the detection of the problem with sentences with a missing VP, as in (8b) which has the same defective structure as (6a).

(8a) The elegant woman that the man I love met moved to Barcelona.

b) * The elegant woman that the man I love moved to Barcelona.

Making missing-VP sentences easier to process makes them less acceptable. This implies that (6a) is not fully parsed. Why then does it seem acceptable? One possibility is that an utterance interpretation is reached by the

[57] Some errors in production and parsing which are revisable have also been called 'grammatical illusions': e.g. most of those discussed in Phillips, Wagers, and Lau (2011, §4).

[58] First published in Montalbetti (1984) and Frazier (1985), respectively.

[59] For the ungrammaticality of (7) see Wellwood et al. (2018, 544–545). On the missing-VP illusion see §3.2 of Fodor, Nickels, and Schott (2017), which provides a helpful introduction to the different accounts that have been proposed.

[60] As this demonstrates, prosody matters even in silent reading. Because of the way prosodic boundaries connect with syntactic ones, the formula for success is that NP1 and VP3 must be roughly equal in length to each other and to everything in between them, where length in English phonology is a matter of number of stresses, not number of syllables.

pragmatic faculty. Anecdotally, some people report a conjunctive interpretation: the man died and the cat scratched the girl. But this interpretation is revisable on reflection and (anecdotally) not everyone gets it. This suggests that the relative acceptability of sentences like (6a) – which seems both unrevisable and widely shared – is not due to an illusory utterance interpretation. A possible explanation for the intuitions is that some kind of internal 'success' signal is wrongly generated whenever the parser attempts to process such sentences.

The comparative illusion (or 'CI') illustrated by (7) has generated much informal discussion and several conference papers, but only one published study (Wellwood et al. 2018).

(7 repeated) * More people have been to Berlin than I have.

A feature of particular interest here is that participants report finding sentences like (7) meaningful. Anecdotally, for some people this intuition of meaningfulness survives considerable reflection (Wellwood et al. 2018, 546). Is this then an unrevisable illusory utterance interpretation?

Probably not. While 'speakers tend to believe sentences like [(7)] are acceptable and have a coherent interpretation, […] they struggle to articulate that interpretation.' (Wellwood et al. 2018, 546) So it is not clear that most people reach any stable interpretation of the utterance. Further anecdotal evidence suggests that some people are sure they get no interpretation, but still have the intuitions as of acceptability and meaningfulness. (The current author is in this population.)

The explanation for the comparative illusion is rather technical. Wellwood et al. argue that in processing CI sentences word by word, 'speakers' semantic knowledge leads them to consider [early on] an event-counting reading licensed by the normal syntactic rules in fully grammatical comparatives, but [they] fail to notice when that interpretation is no longer available.' (Wellwood et al. 2018, 546) In essence, a hearer/reader is led up a semantic garden path, but – unlike with well-known syntactic garden-path sentences – does not notice that no stable representation of the sentence is reached. Apparently, this processing trajectory is what generates intuitions of acceptability and meaningfulness here.

Sentences like the one in (9) might be thought better candidates to demonstrate unrevisable utterance interpretations:

(9) No head injury is too trivial to ignore.[61]

[61]From Wason and Reich (1979).

It is grammatical, but its compositional meaning is pragmatically strange: *All head injuries should be ignored, even the most trivial ones*. Hearers/readers generally do not understand it that way, but rather as though it said that no head injuries should be ignored, i.e. *all should be treated, even the most trivial ones*: hence the label 'inversion illusion' sometimes used for such cases (O'Connor 2015, 6, 16, 142ff., 186). There seem to be three factors: the sentence has multiple negations (explicit and implicit) which makes it hard to process; the semantics of the phrase *too trivial to ignore* implies or presupposes that we are more likely to ignore head injuries when they are not trivial (i.e. when they are serious);[62] and the sentence urges us to ignore head injuries (Wason and Reich 1979; O'Connor 2015, 142–150). The second and third of these obviously clash with speakers' background knowledge.

It is generally assumed that in cases like (9) pragmatics overrides the syntax-semantics system (which is struggling with the stacked negations), delivering the interpretation noted above, which is a better fit with world knowledge (Sanford and Emmott 2012, 28). On this view:

> the grammar is encoding a message that a reasonable person would probably not utter, and people appear to rather systematically 'override' that message by inverting it into something that falls in line with their own world knowledge. (O'Connor 2015, 147)

Anecdotally and intuitively, the interpretation reached is hard to revise. As Carston writes:

> The rigidity of this example is striking; even after having carefully worked out (by some general problem-solving process) what it actually means, renewed confrontation with it inevitably gives the mistaken (but natural, plausible, relevant) interpretation. (1997b, 21)

Is this evidence for cognitive impenetrability of utterance interpretation?[63] I suggest that in fact it shows the opposite. Here world knowledge is pushing towards the illusory interpretation. Contrary to the semantics of the sentence, we know that, in general, head injuries should not be ignored and we believe that head injuries are more likely to be ignored when they *are* trivial.

To test whether there is cognitive impenetrability here, it is no good to observe that processing is unrevisable in a context in which speakers have these beliefs. It has been argued that sentences presented as though out

[62]Cf. ??'It's too cloudy to rain' (O'Connor 2015, 145).
[63]See also Carston (1997b, 21–22) for brief discussion of this question.

of context are actually processed in a context the hearer brings to them (Breheny 2011, 565), and this is in fact assumed by standard explanations of this illusion.[64] A crucial test of cognitive impenetrability here is a context which changes the crucial beliefs to ones congruent with the compositional meaning of the sentence. For the sentence in (9) it would be hard to get participants to 'internalise' such a strange context. However, O'Connor (2015) reports a series of experiments that varied the lexical items in similar sentences to test the contributions of internal anomaly and plausibility in the light of background knowledge. E.g. in one experiment, test items included 'No social program is too wasteful to oppose,' 'No social program is wasteful enough to oppose,' and 'No social program is too wasteful to support' (192). She found that:

> comprehenders were clearly not very good at interpreting [inversion] illusion sentences even in the best of circumstances — even items that had highly plausible veridical meanings and relatively implausible inverted meanings were still processed correctly only approximately at chance (195)

In other words, the form of these sentences makes them hard to process, and that is enough to cause trouble, even without a clash between the sentence semantics and background knowledge. O'Connor argues that her results support a new explanation: perhaps due to the great processing difficulty they cause, we parse sentences of the form 'No X is too Y to Z' as structurally ambiguous in a way that is not predicted by the syntax of each part. On this account, plausibility in the light of background knowledge plays a smaller role, but still influences the interpretation reached, mainly by determining which of the linguistic structures is chosen (229).

4.2 Semantic illusions

We are also susceptible to so-called semantic illusions (Erickson and Mattson 1981[65]), such as the one induced by (10):

(10) How many animals of each kind did Moses take into the ark?

Most people answer 'Two', even though they know that it was Noah who stocked the ark. The cause is debated (Park and Reder 2004, 278ff.). There is some similarity to the comparative and missing-VP illusions: readers/hearers do not notice that the question asks about Moses, not Noah.[66]

[64]Specifically, it is assumed by what O'Connor calls the 'Channel Capacity Hypothesis' and the 'Change Blindness Hypothesis', as well as by her alternative 'Hypernegation Hypothesis' (O'Connor 2015, 157, 160, 170).

[65]For further references and discussion see Park and Reder (2004); McCourt (2021).

The interpretation of examples like (10) is revisable, *inter alia* by simply drawing attention to the anomalous word. So such cases appear to show cognitive penetrability of utterance interpretation. However, it is not clear whether the revision of the interpretation is performed by the faculty that normally performs utterance interpretation. Given that the output of the pragmatic faculty is conceptual, it is not surprising that it can be rationally revised in the light of further evidence.[67] One can reflect: did the speaker really mean that?

It is an interesting open question whether reflective utterance interpretation is a slower, more directed use of the same mental capacity that performs fast spontaneous utterance interpretation, a fully distinct kind of processing, or something in between. In experimental work it is generally assumed that off-line measures such as questionnaires may tap into qualitatively different processing from on-line measures like eye-tracking and real-time FMRI.[68] Also, as noted in section 2.2, there seem to be two modes of utterance interpretation, one fast and intuitive, the other reflective. It may be the second kind of processing that performs revisions of utterance interpretations.

A third possibility is that what happens in revision of cases like (10) is that the normal process is re-run with shifted attention, and the difference in attention leads to different results. There is some evidence that attention is a factor: the illusion is facilitated when the target item (here 'Moses') is not in a linguistic focus position where it would typically receive more attention (Bredart and Modolo 1988). Although attention is to some extent under conscious control, that a process is affected by attention is not necessarily indicative of cognitive penetration since it may change the input to processing rather than the way that processing is carried out. (See Firestone and Scholl 2016, 3, 13–15, 61–62; Gross 2017; Nes, Sundberg, and Watzl 2023, 179–180, for discussion of this point as it applies to perception.)

Conclusion

This paper has argued that utterance interpretation is unencapsulated, on the familiar basis of examples that show that reaching the right

[66]The illusion depends on certain similarities of 'Moses' to 'Noah' which are 'semantic' in a loose sense (van Oostendorp and De Mul 1990, 37): 'Nixon' (Erickson and Mattson 1981) and even 'Adam' (van Oostendorp and De Mul 1990) work less well than 'Moses'.

[67]Cf Helton (2018, 245).

[68]E.g. Breheny, Katsos, and Williams (2006); Noveck (2018), 91–92.

interpretation requires[69] the use of information that (a) is not in the input and (b) is neither purely linguistic nor purely stereotypical knowledge. That is, utterance interpretation is isotropic in Fodor's sense. I have discussed evidence showing that this applies to arriving at the proposition expressed as well as to implicatures and other parts of speaker meaning. In addition, the input underdetermines the output, and the process aims at a representation that explains the stimulus. I have suggested that this combination is why most pragmatic theorists take utterance interpretation to be inferential in a substantive sense.

I have also looked at some 'grammatical' and 'semantic' illusions and argued that they show less than one might think about whether utterance interpretation is encapsulated. The – unrevisable – comparative and missing-VP illusions do not depend on reaching any stable interpretation of the utterance. The inversion illusion, which is *prima facie* unrevisable, is thought to be due in part to background knowledge either overriding parsing or choosing between structures. It therefore tends to support the view that utterance interpretation is cognitively penetrated. Conversely, the Moses illusion is clearly revisable, but such revision may be performed by a distinct processing capacity or be due to shifted attention altering the input, so is not clear evidence for cognitive penetrability of the normal fast and automatic utterance interpretation process.

The main conclusion of this paper is thus that normal fast, unreflective utterance interpretation involves unencapsulated processing. But how can utterance interpretation – or any non-demonstrative inference process – be fast, automatic *and* unencapsulated? This question was discussed by Sperber and Wilson (1986, 66–68). They argue that in spontaneous non-demonstrative inference, 'however much evidence *might* have been taken into account, however many hypotheses *might* have been considered, in practice the only evidence and hypotheses considered are those that are immediately accessible.' (Sperber and Wilson 1986, 66–67, their italics.)[70]

I would like to finish by suggesting that the Moses illusion provides some evidence for their view. Some participants fall repeatedly for Moses-illusion cases even after being taught about the illusion and warned that similar sentences are coming up (Erickson and Mattson

[69]My claim here is not that every successful episode of utterance interpretation involves use of non-linguistic (etc.) information; rather, that any episode might require such information, and thus that a theory of utterance interpretation has to account for this. I intend my claim to be similar in this respect to the Antony/Carruthers/Fodor claim discussed above about the isotropy of scientific explanation. (My thanks to Solveig Aasen for suggesting I could be clearer on this point.)

[70]Compare Carruthers' (2007) 'wide-scope encapsulation'.

1981). This susceptibility needs to be distinguished from unrevisability. The latter involves sustained exposure to one stimulus, while here participants are fooled by a series of distinct (but similar) stimuli. My suggestion is that this continuing susceptibility to the illusion may be evidence that the normal utterance interpretation procedure is fairly rigid, where this is a distinct property from encapsulation.[71]

The key idea is that in spontaneous processing there is little or no top-down control over the *way* certain inputs are processed: 'Automatic processes are often rigid and stereotypic' (Palmeri 2003, 291).[72] Similar inputs in similar contexts will generally receive similar processing. Suppose input in a certain domain activates a particular heuristic or set of heuristics. These heuristics may in principle have access to any and all information (i.e. they are unencapsulated) but what information is actually processed may be largely determined by the context and the input. Suppose the order in which information is accessed is largely driven by its accessibility in context. Given that they work quickly, it is plausible that such heuristics operate according to rules that rapidly curtail search.[73] If so, even if the process is informationally unencapsulated, any particular episode of processing will access only a tiny subset of the information that is potentially available to it. This may explain how, at the cost of susceptibility to the Moses illusion, utterance interpretation can be both unencapsulated and fast.

Acknowledgements

My thanks for very helpful comments and questions to Solveig Aasen, Georges Rey, Deirdre Wilson, two anonymous reviewers for this journal, and participants at two workshops at the University of Oslo in 2021 including Anna Drożdżowicz, Steven Gross, Anders Nes, Kim Pedersen Phillips, and Joey Pollock.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by Norges Forskningsråd: [Grant Number 302083].

[71] See Allott (2008, 270–271).
[72] In contrast, '[c]ontrolled processes can be reconfigured to deal with novel events, allowing for a far greater degree of flexibility'. (Palmeri 2003, 291).
[73] See Allott (2008, 179, 183–188, 259ff).

ORCID

Nicholas Allott http://orcid.org/0000-0001-9742-080X

References

Allott, N. 2008. "*Pragmatics and Rationality*." PhD diss., University of London.

Allott, N. 2019. "Scientific Tractability and Relevance Theory." In *In Relevance: Pragmatics and Interpretation*, edited by K. Scott, R. Carston, and B. Clark, 29–41. Cambridge: Cambridge University Press.

Allott, N., and P. Rubio Fernández. 2002. "This Paper Fills a Much-Needed gap." *Actes de L'atelier des Doctorants en Linguistique, Université Paris* 7: 97–102.

Allott, N., and D. Wilson. 2021. "Chomsky and Pragmatics." In *A Companion to Chomsky*, edited by N. Allott, T. Lohndal, and G. Rey, 433–447. Hoboken, NJ: Wiley Blackwell.

Antony, L. 2003. "Rabbit-pots and Supernovas: On the Relevance of Psychological Data to Linguistic Theory." In *Epistemology of Language*, edited by A. Barber, 47–68. Oxford: Oxford University Press.

Arnold, J. E. 2010. "How Speakers Refer: The Role of Accessibility." *Language and Linguistics Compass* 4: 187–203. https://doi.org/10.1111/j.1749-818X.2010.00193.x.

Bach, K. 2006. "The top 10 Misconceptions About Implicature." In *Drawing the Boundaries of Meaning: Neo-Gricean Studies in Pragmatics and Semantics in Honor of Laurence R. Horn*, edited by B. J. Birner, and G. L. Ward, 21–30. Amsterdam: John Benjamins.

Baillargeon, R. 2004. "Infants' Physical World." *Current Directions in Psychological Science* 13 (3): 89–94. https://doi.org/10.1111/j.0963-7214.2004.00281.x.

Bezuidenhout, A. 2001. "Metaphor and What is Said: A Defense of a Direct Expression View of Metaphor." In *Figurative Language*, edited by P. A. French, and H. K. Wettstein, 156–186. Oxford: Blackwell.

Bhalla, M., and D. R. Proffitt. 1999. "Visual–Motor Recalibration in Geographical Slant Perception." *Journal of Experimental Psychology: Human Perception and Performance* 25 (4): 1076–1096. https://doi.org/10.1037/0096-1523.25.4.1076.

Block, N. 2016. "Tweaking the Concepts of Perception and Cognition." *Behavioral and Brain Sciences* 39: e232. https://doi.org/10.1017/S0140525X15002733.

Bloom, P. 2002. "Mindreading, Communication and the Learning of Names for Things." *Mind & Language* 17 (1&2): 37–54. https://doi.org/10.1111/1468-0017.00188.

Boghossian, P. 2014. "What is Inference?" *Philosophical Studies* 169: 1–18. https://doi.org/10.1007/s11098-012-9903-x.

Borg, E. 2004. *Minimal Semantics*. Oxford: Oxford University Press.

Bredart, S., and K. Modolo. 1988. "Moses Strikes Again: Focalization Effect on a Semantic Illusion." *Acta Psychologica* 67 (2): 135–144. https://doi.org/10.1016/0001-6918(88)90009-1.

Breheny, R. 2011. "Experimentation-based Pragmatics." In *Handbook of Pragmatics: Volume 1 Foundations of Pragmatics*, edited by W. Bublitz, and N. Norrick, 561–586. Berlin: De Gruyter Mouton.

Breheny, R., H. J. Ferguson, and N. Katsos. 2013. "Taking the Epistemic Step: Toward a Model of on-Line Access to Conversational Implicatures." *Cognition* 126 (3): 423–440. https://doi.org/10.1016/j.cognition.2012.11.012.

Breheny, R., N. Katsos, and J. Williams. 2006. "Are Generalised Scalar Implicatures Generated by Default? An On-Line Investigation into the Role of Context in Generating Pragmatic Inferences." *Cognition* 100 (3): 434–463. https://doi.org/10.1016/j.cognition.2005.07.003.

Brogaard, B. 2020. "Seeing and Hearing Meanings. A non-Inferential Approach to Utterance Comprehension." In *Inference and Consciousness*, edited by A. Nes, and T. H. W. Chan, 99–124. New York: Routledge.

Brown-Schmidt, S., and D. Heller. 2018. "Perspective-taking During Conversation." In *Oxford Handbook of Psycholinguistics*, edited by S.-A. Rueschemeyer, and M. G. Gaskell, 548–572. Oxford: Oxford University Press.

Camp, E. 2006. "Contextualism, Metaphor, and What is Said." *Mind & Language* 21 (3): 280–309. https://doi.org/10.1111/j.1468-0017.2006.00279.x.

Carruthers, P. 2003. "On Fodor's Problem." *Mind & Language* 18 (5): 502–523. https://doi.org/10.1111/1468-0017.00240.

Carruthers, P. 2006. *The Architecture of the Mind: Massive Modularity and the Flexibility of Thought*. Oxford: Clarendon Press; Oxford University Press.

Carruthers, P. 2007. "Simple Heuristics Meet Massive Modularity." In *The Innate Mind: Culture and Cognition*, edited by P. Carruthers, S. Laurence, and S. Stich, 181–198. Oxford: Oxford University Press.

Carston, R. 1997a. "Enrichment and Loosening: Complementary Processes in Deriving the Proposition Expressed." *Linguistische Berichte* 8: 103–127.

Carston, R. 1997b. "Relevance-theoretic Pragmatics and Modularity." *UCL Working Papers in Linguistics* 9: 1–21.

Carston, R. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Carston, R. 2007. "How Many Pragmatic Systems are There?" In *Saying, Meaning, Referring: Essays on the Philosophy of Francois Recanati*, edited by M.-J. Frapolli, 18–48. New York: Palgrave.

Carston, R. 2010. "XIII-Metaphor: Ad hoc Concepts, Literal Meaning and Mental Images." *Proceedings of the Aristotelian Society* 110 (3): 295–321. https://doi.org/10.1111/j.1467-9264.2010.00288.x.

Carston, R. 2012. "Word Meaning and Concept Expressed." *The Linguistic Review* 29 (4): 607–623. https://doi.org/10.1515/tlr-2012-0022.

Chiappe, D. L., and A. Kukla. 1996. "Context Selection and the Frame Problem." *Behavioral and Brain Sciences* 3 (19): 529–530. https://doi.org/10.1017/S0140525X00082029.

Clark, E. V., and H. H. Clark. 1979. "When Nouns Surface as Verbs." *Language* 55 (4): 767–811. https://doi.org/10.2307/412745.

Colin, C., M. Radeau, and P. Deltenre. 2005. "Top-down and Bottom-up Modulation of Audiovisual Integration in Speech." *European Journal of Cognitive Psychology* 17 (4): 541–560. https://doi.org/10.1080/09541440440000168.

Cruse, D. A. 1986. *Lexical Semantics*. Cambridge: Cambridge University Press.

Currie, G., and K. Sterelny. 2000. "How to Think About the Modularity of Mind-Reading." *The Philosophical Quarterly* 50 (199): 145–160. https://doi.org/10.1111/1467-9213.00176.

Dehaene, S. 1998. *The Number Sense: How the Mind Creates Mathematics*. London: Allen Lane.

Drożdżowicz, A. 2023. "Do we Hear Meanings? – Between Perception and Cognition." *Inquiry* 66 (2): 196–228. https://doi.org/10.1080/0020174X.2019.1612774.

Erickson, T. D., and M. E. Mattson. 1981. "From Words to Meaning: A Semantic Illusion." *Journal of Verbal Learning and Verbal Behavior* 20 (5): 540–551. https://doi.org/10.1016/S0022-5371(81)90165-1.

Evans, J. S. B. T. 2003. "In Two Minds: Dual-Process Accounts of Reasoning." *Trends in Cognitive Science* 7 (10): 454–459. https://doi.org/10.1016/j.tics.2003.08.012.

Ferreira, F., and J. Nye. 2017. "The Modularity of Sentence Processing Reconsidered." In *On Concepts, Modules, and Language: Cognitive Science at Its Core*, edited by R. G. De Almeida, and L. R. Gleitman, 64–85. Oxford: Oxford University Press.

Fillmore, C. J. 1976. "Frame Semantics and the Nature of Language." *Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech* 280 (1): 20–32.

Firestone, C., and B. J. Scholl. 2016. "Cognition Does not Affect Perception: Evaluating the Evidence for "Top-Down" Effects." *Behavioral & Brain Sciences* 39 (e229): 1–77.

Fodor, J. A. 1983. *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, MA: MIT Press.

Fodor, J. A. 1984. "Observation Reconsidered." *Philosophy of Science* 51 (1): 23–43. https://doi.org/10.1086/289162.

Fodor, J. D., S. Nickels, and E. Schott. 2017. "Center-embedded Sentences: What's Pronounceable is Comprehensible." In *On Concepts, Modules, and Language: Cognitive Science at Its Core*, edited by R. G. De Almeida, and L. R. Gleitman, 139–168. Oxford: Oxford University Press.

Frazier, L. 1985. "Syntactic Complexity." In *Natural Language Parsing: Psychological, Computational, and Theoretical Perspectives*, edited by D. R. Dowty, L. Karttunen, and A. M. Zwicky, 129–189. Cambridge: Cambridge University Press.

Gould, S. J. 1991. *Bully for Brontosaurus: Reflections in Natural History*. London: Hutchinson Radius.

Grice, P. 1957. "Meaning." *The Philosophical Review* 66: 377–388. https://doi.org/10.2307/2182440.

Grigoroglou, M., and A. Papafragou. 2022. "Contributions of Pragmatics to Word Learning and Interpretation." In *The Oxford Handbook of the Mental Lexicon*, edited by A. Papafragou, J. C. Trueswell, and L. R. Gleitman, 421–437. Oxford University Press.

Gross, S. 2017. "Cognitive Penetration and Attention." *Frontiers in Psychology* 8 (221): 1–12.

Gross, Steven. 2022. "Is There an Empirical Case for Semantic Perception?." *Inquiry* 1–26. http://dx.doi.org/10.1080/0020174X.2022.2074099.

Harris, P. 1996. "Desires, Belief and Language." In *Theories of Theories of Mind*, edited by P. Carruthers, and P. K. Smith, 200–220. Cambridge: Cambridge University Press.

Helton, G. 2016. "Recent Issues in High-Level Perception." *Philosophy Compass* 11 (12): 851–862. https://doi.org/10.1111/phc3.12383.

Helton, G. 2018. "Visually Perceiving the Intentions of Others." *The Philosophical Quarterly* 68 (271): 243–264. https://doi.org/10.1093/pq/pqx051.

Kahneman, D. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
Katz, J. J., and J. A. Fodor. 1963. "The Structure of a Semantic Theory." *Language* 39 (2): 170–210. https://doi.org/10.2307/411200.
Kempson, R. M., and A. Cormack. 1981. "Ambiguity and Quantification." *Linguistics and Philosophy* 4 (2): 259–309.
Kissine, M., and O. Klein. 2013. "Models of Communication, Epistemic Trust, and Epistemic Vigilance." In *Social Cognition and Communication*, edited by J. P. Forgas, O. Vincze, and J. László, 139–154. Hove, East Sussex: Psychology Press.
Lederman, H. 2018. "Uncommon Knowledge." *Mind; A Quarterly Review of Psychology and Philosophy* 127 (508): 1069–1105.
Levin, D. T., and M. R. Banaji. 2006. "Distortions in the Perceived Lightness of Faces: The Role of Race Categories." *Journal of Experimental Psychology: General* 135 (4): 501–512. https://doi.org/10.1037/0096-3445.135.4.501.r
Mandelbaum, E. 2013. "Against Alief." *Philosophical Studies* 165 (1): 197–211. https://doi.org/10.1007/s11098-012-9930-7.
Mandelbaum, E. 2016. "Attitude, Inference, Association: On the Propositional Structure of Implicit Bias." *Noûs* 50 (3): 629–658. https://doi.org/10.1111/nous.12089.
Mazzarella, D. 2014. "Is inference Necessary to Pragmatics?" *Belgian Journal of Linguistics* 28: 71–95. https://doi.org/10.1075/bjl.28.04maz.
Mazzarella, D. 2015. "*Inferential Pragmatics and Epistemic Vigilance*." PhD diss., University College London.
Mazzarella, D., and I. Noveck. 2021. "Pragmatics and Mind Reading: The Puzzle of Autism (Response to Kissine)." *Language* 97 (3): e198–e210. https://doi.org/10.1353/lan.2021.0037.
Mazzone, M. 2011. "Schemata and Associative Processes in Pragmatics." *Journal of Pragmatics* 43 (8): 2148–2159. https://doi.org/10.1016/j.pragma.2011.01.009.
Mazzone, M. 2018. *Cognitive Pragmatics: Mindreading, Inferences, Consciousness*. Boston: De Gruyter Mouton.
McCourt, M. 2021. "Semantics and Pragmatics in a Modular Mind." PhD diss., University of Maryland.
Mcgurk, H., and J. Macdonald. 1976. "Hearing Lips and Seeing Voices." *Nature* 264 (5588): 746–748.
Melnikoff, D. E., and J. A. Bargh. 2018. "The Mythical Number two." *Trends in Cognitive Sciences* 22 (4): 280–293. https://doi.org/10.1016/j.tics.2018.02.001.
Mercier, H., and D. Sperber. 2009. "Intuitive and Reflective Inferences." In *Two Minds: Dual Processes and Beyond*, edited by J. S. B. T. Evans, and K. Frankish, 149–170. Oxford: Oxford University Press.
Mercier, H., and D. Sperber. 2017. *The Enigma of Reason*. Cambridge, MA: Harvard University Press.
Mikhail, J. M. 2011. *Elements of Moral Cognition: Rawl's Linguistic Analogy and the Cognitive Science of Moral and Legal Judgment*. Cambridge: Cambridge University Press.
Millikan, R. G. 1984. *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
Montalbetti, M. M. 1984. "*After Binding: On the Interpretation of Pronouns*." PhD diss., Massachusetts Institute of Technology.

Nadig, A. S., and J. C. Sedivy. 2002. "Evidence of Perspective-Taking Constraints in Children's On-line Reference Resolution." *Psychological Science* 13 (4): 329–336.

Neale, S. 2007. "Heavy Hands, Magic, and Scene-Reading Traps." *European Journal of Analytic Philosophy* 3 (2): 77–132.

Nes, A., K. Sundberg, and S. Watzl. 2023. "The Perception/Cognition Distinction." *Inquiry* 66 (2): 165–195. https://doi.org/10.1080/0020174X.2021.1926317.

Noveck, I. A. 2018. *Experimental Pragmatics: The Making of a Cognitive Science*. Cambridge: Cambridge University Press.

O'Connor, E. 2015. "*Comparative Illusions at the Syntax-Semantics Interface*." PhD diss., University of Southern California.

Ogilvie, R., and P. Carruthers. 2016. "Opening up Vision: The Case Against Encapsulation." *Review of Philosophy and Psychology* 7 (4): 721–742. https://doi.org/10.1007/s13164-015-0294-8.

O'grady, C., C. Kliesch, K. Smith, and T. C. Scott-Phillips. 2015. "The Ease and Extent of Recursive Mindreading, Across Implicit and Explicit Tasks." *Evolution and Human Behavior* 36 (4): 313–322.

Ortony, A., D. L. Schallert, R. E. Reynolds, and S. J. Antos. 1978. "Interpreting Metaphors and Idioms: Some Effects of Context on Comprehension." *Journal of Verbal Learning and Verbal Behavior* 17 (4): 465–477. https://doi.org/10.1016/S0022-5371(78)90283-9.

Palmeri, T. J. 2003. "Automaticity." In *Encyclopedia of Cognitive Science*, edited by L. Nadel, 290–301. London: Nature Publishing Group.

Park, H., and L. M. Reder. 2004. "Moses Illusion." In *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory*, edited by R. Pohl, 275–291. Hove: Psychology Press.

Phillips, C., M. W. Wagers, and E. F. Lau. 2011. "Grammatical Illusions and Selective Fallibility in Real-Time Language Comprehension." In *In Experiments at the Interfaces*, edited by J. T. Runner, 147–180. Bingley: Emerald.

Pickering, M. J., and R. P. G. van Gompel. 2006. "Syntactic Parsing." In *Handbook of Psycholinguistics*, edited by M. J. Traxler, and M. A. Gernsbacher, 455–503. Boston: Elsevier.

Pohl, R., ed. 2004. *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory*. Hove: Psychology Press.

Pylyshyn, Z. 1999. "Is Vision Continuous with Cognition?: The Case for Cognitive Impenetrability of Visual Perception." *Behavioral & Brain Sciences* 22 (3): 341–423. https://doi.org/10.1017/S0140525X99002022.

Recanati, F. 1989a. "Referential/Attributive: A Contextualist Proposal." *Philosophical Studies* 56 (3): 217–249. https://doi.org/10.1007/BF00354363.

Recanati, F. 1989b. "The Pragmatics of What is Said." *Mind & Language* 4: 295–329. https://doi.org/10.1111/j.1468-0017.1989.tb00258.x.

Recanati, F. 1993. *Direct Reference: From Language to Thought*. Oxford: Blackwell.

Recanati, F. 2002. "Does Linguistic Communication Rest on Inference?" *Mind & Language* 17 (1&2): 105–126. https://doi.org/10.1111/1468-0017.00191.

Recanati, F. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.

Rey, G. 2020. *Representation of Language: Philosophical Issues in a Chomskyan Linguistics*. Oxford: Oxford University Press.

Robbins, P. 2007. "Minimalism and Modularity." In *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics*, edited by G. Preyer, and G. Peter, 303–319. Oxford: Oxford University Press.

Sanford, A. J., and C. Emmott. 2012. *Mind, Brain and Narrative*. Cambridge: Cambridge University Press.

Sloman, S. A. 1996. "The Empirical Case for two Systems of Reasoning." *Psychological Bulletin* 119 (1): 3–22. https://doi.org/10.1037/0033-2909.119.1.3.

Sperber, D. 1994. "The Modularity of Thought and the Epidemiology of Representations." In *Mapping the Mind: Domain Specificity in Cognition and Culture*, edited by L. A. Hirschfeld, and S. A. Gelman, 39–67. Cambridge: Cambridge University Press.

Sperber, D. 2000. "Metarepresentations in an Evolutionary Perspective." In *Metarepresentations: A Multidisciplinary Perspective*, edited by D. Sperber, 117–137. New York: Oxford University Press.

Sperber, D. 2001. "In Defense of Massive Modularity." In *Language, Brain and Cognitive Development: Essays in Honor of Jacques Mehler*, edited by E. Dupoux, 47–57. Cambridge, MA: MIT Press.

Sperber, D., and D. Wilson. 1986. *Relevance: Communication and Cognition*. Oxford: Blackwell.

Sperber, D., and D. Wilson. 1987. "Authors' Response: Presumptions of Relevance." *Behavioral and Brain Sciences* 10: 736–754.

Sperber, D., and D. Wilson. 1996. "Fodor's Frame Problem and Relevance Theory (Reply to Chiappe & Kukla)." *Behavioral and Brain Sciences* 19 (3): 530–532. https://doi.org/10.1017/S0140525X00082030.

Sperber, D., and D. Wilson. 1998. "The Mapping Between the Mental and the Public Lexicon." In *Language and Thought: Interdisciplinary Themes*, edited by P. Carruthers, and J. Boucher, 184–200. Cambridge: Cambridge University Press.

Sperber, D., and D. Wilson. 2002. "Pragmatics, Modularity and Mind-Reading." *Mind & Language* 17 (1&2): 3–23. https://doi.org/10.1111/1468-0017.00186.

Stanley, J. 2005. "Hornsby on the Phenomenology of Speech." *Aristotelian Society* 79 (1): 131–145. https://doi.org/10.1111/j.0309-7013.2005.00129.x.

Swinney, D. A. 1979. "Lexical Access During Sentence Comprehension: (Re)Consideration of Context Effects." *Journal of Verbal Learning and Verbal Behavior* 18 (6): 645–659. https://doi.org/10.1016/S0022-5371(79)90355-4.

Tanenhaus, M. K., M. J. Spivey-Knowlton, K. M. Eberhard, and J. C. Sedivy. 1995. "Integration of Visual and Linguistic Information in Spoken Language Comprehension." *Science* 268 (5217): 1632–1634. https://doi.org/10.1126/science.7777863.

Tversky, A., and D. Kahneman. 1982. "Judgments of and by Representativeness." In *Judgment Under Uncertainty: Heuristics and Biases*, edited by D. Kahneman, P. Slovic, and A. Tversky, 84–98. Cambridge: Cambridge University Press.

Vaccargiu, E. 2023. "*Fitting Pragmatics Into the Human Mind: A Philosophical Investigation of the Pragmatics Module Hypothesis*." PhD diss., Università di Genova.

van Oostendorp, H., and S. De Mul. 1990. "Moses Beats Adam: A Semantic Relatedness Effect on a Semantic Illusion." *Acta Psychologica* 74 (1): 35–46. https://doi.org/10.1016/0001-6918(90)90033-C.

von Helmholtz, H. 1924. *Helmholtz's Treatise on Physiological Optics*. Rochester, N.Y.: The Optical Society of America.

Wason, P. C., and S. S. Reich. 1979. "A Verbal Illusion." *The Quarterly Journal Of Experimental Psychology* 31 (Pt 4): 591–597. https://doi.org/10.1080/14640747908400750.

Wellwood, A., R. Pancheva, V. Hacquard, and C. Phillips. 2018. "The Anatomy of a Comparative Illusion." *Journal of Semantics* 35: 543–583. https://doi.org/10.1093/jos/ffy014.

Wilson, D., and R. Carston. 2007. "A Unitary Approach to Lexical Pragmatics: Relevance, Inference and *ad hoc* Concepts." In *Pragmatics*, edited by N. Burton-Roberts, 230–259. Basingstoke: Palgrave Macmillan.

Winograd, T. 1972. "Understanding Natural Language." *Cognitive Psychology* 3 (1): 1–191. https://doi.org/10.1016/0010-0285(72)90002-3.

Zbrodoff, N. J., and G. D. Logan. 1986. "On the Autonomy of Mental Processes: A Case Study of Arithmetic." *Journal of Experimental Psychology: General* 115 (2): 118–130. https://doi.org/10.1037/0096-3445.115.2.118.