Western University

From the Selected Works of Robert J. Stainton

2017

Contextualism in Epistemology and Relevance Theory

Mark Jary

Robert J. Stainton, *University of Western Ontario*

37

CONTEXTUALISM IN EPISTEMOLOGY AND RELEVANCE THEORY¹

Mark Jary and Robert J. Stainton

I Introduction

In the context of this volume, we will assume that readers already have a reasonable understanding of what Contextualism in Epistemology says, and what work it does in the theory of knowledge. By way of a briefest reminder, and very crudely, one can read it as the conjunction of two clauses: (i) knowledge attributions may shift their truth conditions according to the standards at play in the utterance context; and as a result, (ii) one can reconcile the allure of skeptical arguments with ordinary intuition by holding that there is cross-talk. Goes the idea, the skeptic, who says we almost never know, and the anti-skeptic, who says we very frequently know, are not really disagreeing, but talking past each other because they state different things with ‘know’. Granting that this would be a happy result for epistemology, notice that Contextualism encompasses a descriptive empirical claim about linguistic meaning in context: they are not making a point (directly) about knowledge, but about the word ‘know’.

The point can be illustrated with example (1), to which we will frequently recur. Ajay is renowned for boasting about his hometown of London, Ontario: it used to have a Guy Lombardo museum; its junior hockey team, the Knights, regularly wins the Division Trophy; and it is located only two hours from Toronto. Suppose that Ajay is part of a group of students on the very first day of their very first philosophy class. The instructor asks whether one should say, of Ajay over there, that he knows that he lives in London, Ontario. The unsuspecting students agree that one should definitely assent to:

- (1) Ajay knows that he lives in London, Ontario.

The instructor then floats a variety of bizarre doubt-inducing scenarios: Descartes’s Evil Genius thought experiment, worries about dreams and brains in vats, “the impenetrable veil of perception”, etc. She asks again whether it remains correct to attribute to Ajay the knowledge that he lives in London, Ontario – and many students now respond, very plausibly, in the negative. For instance, they may suppose that Ajay lacks therein sufficient justification for the belief that he lives in London, or that he cannot rule out defeater alternatives. Now suppose that right after class, Ajay trips in the hallway, hitting his head. Paramedics arrive, and one administers tests of Ajay’s cognitive faculties. He not only successfully answers the question ‘Where do you live?’, but gives

evidence for it: that he grew up in London, his family home is still there, etc. The paramedic thus reports to her colleague that Ajay knows that he lives in London, Ontario.

Contextualism in Epistemology has this to say about the example. It does seem that attribution of knowledge goes from true to false and back to true. Yet this does not seem to arise from Ajay's evidence going down and then up again: he didn't, e.g., learn relevant new facts as he left the classroom. Must philosophers hold that (at least) one of the speakings was mistaken? No, says the Contextualist, because there is a linguistic illusion at play: the changing standards in the utterance contexts yield only *seeming* disagreement. In fact, there isn't one attribution which changes truth values; instead, there are different contents, two of which are true of the (constant and unchanging) epistemic situation, one is which is not.

Having briefly rehearsed the view, the remainder of our chapter will address a linguistic dilemma facing it. We first explain the dilemma, and then provide a response which draws on Relevance Theory and its notion of explicature (Sperber and Wilson 1986).

II A linguistic dilemma for Contextualism in Epistemology

To explain the dilemma, we need to introduce twin terms of art. Both trace to H. Paul Grice (1975) and his famous contrast between what a speaker says versus what she conversationally implicates. However, we employ our own coinage to evade distracting issues of exegesis. Simplifying for now, let us contrast two kinds of meaning: *what is said by a sentence at a context* and *what the speaker merely conveys* in addition. The former, by our stipulative definition, is arrived at by taking the standing-content of the sentence in the language, resolving ambiguities where necessary, and assigning reference to any context-sensitive indexicals such as 'I', 'there', 'now', 'today', etc. The latter kind of meaning is anything which the speaker means which extends beyond that. For example, suppose that it is July 4th, 2016. Bruce, a Seattleite, has recently broken up with his girlfriend Corinne, and she has moved away. A bit distraught and lonely, he calls Corinne's old cell phone number, reaches her, and asks whether she might consider getting back together. After some pleasantries, Corinne eventually utters the words:

(2) I live in New York now.

What is said by the sentence would be that *Corinne* lives in *New York City* on *July 4th, 2016* – taking 'New York' to designate the city rather than the state, and assigning contextually-fixed reference to 'I', 'now', and the present-tense-marker on 'live'. In addition, however, Corinne conveys-without-saying that it's not feasible to revive their romantic relationship.

Given this terminology, the central focus of the chapter can be couched in three simple premises:

- P1:** Appeal to *what is said by a sentence at a context* cannot underwrite the required shift in the truth conditions of knowledge attributions.
- P2:** Appeal to *what the speaker merely conveys* cannot underwrite the required shift in the truth conditions of knowledge attributions.
- P3:** If P1 and P2 are true, then nothing can underwrite the required shift in the truth conditions of knowledge attributions.

We will not undertake to defend P1 and P2. Their plausibility is debated in detail elsewhere in this volume. Instead, we will try to underscore what they mean by noting the kind of evidence which supports them. We begin with P1.

Given how we have defined ‘what is said by a sentence at a context’, it could underwrite a shift in truth conditions in exactly two ways, namely, if ‘know’ is ambiguous or if ‘know’ is an indexical. Now, there does seem to be a certain kind of ambiguity which infects English’s ‘know’. It has at least a propositional and an objectual use: ‘know *that* chickens lay eggs’ versus ‘know Denis’. In discussing the viability of Contextualism, the issue is whether the former itself admits of an “ordinary standards” versus “skeptical standards” ambiguity. To be clear, what is meant by “the word ‘know’ is ambiguous” herein is not that knowledge comes in degrees (as wealth and youth do); nor that there are different kinds of propositional knowledge (as there are different kinds of love and food). The issue, as we stressed at the outset, is about the word ‘know’. More specifically, to use the nomenclature of theoretical linguistics, P1 pertains to whether there is *homophony* at work: one sound corresponding to two or more words. The terms ‘lock’, ‘bank’, and ‘mine’ really are homophonous in just this way: rather than there being one term which applies to a variety of things, there are various terms which just happen to be pronounced /la:k/, /bæŋk/, and /maɪn/, respectively. (Think also of ‘to’, ‘two’, and ‘too’: this is another case of three distinct words, three homonyms – which also happen, nowadays, to have different spellings.) In this strict sense of the term, the hearer’s job, when faced with an ambiguous sound, is to discover which linguistic item was actually produced.

Is propositional ‘know’ ambiguous in this sense? One piece of evidence that it is not is that homophony is accidental. It is thus not repeated across unrelated languages. (For instance, the words for clump-of-hair and door-fastener are not pronounced the same cross-linguistically. This provides evidence that English’s ‘lock’ really is homophonous.) In contrast, the alleged accident of “ordinary standards” and “skeptical standards” being encoded in the same term is widespread across human tongues. Another bit of evidence is this: there are not two standards, but many; and, as Peter Unger (1986) stressed, they vary along a range of axes. Variation might pertain to: degree of felt certainty required; how reliable the relevant epistemic mechanism needs to be; how discriminating/sensitive the informational state must be to the environmental circumstances; the kind of evidence demanded; the alternatives to be rejected; the practical interests at play; etc. So, there would be not two lexical items corresponding to the one sound /noʊ/, but a massive number. (Compare here the propositional versus objectual uses of English ‘know’. Only two terms for the one sound are required, and the accident is not replicated cross-linguistically. In other languages, there *are* two words: ‘saber’ vs. ‘conocer’ in Spanish; ‘savoir’ vs. ‘connaître’ in French; ‘vissen’ vs. ‘kennen’ in German.)

Further evidence for P1, this time pertaining to both (alleged) ambiguity and indexicality, is that, as Stephen Schiffer (1996) has stressed, people tend to notice when truth-conditions shift because of these. If Corinne says, ‘I live near a bank’, referring to Chase-Manhattan, hearers are not tricked into supposing that ‘I don’t live near a bank’, said by our Seattleite Bruce about the Duwamish River, will be inconsistent with what Corinne stated. Indeed, we all recognize right off that conjunctions whose sound pattern is of the form “*p*, but not *p*” can be unproblematically true when the sentence *p* contains either an indexical or a homonym. Patently, ‘know’ behaves very differently: it certainly doesn’t seem to the skeptic and the anti-skeptic that they are talking past one another.² (Readers may be thinking that there are other “kinds of ambiguity” or “kinds of indexicality” which might rescue Contextualism. Quite right. We will revisit the issue in the final section of the chapter.)

Consider now P2, recalling the twin clauses of Contextualism with which we began. The first says that different knowledge attributions are made as the utterance context changes, while the second says that, as a result, there is no real disagreement regarding the attribution. Thus, both the skeptic’s and the dogmatist’s attributions can be true. It is essential to recognize that these are not merely two conjuncts of a philosophical view. They are interdependent. In particular,

and emphasizing the word ‘required’ in the linguistic dilemma, what the Contextualist needs is for knowledge attributions to shift their so-called *strict and literal truth conditions* according to the standards at play in the utterance context. This is to be sharply contrasted, for instance, with the proposal that (most) knowledge attributions are strictly speaking false (as per skepticism), but that ordinary people often convey—without saying a range of truths – such as that the addressee has reasonable grounds for believing, would not be blameworthy for acting on the belief, can be relied upon in practice as an authority, etc. Such a proposal would not reconcile the philosophical skeptic and the ordinary person. It would grant total victory to the skeptic, while offering a “pragmatic consolation prize” to common sense. (Following Keith DeRose [1995], this sort of proposal is often denigrated as a “warranted assertability maneuver”.) Yet what the speaker merely conveys precisely cannot deliver a difference in the “fully on-the-record”, “full-blown stated”, truth conditions.³

The remainder of the chapter will focus on P3. We will introduce Relevance Theory, a cognitive-representational, modular, and intention-oriented approach to linguistic communication. We will explain its notion of explicature, urging that it is an empirically supported level of meaning in between what is said by a sentence at a context and what a speaker merely conveys in addition. The existence of explicatures, we will claim, offers a way out of the linguistic dilemma. We end by clarifying the Relevance Theoretic proposal by contrasting it with other ways to resist P3.

III Overview of Relevance Theory

Relevance Theory, originally developed by Dan Sperber and Deirdre Wilson in the early 1980s, may be thought of as lying at the confluence of two 20th century philosophical traditions. Drawing on H. Paul Grice, it emphasizes meaning intentions and their recognition in communication. Human linguistic interaction is a species of ‘ostensive-inferential communication’: ostensive in that the speaker *overtly* produces a stimulus (‘to ostend’ means to exhibit, to show, e.g., by intentionally pointing); inferential in that the hearer uses the fact that this stimulus has been overtly produced as the basis on which to *surmise* the underlying reasons for speaking, thereby identifying the speaker’s intended meaning. Drawing on Noam Chomsky and Jerry Fodor, Relevance Theory takes its task to be the empirical description of the actual linguistic and psychological mechanisms at work in human communication; and the description it favours posits unconscious mental operations over representations, some of these performed in a modular way by a special-purpose language faculty. The output of that faculty is a schematic mental representation (rather than a worldly entity) which grossly underdetermines what the utterance means. This lexically encoded content, the schematic mental representation, is itself associated with other conceptual representations; and it is by means of this association that the linguistic form serves as a stimulus, activating encyclopaedic information stored at those conceptual addresses. Thus decoding, on this view, is not a method of accessing the whole intended message, but merely a means of activating, via an interaction effect of various mental modules, propositions which will serve as premises in the inferential process of identifying the speaker’s intended meaning.

Relevance Theory sets itself in opposition to the *Code Model* of linguistic communication. According to the latter, linguistic communication is achieved by the encoding of a thought, a precise message by a speaker, and its subsequent decoding by a hearer who shares that code. In contrast, Relevance Theory holds that linguistic understanding is an interaction effect, with inference always at play. Now, a wink, a shrug, or an intentionally produced grimace could serve as a stimulus in ostensive-inferential communication. But in the case of linguistic communication, the ostensive stimulus is an act of producing a token of a natural language expression, e.g.,

a sentence of English, Swahili, or Urdu. In such cases, Relevance Theorists do indeed take the linguistic signals to be decoded, yielding some communality with the Code Model. However, decoding is far from sufficient.

There are a number of ways to integrate a Gricean perspective with the sort of cognitive scientific picture found in Chomsky and Fodor. One feature which sets Relevance Theory's apart is its emphasis on cognitive efficiency. This is essential because of a pressing issue, namely how the addressee arrives at the correct interpretation, if not by a preset code. Partly, the answer involves denying the presupposition that correct interpretation always involves finding the one precise message meant. Still, how does Relevance Theory ensure that the open-ended inferences drawn are appropriate ones? This coordination problem is solved by first making a simple assumption about human cognition, and then exploiting it in the model of communication. The assumption is that human cognition is geared towards *maximising relevance* (in a technical sense of those words) by balancing cognitive benefit and effort: for any given level of effort, the greater the cognitive benefit, the greater the relevance; for any given level of cognitive benefit, the lower the effort involved in achieving that benefit, the greater the relevance. (Cognitive benefit is to be understood in terms of various sorts of improvements in the individual's representation of the world: e.g., a brand new belief may be added, or its strength may be increased; an old belief may be discarded, or its credence level decreased.) This tendency to balance benefit and effort is referred to as 'The Cognitive Principle of Relevance'. This extremely general principle about the efficiency of human cognition gets applied specifically to cognition about interlocutors. The idea is that every act of ostensive-inferential communication conveys a presumption of its own optimal relevance. This justifies the hearer in seeking to balance benefit and effort by following a path of least effort in interpreting the utterance, until he reaches a level of benefit that merits the effort taken. Once this level has been achieved, he should stop – or at least he should not assume that any further benefits derived are intended by the communicator. The assumption that each act of ostensive communication conveys a presumption of its own optimal relevance is known as 'The Communicative Principle of Relevance'.

IV Explicature and the linguistic dilemma

Our brief overview of Relevance Theory has already hinted at two levels of meaning therein, namely the decoded content and total conveyed content, i.e., the overall result of the hearer's inferences. Neither, however, will serve the purposes of refuting P3. For that, we require another crucial contrast from Relevance Theory, a distinction between two kinds of conveyed content: that which is *explicitly* conveyed and that which is *implicitly* conveyed. Relevance Theory draws this contrast in terms of propositional contents which are derived by *developing* the linguistically encoded meaning of the utterance (roughly, modulating the elements within the mental representation as decoded) versus contents which cannot be so derived. The former is called an *explicature*; any other communicated proposition is then an *implicature*. This contrast is well illustrated by an example from Robyn Carston (1988: 261):

- (3) *Denis*: How is Jane feeling after her first year at university?
Ebrahim: She didn't get enough units and can't continue.

What Ebrahim's utterance encodes is something along the lines of the mental representation in (4):

- (4) IDENTIFY A CERTAIN SALIENT FEMALE *F*. IDENTIFY UNITS OF A CERTAIN SALIENT KIND *U*. IDENTIFY A CERTAIN SALIENT ACTIVITY *A*. IDENTIFY A

CERTAIN MODULATION AND* OF CONJUNCTION. THE SPEAKER IS SAYING THAT *F* FAILED TO OBTAIN SUFFICIENT *U*'S AND* *F* IS UNABLE TO CONTINUE *A*.

Ebrahim conveys explicitly (5), while what he conveys implicitly is (6):

- (5) JANE DID NOT PASS ENOUGH UNIVERSITY COURSE UNITS TO QUALIFY FOR ADMISSION TO SECOND YEAR STUDY AND, AS A RESULT, JANE CANNOT CONTINUE WITH UNIVERSITY STUDY.
- (6) JANE IS NOT FEELING VERY HAPPY.⁴

Although the schema in (4) falls a long way short of anything Ebrahim conveys, (5) is much closer to (4) than (6) is. According to Relevance Theory, (5) is a development of (4): that is, we can get from (4) to (5) by specifying various more precise contents. In contrast, to arrive at (6), it is not enough to modulate the elements within (4). So, (6) is an implicature.

How does development to the explicature work? A key idea in recent research is *lexical adjustment* of encoded conceptual content: the encoded meaning of a lexical item can be inferentially modified in the interpretation process so that its denotation is broadened or narrowed. These processes are said to result in a thought that contains an *ad hoc* concept. Importantly for us, this *ad hoc* concept contributes towards the strictly stated truth conditions of the utterance, rather than merely contributing to what the speaker got across.

Here is an example of broadening from Deirdre Wilson (2003: 276):

- (7) The water is boiling.

The sentence in (7) might be used to claim that the water is at exactly 100 degrees centigrade, that it is at approximately 100 degrees centigrade, that it is too hot to bathe in (without any commitment to its specific temperature), or that it is behaving like boiling water in that it is bubbling ferociously. These would be, respectively, exact, approximate, hyperbolic, and metaphoric uses of 'boiling', and in each case the denotation of the term would be broader than in the previous. Now consider a case of narrowing (Wilson 2003: 278):

- (8) Churchill was a man.

An utterance of (8) might be a claim that Churchill was a male human being, but it is more likely to be a claim that he was a special type, or perhaps typical type, thereof. In these latter cases, the denotation of 'man' is narrowed, so as to exclude some male humans.⁵ Crucially, the strictly stated truth conditions of the utterance will be distinct in each case.

This is a very powerful account, and, to a large degree, its plausibility rests on how the process of lexical adjustment is constrained – so that, e.g., a speaker cannot use (1) to assert that *Queen Elizabeth fears that she'll be crushed by an asteroid*, with 'Ajay', 'know', 'live', etc., merely being "adjusted". We must therefore say a word about what constrains it. The first has been mentioned already: something is an explicature only if it traces to the decoded mental representation, so that wholly novel additions do not count. That is why (6) is not a candidate explicature. Second, the mere activation of thoughts by a linguistic stimulus is insufficient reason for the hearer to take these as intended by the hearer, and the Communicative Principle of Relevance requires them to have been so intended to be conveyed, explicitly or implicitly. Third, this principle holds not merely that the act is relevant, but that it is the most relevant stimulus available given the

communicator's abilities and preferences. (This very important point is often missed by critics.) These first three will resonate immediately with those who know Grice or have ever encountered Relevance Theory. A fourth is this: the process is constrained by the need for a development of the linguistically encoded meaning of the utterance to provide *warrant* for the implicatures of the utterance, i.e. for the explicature of the utterance to imply the implicatures. The adjustment of lexically encoded material is an outcome of this need for warrant. To return to one of our examples, consider a child uttering (7) in a complaint that her bath water is too hot for her to get in. Her use of 'boiling', predicated of water and in the scenario described, will give rise to thoughts about the discomfort and scalding associated with immersion in very hot water. This will lead to further thoughts such as *The child does not want to enter the water*. In order to get warrant to ascribe this thought the status of an implicature, the hearer need only assume that the child intends the weaker claim that the water is too hot to bathe in, rather than the stronger claim that it is at 100 degrees. Thus, despite being hyperbolic, this is the thought the hearer will treat as being strictly stated by the child's utterance; and he will only accuse the child of asserting an untruth if the water is patently not too hot to bathe in. The question of whether it is at exactly 100 degrees will not arise.⁶

In brief, the overall pattern, which keeps explicatures from "over-generating", is this: thoughts activated by the decoding of a linguistic stimulus are *candidates* for things meant by the utterance; but, the theory has it, mere activation does not warrant them as intended by the speaker; rather, warrant must be sought by developing the linguistically encoded meaning of the utterance until it implies those thoughts activated by decoding, thereby conferring on them that status:

the hearer's expectations of relevance warrant the assumption that the speaker's explicit meaning will contextually imply a range of specific consequences (made easily accessible, though not yet implied, by the linguistically encoded sentence meaning). Having identified these consequences, he may then, by a process of backwards inference, enrich his interpretation of the speaker's explicit meaning to a point where it does carry these implications.

(Wilson and Sperber 2002: 616)

The upshot of this section is that, within Relevance Theory, there exists a three-way contrast among kinds of content. There is the linguistically encoded meaning of the utterance. This is a highly schematic mental representation which typically will not express a proposition. (In itself, it does not represent part of the communicated content of the utterance.) There are the implicatures of the utterance: any communicated propositions which are not arrived at by modulating the elements in the utterance's linguistically encoded meaning. Finally, there is the explicature: an explicitly conveyed proposition which is a development of what is decoded. The essential point of our chapter is that, given this three-way contrast, Contextualism in Epistemology may be rescued from the linguistic dilemma by taking the explicature to be what shifts across utterance contexts. For this gambit to succeed, the explicature must be sharply contrasted with the notion of what is said by the sentence at the context. To reinforce their dissimilarity, notice that what is said by sentence (5) at the context will not include what kind of units Jane failed to receive, and which activity she cannot continue. Or again, consider what Relevance Theory would say about (2) above, 'I live in New York now.' Yes, the explicature of Corinne's utterance has to do with her present domicile, while the implicature pertains to the end of a romantic relationship. This is, we agree, similar to the terminological contrast introduced on P2. However, according to Relevance Theory, the explicature of Corinne's utterance of sentence (2) will extend well beyond what results from assigning reference to 'I' and 'now' and disambiguating 'New York'. For instance, the

in-context use of the linguistic expression ‘in ___ now’ will, for the Relevance Theorist, yield a specific modulation: what satisfies IN and NOW is different for LIVE IN NEW YORK NOW as opposed to BE IN BANKRUPTCY NOW or BE SINGING IN FRENCH NOW.

V Compare and contrast

We end by contrasting this Relevance Theoretic proposal to other means of rejecting P3. This will both reinforce certain points made above, and fill in some gaps.

Relevance Theory maintains that linguistic communication is inevitably modular. In particular, even fixing the explicature is an interaction effect, because an *ad hoc* concept must be developed from the decoded content. A certain kind of “context sensitivity” will thus prove ubiquitous, occurring with lexical items generally. (For instance, what counts as a satisfier of ‘weapon’, ‘dog’, or ‘water’ will vary according to context: if a son kills his mother with a knitting needle, and it is placed in a city’s museum, does that mean that the museum allows weapons on its grounds? Does a taxidermied wolf count as a dog? When you order water at the bar, would it satisfy your request for the barkeep to bring just an ice cube on a saucer?) An important result is that Relevance Theory can treat ‘know’ as a plain-old lexical item. It need not be in any respect morphosyntactically or semantically peculiar. To the contrary, if the statements/claims that can be performed with ‘know’ were not pragmatically flexible, that would be an extraordinary exception.

This contrasts with means of rejecting P3 which grant that ‘know’ in its propositional use isn’t homophonous, and grant that it is unlike ‘I’ or ‘now’, but maintain that ‘know’ is special in that it somehow “triggers” context – specifically, it triggers a search for a standard in the context of utterance. Sub-varieties of this gambit include taking ‘know’ to be a degree-theoretical word, comparable to ‘tall’ or ‘flat’; taking it to be a “contextual” along the lines of ‘enemy’ or ‘neighbor’; or taking it to be implicitly modal like ‘must’ or ‘can’.⁷

A closely related feature of Relevance Theory is that it embraces a Mentalistic “idea theory” of meaning. Semantics involves not assignment of worldly referents, but rather translation from a natural language expression to a corresponding mental representation. Indeed, this is precisely what modular decoding achieves. An upshot of this is that meanings are mere schemata, and partly procedural ones at that. Hence, though for a much more general reason, the word ‘know’ itself, the type, no more has an extension than ‘enough’ or ‘appropriate’ do: it is only the contextually-modulated *ad hoc* concept which does so. Turning to sentences, in general they are not truth-evaluable – not even relative to a limited set of contextual parameters. Thus, sentences of the form “S knows that *p*”, even after resolving ambiguities and assigning reference to indexicals, do not express propositions – they certainly do not “divide the space of worlds” into those which satisfy the sentence and those which fail to.

This contrasts with attempts to rebut P3 according to which a sentence of the form “S knows that *p*” semantically encodes a standard, loose or strict, though context can escalate/reduce that standard in such a way to change what is “full-blown stated”. This would be a variant on Minimalism about the semantics of ‘know’, conjoined with Speech Act Pluralism regarding what ‘know’ can be used to assert/claim/state. Another competitor idea is that instances of “S knows that *p*” typically do express propositions, but assigning an actual truth value requires a (varying) context of assessment for the proposition. This would be Relativism about knowledge attributions.⁸

The final contrast we wish to highlight concerns the degree of psychological commitment in play. Relevance Theory understands semantics and pragmatics to be part of cognitive psychology. It thus reads clause (i) of Contextualism as a claim in cognitive science, such that verifying

it requires empirical investigation of actual neurocognitive mechanisms, drawing on whatever evidence one can find. As we have stressed, Relevance Theory also takes a strong and specific empirical stance. Knowledge attributions may indeed shift their truth conditions according to the standards in the utterance context, because of: activation of concepts linked to the decoded mental schema; the unconscious and automatic development thereof, via the maximization of relevance; constrained by backwards inference from what would be implicated.

An overarching alternative is to de-psychologize. A moderate sub-option, which tempts one of the co-authors (namely, Jary), is to rebut P3 along broadly Relevance Theoretic lines, but without buying the entire picture. The Relevance Theoretic model outlined above has two elements: one is designed to show that it is reasonable for a hearer to interpret an utterance as he does, given the data made available to him by that act of speaking; the other is to suggest how, as a matter of cognitive psychology, he might reliably reach the intended interpretation (in a seemingly instantaneous and unreflective manner). The first element concerns the relationship between the thoughts activated by decoding the linguistic form of the utterance and the development of this form into a proposition that provides inferential warrant for those thoughts such that they are implied by the utterance. The second element involves the balance of neurocognitive effort and effect that justifies treating the thoughts activated as (being intended by the speaker as) having a role in the interpretation process. It is, we think, the *first* element that does the main work in rebutting P3, even pursuing a broadly Relevance Theoretic approach – for it is this which results in the modulation of the verb's content in the knowledge attribution. Consider how this approach would apply to the various utterances of (1). The variation in the concepts arrived at will trace to warranting potential implicatures. In the skepticism case, what is at issue is whether Ajay is justified in believing that he lives in London in the face of bizarre error-inducing scenarios. In the bang-on-the-head case, it is whether Ajay both believes that he lives in London and can deploy evidence the way that normally functioning people can. Therefore, an implicature of 'Ajay knows that he lives in London' in the skepticism scenario is that Ajay's evidence, degree of certainty, etc., defeat skeptical considerations; whereas this is not an implicature of 'Ajay knows that he lives in London' in the bang-on-the-head case, because it is irrelevant to the concerns of the speakers. Rather, the implicature is that Ajay's mental faculties are functioning normally. What the cautious shopper might do, then, is borrow from Relevance Theory the idea that the explicit content of an utterance is that proposition which is both a development of the linguistically encoded meaning of the utterance and provides warrant for its implicatures. Hence, in analysing an utterance, one would consider its overall import and seek the weakest proposition that (a) can be developed from the linguistically encoded meaning of the utterance and (b) will also warrant those propositions that contribute to the utterance's import that cannot be derived by developing its encoded meaning.

A more radical de-psychologizing alternative, consonant with Stainton (2016), would be to urge that Contextualism hardly owes debts in empirical cognitive science at all. What matters is *whether* pragmatics can impact upon what is stated with 'know', not *how*. And one can show that it does, without addressing linguistico-cognitive mechanisms, by taking the notion of the literal truth conditions of knowledge attributions to pertain to the kind/degree of commitment the attributor undertakes. This amounts to recasting the Relevance Theoretic contrast between content conveyed explicitly versus implicitly in forensic/normative terms. A shift in truth conditions is explicit, goes the idea, if it is lie-prone rather than merely misleading-prone. Consider an example. Corinne, having moved to New York, and having struggled to find a job, is on trial for attempting to rob a bank. She was caught with a loaded gun in her purse as she entered the nearby Chase-Manhattan branch, and was arrested immediately. In an attempt to convince

the grand jury that there was a giant misunderstanding, an embarrassing oversight on her part, Corinne said on the witness stand:

(9) I didn't know that there was a gun in my purse.

Suppose further that Corinne placed the firearm in her purse just a minute before entering the bank, even double-checking to make sure that it was loaded. Could she show that she did not commit perjury with (9), because an Evil Demon might have tricked her, sensory perception is always fraught, etc.? Insofar as the answer is 'No', that alone establishes that "ordinary standards" are built into the content of what Corinne *stated* with (9). Hence, even if we don't understand the mechanisms, additional empirical work in cognitive science is not required to resolve the "whether issue".

On the one hand, this thoroughly agnostic approach to the "how issue" can claim the advantage of consistency with the truth of the "cautious shopper's" view: it's simply non-committal about it. Indeed, this radical option is consistent with the truth of any of the foregoing defenses of Contextualism in Epistemology, including the full-blown Relevance Theoretic one. On the other hand, a disadvantage is that an independently motivated, detailed account of how the shift in truth conditions occurs makes it more plausible that P3 really is false. Another disadvantage of de-psychologizing, whether moderately or radically, is that one forfeits a promising reply to Schiffer's objection, viz., that interlocutors would notice the (alleged) meaning-shift. When speakers are talking past each other with indexicals and homophones, they recognize this quite quickly. Not so with "knows that —". In embracing the full Relevance Theoretic picture, one has an easy time explaining one aspect of this: as stressed above, Relevance Theory simply does not take 'know' to belong to a special class of ambiguous or context-sensitive lexical items. One does face the issue of why the interlocutors do not readily recognize that they are *lexically adjusting* in different ways. Even here, however, Relevance Theory has advantages. Development, as explained, is a modular, unconscious process which occurs automatically and rapidly. Not being a matter of personal-level reflection, it is more plausible that interlocutors might, under the right circumstances, become "blind" to it. The tricky and unnoticed shift of standards in terms of whether a use of (1) is true would be more comparable, say, to a shift wherein a jeweler, informed by a passionate youth that he simply must buy his boyfriend a gemstone ring despite limited funds, is first shown a large and lovely cubic zirconia. Having laid the trap, the salesperson then suggests that cubic zirconias "aren't really gemstones", and induces the young lover to buy a diamond. The decoded content GEMSTONE gets different lexical adjustments across the discourse, without the buyer noticing.

VI Conclusion

To sum up, we introduced the twin clauses of Contextualism in Epistemology: one says that there is a shift in truth conditions, one says that a philosophical reconciliation can be achieved as a result. We then presented a linguistic dilemma which it faces, to the effect that one cannot satisfy the first clause consistent with both the empirical facts and the philosophical demands of the second clause. Without defending them in detail, we sketched the kind of evidence in favour of two key premises of that dilemma, and then proposed a means of rejecting its third premise by drawing on Relevance Theory. Specifically, we characterized a kind of truth-conditional content, the explicature, which is both strictly stated yet modulated in a novel pragmatic way. We ended by contrasting this approach with three other strategies for resisting P3: appealing to another kind of special context-sensitivity for 'know'; assigning minimal truth conditions to

knowledge-attributing sentences while denying that they are the (only) ones asserted; and welcoming some of Relevance Theory's insights, all the while de-psychologizing them to a lesser or greater degree. The general lesson is that the linguistic dilemma can indeed be overcome.

Notes

- 1 A draft of this chapter was presented to Linguistic Talks at Western, in London, Ontario, on October 23rd, 2015. We are grateful to the audience members for discussion. We also thank Jonathan Jenkins Ichikawa, Martin Montminy, Ram Neta, Geoff Pynn, Patrick Rysiew, John Turri, and Deirdre Wilson for helpful comments. The project was supported by grants from: Research Project: Lexical Meaning and Concepts, FFI2014-52196-P, funded by the Spanish Ministry of Economy and Competitiveness, via Mark Jary; and Linguistic Pluralism, funded by the Social Sciences and Humanities Research Council of Canada, via Robert Stainton.
- 2 For additional linguistic evidence relevant to P1, see Pynn (2015a, section 5). It affords an excellent overview. For some of the original sources, see: Cappelen and Lepore (2005), Hawthorne (2004), and Stanley (2004). (Interestingly, both Cappelen and Hawthorne ultimately change their minds about how compelling some of the linguistic data is. See their 2009 book.) See also Ludlow (2005) for a novel corpus-based defense of context sensitivity of knowledge attributions.
- 3 Here is another way at the main points, in terms of recurring example. According to the Contextualist, though Ajay's epistemic state remains constant, the attributions using (1) go from true (as said by the students at the outset of the philosophy class) to false (after the introduction by the instructor of the skeptical scenarios) and back to true (as said by the paramedic in the hallway), because the relevant standards in the utterance context have risen and then dropped again. Thus, continues the idea, the dogmatic philosopher who says that all the attributions are true, and the skeptical philosopher who says that all are false, are equally mistaken. This middle-path is to be achieved by having the truth conditions of the knowledge-attribution shift. The problem posed by the dilemma is about how this can occur. P1 says, in effect, that standards cannot enter in at the level of what is said by the sentence 'Ajay knows that he lives in London, Ontario.' That's because the English word 'know' simply does not function as required: the only two linguistic mechanisms for shifting that kind of meaning, namely ambiguity and indexicality, seem to be blocked with respect to 'know'. P2 says, in effect, that standards cannot enter in at the level of what the speaker merely conveys. That's because, though this might be empirically plausible, it would not actually yield the result that the Contextualist wants. Granting that there are no other options, nothing can underwrite the required shift. Hence Contextualism as a whole is false.
- 4 To understand Relevance Theory, it is very important to distinguish representations within a spoken natural language from representations inside the mind. A central plank of the theory is that these need not correspond closely to one another. To reflect this important contrast, we here follow the usual notational custom of using lowercase letters for expressions of natural language and uppercase ones for mental representations. Thus, what the typefaces mean here is that Ebrahim spoke the English sentence 'She didn't get enough units and can't continue,' meaning to get across the mental content that Jane did not pass enough university course units to qualify for admission to second year study and, as a result, Jane cannot continue with university study.
- 5 Narrowing and broadening are not mutually exclusive. Consider 'Jane is a real princess,' which might be uttered to claim that Jane is royalty, but might also be uttered to claim that she is high maintenance, given to tantrums, and so on, without also claiming that she is blue-blooded. In the latter case, the denotation of the English word 'princess' is both broadened to include non-royalty, but also narrowed to exclude temperamentally well-balanced princesses (Carston 2002: 347).
- 6 To mention another much-discussed example, suppose Georgina asks Henri whether he is hungry. Henri replies, 'I have had breakfast.' Here, Georgina does not first arrive definitively at the explicature, find something amiss with it, and then seek out an implicature. Rather, the explicature is mutually adjusted with potential implicatures, by a "backwards inference". The potential implicature that Henri is not in fact hungry would not be licensed by the explicature HENRI HAS EATEN BREAKFAST AT SOME TIME IN HIS LIFE. That is why this explicature would be rejected.
- 7 Theorists who have been tempted to treat 'know' as context-sensitive in a special way include Cohen (1999) and Lewis (1996). Other theorists who, with Relevance Theorists, countenance full-bore, non-triggered "pragmatic determinants of literal speech act content", "free pragmatic enrichment"

- of assertions/statements, or something along these lines include: Bach (1994), Bezuidenhout (1997), Moravcsik (1998), Perry (1986), Pietroski (2005), Recanati (2002), Searle (1978), and Travis (1985).
- 8 For the general framework of Semantic Minimalism plus Speech Act Pluralism, see Borg (2004) and Cappelen and Lepore (2004, 2005). For a very clever and detailed application of the general framework to Contextualism in Epistemology, see Pynn (2015b). The Relativist approach to pragmatic enrichment has been developed in Brogaard (2008) and MacFarlane (2009). See also Montminy (2007) and Rysiew (2011). An aside. There is a line of thought according to which what is at stake is not merely how best to rebut the linguistic dilemma, but something much more pressing. Goes the idea, on the ‘knows’-has-an-extension sort of view, there will be *something* which the skeptic and the dogmatist cannot agree upon. To see why, recall both example (1) and the notion of what is said by a sentence at a context. In contrast to Relevance Theory, the ‘knows’-has-an-extension sort of view maintains that the latter is either true in all three cases, or false in all three cases. If it’s true, the dogmatist wins; if it’s false, the skeptic does. In fact, we think this introduces a red herring, because the debate is about the truth of knowledge attributions – i.e., statements, claims, assertions of knowledge – not about some *recherché* formal semanticists’ construct, what is said by the sentence at the context. Instead, to employ terminology in a non-Relevance Theoretic way, Contextualism pertains to whether the *strict and literal* content of speech acts can vary according to standards in the utterance context, and whether one can thus reconcile seemingly conflicting ones. Everyone being discussed here answers ‘yes’. See Stainton (2010) for discussion.

References

- Bach, Kent. 1994. “Conversational Implicature”. *Mind and Language* 9: 124–162.
- Bezuidenhout, Anne. 1997. “Pragmatics, Semantic Underdetermination and the Referential/Attributive Distinction”. *Mind* 106: 375–409.
- Borg, Emma. 2004. *Minimal Semantics*. Oxford: Oxford University Press.
- Brogaard, Berit. 2008. “In Defense of a Perspectival Semantics for ‘Know’”. *Australasian Journal of Philosophy* 86(3): 439–459.
- Cappelen, Herman and John Hawthorne. 2009. *Relativism and Monadic Truth*. Oxford: Oxford University Press.
- Cappelen, Herman and Ernie Lepore. 2005. *Insensitive Semantics*. Oxford: Blackwell.
- . 2004. “A Tall Tale: In Defense of Semantic Minimalism and Speech Act Pluralism”. In M. Ezcurdia, R.J. Stainton and C. Viger (eds.) *New Essays in the Philosophy of Language and Mind*. Calgary: University of Calgary Press, pp. 3–28.
- Carston, Robyn. 2002. *Thoughts and Utterances*. Oxford: Blackwell.
- . 1988. “Implicature, Explicature and Truth-Theoretic Semantics”. In R. Kempson (ed.) *Mental Representations: The Interface between Language and Reality*. Cambridge: Cambridge University Press, pp. 155–181. Reprinted in M. Ezcurdia and R.J. Stainton (eds.) (2013) *The Semantics-Pragmatics Boundary in Philosophy*. Peterborough, ON: Broadview Press, pp. 261–283.
- Cohen, Stewart. 1999. “Contextualism, Skepticism, and the Structure of Reasons”. *Philosophical Perspectives* 13: 57–89.
- DeRose, Keith. 1995. “Solving the Skeptical Problem”. *Philosophical Review* 104: 1–52.
- Grice, H. Paul. 1975. “Logic and Conversation”. In P. Cole (ed.) *Syntax and Semantics, Vol. 3: Speech Acts*. New York: Academic, pp. 41–58.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Lewis, David. 1996. “Elusive Knowledge”. *Australasian Journal of Philosophy* 74(4): 549–567.
- Ludlow, Peter. 2005. “Contextualism and the New Linguistic Turn in Epistemology”. In G. Preyer and G. Peters (eds.) *Contextualism in Philosophy: Knowledge, Meaning and Truth*. Oxford: Oxford University Press, pp. 11–50.
- MacFarlane, John. 2009. “Nonindexical Contextualism”. *Synthese* 166(2): 231–250.
- Montminy, Martin. 2007. “Epistemic Contextualism and the Semantics-Pragmatics Distinction”. *Synthese* 155(1): 99–125.
- Moravcsik, Julius M. 1998. *Meaning, Creativity, and the Partial Inscrutability of the Human Mind*. Chicago: University of Chicago Press.
- Perry, John. 1986. “Thought without Representation”. *Proceedings of the Aristotelian Society* 60: 137–152.
- Pietroski, Paul. 2005. “Meaning before Truth”. In G. Preyer and G. Peters (eds.) *Contextualism in Philosophy*. Oxford: Oxford University Press, pp. 253–300.

- Pynn, Geoff. 2015a. "Pragmatic Contextualism". *Metaphilosophy* 46(1): 26–51.
- . 2015b. "Contextualism in Epistemology". *Oxford Handbooks Online*. New York: Oxford University Press. <<http://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199935314.001.0001/oxfordhb-9780199935314-e-12>>
- Recanati, François. 2002. "Unarticulated Constituents". *Linguistics and Philosophy* 25: 299–345.
- Rysiew, Patrick. 2011. "Epistemic Contextualism". In E. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*. <<http://plato.stanford.edu/entries/contextualism-epistemology/>>
- . 2001. "The Context-Sensitivity of Knowledge Attributions". *Nous* 35: 477–514.
- Schiffer, Stephen. 1996. "Contextualist Solutions to Skepticism". *Proceedings of the Aristotelian Society* 96: 317–333.
- Searle, John. 1978. "Literal Meaning". *Erkenntnis* 13: 207–224.
- Sperber, Dan and Deirdre Wilson. 1986. *Relevance*. Cambridge, MA: Harvard University Press.
- Stainton, Robert. 2016. "Full-On Stating". *Mind and Language* 31(4), 395–413.
- . 2010. "Contextualism in Epistemology and the Context Sensitivity of 'Knows'". In M. O'Rourke and H. Silverstein (eds.) *Knowledge and Skepticism*. Cambridge, MA: MIT Press, pp. 113–139.
- Stanley, Jason. 2004. "On the Linguistic Basis for Contextualism". *Philosophical Studies* 119: 119–146.
- Travis, Charles. 1985. "On What Is Strictly Speaking True". *Canadian Journal of Philosophy* 15(2): 187–229.
- Unger, Peter. 1986. "The Cone Model of Knowledge". *Philosophical Topics* 14(1): 125–178.
- Wilson, Deirdre. 2003. "Relevance and Lexical Pragmatics". *Italian Journal of Linguistics* 15: 273–292.
- Wilson, Deirdre and Dan Sperber. 2002. "Truthfulness and Relevance". *Mind* 111: 583–632.

Taylor & Francis
Not for distribution