

1

THE VARIABILITY OF ‘KNOWS’

An opinionated overview

Crispin Wright

I The variabilist reaction against traditional epistemology

It is fair to say that from the time of the *Theaetetus* until relatively recently, theorists of knowledge tended to conceive their central task as being to explain in what knowledge consists; more exactly, to explain what further conditions need to be satisfied by a true belief if it is to count as knowl-edgeable. The widely accepted failure of the post-Gettier debates to execute this task convinc-ingly has motivated a very different tendency in mainstream contemporary epistemology. This, influentially promoted by Timothy Williamson in particular, is *epistemic primitivism*: to concede that knowledge is, as Williamson puts it, ‘prime’ – that it is a fundamental, irreducible cognitive relation. Knowledge, on the primitivist view, is a basic epistemological kind, and to know is to be in a basic, *sui generis* attitudinal state. There can therefore be no correct analysis of it in terms of other, supposedly constitutive or more fundamental cognitive states (true belief + X). The post-Gettier “X knows that P if and only if . . .” cottage industry was doomed to disappointment for this reason. To the contrary, it is in terms of knowledge that other epistemic notions – justifica-tion, evidence, warranted assertion and rational action – are to be understood.¹

This primitivism, however, still shares three traditional assumptions with the reductionism it is set against. They can be wrapped together as the compound idea that knowledge is a unique, objective, purely cognitive type of state – hence something at which the aspiration of reductive analysis could be sensibly (even if mis-) directed. If we unpack that, however, we find the follow-ing three distinct thoughts. First, ascriptions of knowledge, that X knows that P, are *contentually invariant* as far as the semantic contribution of ‘knows’ is concerned. More specifically, once the referent of ‘X’, the identity of the proposition that P and the time reference associated with ‘knows’ are settled, the result is a unique proposition, the same for any competent thinker who considers it. Second – although this would normally be taken to be entailed by the first point – this unique proposition has one and the same truth-value, no matter who asserts or assesses it. Third, this truth-value is determined purely by the cognitive achievements of the subject, irre-spective of what else, other than that part of her total information relevant to the judgement that P is true of X. In particular, such aspects as X’s (or anyone else’s) *interest* in whether P, is true, or *what is at stake* for her in its truth, or the range and specifics of counter-possibilities to P that occur, or are *salient*, to X – in short: such, as they are often described, ‘non-traditional’ or as I shall say *pragmatic* factors – have no bearing on the matter.

The striking recent tendency that provides the subject matter of this chapter is the rejection of one or more of these traditional assumptions in favour of one or another form of *variabilism*: broadly, the notion that whether an ascription of knowledge may correctly be regarded as true may depend on pragmatic factors that pertain to the circumstances of the ascriber, or to those of a third party assessing the ascription, or on pragmatic aspects of the circumstances of the ascribee. Although well short of a consensus, a considerable body of opinion has been developed that agrees that *some* form of epistemic variabilism is called for if justice is to be done to the actual employment of ‘knows’ and its cognates. In what follows, I will review some of the principal considerations that are taken to support that view, critically compare and assess some of the resulting variabilist proposals and recommend a conclusion both about them and about the prospects for primitivism.

II Three types of consideration suggestive of variability

(i) Hume remarked long ago on the contrast between the potency of skeptical doubts, at least in their subtler forms, when developed in the philosophical study and their apparent fatuity when considered in the pub over beer and backgammon.² We may of course address the tension by proposing that one or the other – study or pub – response has to be misconceived; but then we remain in a state of cognitive dissonance until we have given a convincing account of which. No need, however, for such an account if “One normally knows that one has two hands” is false in the study but true in the pub – either because the semantic value of ‘knows’ is context-sensitive (and the context shifts in relevant respects as one moves from the study to the pub), or because the truth-value of the single proposition expressed by tokenings of that sentence in both locations is not absolute or because *we* change in relevant pragmatic respects as we move from the study to the pub.

(ii) John Hawthorne and others have emphasised problems generated by our ordinary practices of knowledge-ascription for the principle of the closure of knowledge across known entailment.³ A range of cases exists where one might naturally self-ascribe knowledge of a premise of what one knows is a trivially valid entailment but might then hesitate to self-ascribe knowledge of its conclusion. Some of the most striking are so-called lottery cases. Suppose you buy a ticket in next week’s UK National Lottery (the first prize has built up to about £50 million). You are under no illusions about the odds and sensibly expect (truly, let’s suppose) that you will not win. But, despite this true belief’s being overwhelmingly strongly justified, there are powerful reasons for denying that it is, strictly, knowledgeable. For one thing, if it is knowledgeable, then your buying the ticket is irrational – but that seems a harsh verdict; indeed if the scale of the prize and the odds suitably combine, the expected utility may actually rationalise the purchase. Moreover having bought a ticket, you will have, if you *know* that it won’t win, no reason not to tear it up. But actually, once having bought a ticket, tearing it up would seem irrational so long as you have every reason to think that the lottery is fair.

Suppose it agreed that, for such reasons, you don’t strictly know that you won’t win the lottery. On the other hand, there are plenty of things that in ordinary contexts you would take yourself to know – for instance, that you won’t be able to afford to buy a new Maserati next week or to retire at the end of the current academic year – that entail that you won’t win the lottery. And in general there are plenty of things we would ordinarily be regarded as in position to know about our future circumstances in all kinds of respects (indeed, had better know if knowledge is the basis of rational practical reasoning to conclusions about what to do) that, in turn, entail that we won’t be the subject of various forms of unlikely happenstance – even in cases, like lotteries, where it is sure that someone will be – which, once contemplated, we will be inclined to acknowledge that we don’t strictly *know* will not occur.

There is the option of regarding such cases as actually challenging the validity of closure, of course. But that is a hard row to hoe.⁴ Variability offers a different recourse. Perhaps the very act of bringing to mind the conclusion of an entailment of relevant kind 'ups the ante' in some way. Maybe the correctness of your self-ascription of knowledge that you will not be able to afford to retire at the end of the current academic year is originally relativised to a range of salient counterpossibilities which do not include lottery wins and which you *are* a position to rule out – and maybe this range enlarges with the purchase of the ticket.

(iii) Perhaps the dominant motivation towards variabilism, however, springs from a range of putative linguistic 'intuitions' concerning proprieties of knowledge-ascription provoked, at least among many of the philosophers who think about them, by imaginary cases of a kind first put forward by Stewart Cohen and Keith DeRose.⁵ We can illustrate by reference to a version of DeRose's famous Bank Case. Suppose it is Friday afternoon, and Ashley and Bobbie are considering whether to bank their salary cheques. There are long queues at all the bank counters. Ashley recalls being at the bank on a Saturday morning two weeks ago and says, "Let's come back tomorrow. **I know the bank will be open tomorrow morning.**" Suppose that the bank will indeed be open on the Saturday morning.

Case 1 (*Low stakes*): Suppose that there is no particular reason to ensure that the cheques are banked sooner rather than later – say, by the following Monday. Then

Invited intuition: Ashley's recollection of Saturday morning opening two weeks ago suffices for her to speak truly.

Contrast that scenario with

Case 2 (*High stakes*): The couple's mortgage lender will foreclose unless the cheques are in the account by Monday to service their monthly repayment. Ashley and Bobbie know this. Bobbie says, "But what if the bank has changed its opening hours? Or what if the Saturday morning opening was some kind of one-off promotion?" Ashley says, "You're right. **I suppose I don't really know that the bank will be open tomorrow** (even though I am pretty confident that it will). We had better join the queue."

Invited intuition: Again, Ashley speaks truly. Too much is at stake to take the risk of e.g. a change in banking hours.

So the suggested conclusion is that "I know the bank will be open tomorrow" uttered by Ashley is true in Case 1 and false in Case 2 even though all that is different between the two are the costs to Ashley and Bobbie of Ashley's being wrong. Only the pragmatic factors have changed. Everything that might be mentioned in a traditional account of knowledge – as we would naturally say, all Ashley's relevant evidence or information – remains the same.

Two further cases may seem to prompt another important conclusion:

Case 3 (*Unknowing high stakes*): The couple's mortgage lender will indeed foreclose unless the cheques are in the account by Monday to service their monthly repayment but Ashley and Bobbie are unaware of this (they habitually leave what looks like circular mail from the mortgage company unopened and have missed the reminder). The dialogue proceeds as first described above, with Ashley asserting, "**I know the bank will be open tomorrow morning.**"

Invited intuition: This time, Ashley speaks falsely.

Compare that with

Case 4 (*Unknowing low stakes*): Ashley and Bobbie actually have no good reason to ensure that the cheques are banked before Monday but, misremembering the notice from the mortgage company, they *falsely* believe that Monday will be too late. The dialogue proceeds as in Case 2.

Invited intuition: This time Ashley's disclaimer, "**I suppose I don't really know that the bank will be open tomorrow**" is false.

The suggested conclusion from Cases 3 and 4 is this: when changes in pragmatic factors convert a true knowledge-ascription into a false one, or vice versa, it is *actual* changes that matter, rather than thinkers' impressions of what changes in such factors may have taken place.

III The varieties of variabilism

We have already, in effect, noted that the space of theoretical options here must include at least three quite different kinds of proposals: one for each of the traditional assumptions distinguished in section I. First, there is the option of maintaining that although knowledge-ascriptions are contentually invariant (in the sense there specified), the proposition thereby expressed may take different truth-values in different circumstances, depending on variation in the pragmatic factors applying to its subject, X. This is the thesis, proposed separately by Stanley and Hawthorne,⁶ that is most often termed *interest-relative invariantism* (IRI).⁷ The details of a proposal of this kind will naturally depend on just what kinds of pragmatic factor are deemed relevant – saliences seemed to be the germane factor for the issue about closure; but variation in stakes is what seems germane in the various scenarios in the Bank Case. IRI allows, apparently, that a pair of subjects may both truly believe that P on the basis of the same evidence or cognitive achievements yet one knows that P, and the other fails to know that P if they suitably differ in pragmatic respects. I'll come back to this.

Second, there is the option of maintaining that the variability in truth-value of knowledge-ascriptions across the kinds of situation illustrated is actually a product of variation in *content*. The specific version of this proposal made by DeRose and Cohen is standardly termed *ascriber contextualism* (henceforward simply 'contextualism'). In its original and basic form, this view holds that the (level of) cognitive achievement that is required of X by the truth of an utterance of "X knows that P" varies as a function of pragmatic aspects – needs, stakes, saliences – *of the speaker*. Thus, in an example like the Bank Case, variation in pragmatic aspects of a *self*-ascriber across actual, or hypothetical, cases may result in (actual or hypothetical) tokenings of "I know that P" demanding different – more or less exigent – levels of cognitive achievement if they are to count as true. The truth-conditions, hence content, of tokens of such an ascription can vary, even though the only differences in their respective contexts of utterance pertain to the situation of the speaker in purely pragmatic respects.

The third option – that of *knowledge-relativism*, fashioned on the model of assessment-sensitivity as developed by John MacFarlane⁸ – shifts the location of the pragmatic factors once again, this time to anyone who evaluates a knowledge-ascription, whether or not they are its original author. So a single token of "X knows that P" may properly be assigned different truth-values in differing contexts of assessment, whether or not distinct assessors are involved, depending on the situation in pragmatic respects of the assessor. Thus, Ashley may again quite correctly return different verdicts on a self-ascription of knowledge that the bank will open on the Saturday in the two contexts described. A smooth account of Hume's observation is likewise in prospect if the knowledge relativist can make a convincing case that travel between the philosophical study and the pub is apt to change the context of assessment in some relevant respect; a relativistic

treatment of lottery cases will require a similar story concerning the potential effects of explicit consideration of certain of a statement's consequences. But I shall not here consider in any detail how such an account might run.

It will not have escaped the attention of the alert reader that the three types of variabilist views distinguished exhibit disagreement in two dimensions. Agreeing that the truth-value of a knowledge-ascription may vary as an effect of variation in non-traditional pragmatic factors, they disagree about the *location* – subject, ascriber or assessor – of the relevant factors; but they also disagree about the *semantic significance* of such variation. For both knowledge-relativism and interest-relative invariantism, variation in pragmatic factors is of no semantic significance at all; rather, one and the same proposition gets to vary in truth-value in tandem with variation in the pragmatic characteristics of the subject or assessors of that proposition. For knowledge contextualism, by contrast, at least in its classic form, it is the proposition expressed by a particular knowledge-ascription that varies in a fashion sensitive to the pragmatic factors. Ashley's tokens of "I know the bank will be open tomorrow morning" express different propositions in the low-stakes and high-stakes scenarios outlined. Thus, conceptual space exists for three further types of views that are the duals in these two dimensions of the three distinguished. There is, first, scope for a kind of contextualism – an instance of *non-indexical* contextualism⁹ – that agrees with classical contextualism on the matter of location but disagrees on the matter of semantic significance. On this view, Ashley's two imaginary tokens of "I know the bank will be open tomorrow morning" express the same proposition in the low- and high-stakes scenarios, but this proposition takes a different truth-value as a function of the difference in what is at stake for the ascriber – Ashley – in those scenarios. Second, there is scope for a view which, like classical contextualism, regards ascriptions of knowledge as varying in their content (truth-conditions) as a function of variation in pragmatic characteristics but holds, like interest-relative invariantism, that the relevant characteristics are those not of the ascriber but of the subject, or subjects, to whom knowledge is ascribed. On such a view, a predicate of the form "... knows that P" will vary in its satisfaction-conditions rather as e.g. "... is sharp enough" so varies depending on whether it is being applied to a bread knife or a surgical scalpel. And finally, there is scope for an example of the view that *content itself* is, locally, assessment-sensitive: that what proposition is expressed by a token knowledge-ascription is itself a function of pragmatic characteristics of an assessor of it, with assessment-sensitivity of truth-value merely a consequence of such assessment-relativity of what is said.¹⁰ I do not know if anyone has ever seriously proposed a view of either of these two latter kinds for the semantics of 'knows' but in any case neither will feature further in the discussion to follow. However, in view of the difficulties, to be touched on below, that classical contextualists have encountered in trying to make good the claim that 'knows' is indeed semantically context-sensitive, its non-indexical counterpart presents as worthy of serious consideration. We'll touch on it from time to time below.

IV The location question

So, *whose* standards (salience, interests, etc.) count? The cases considered to this point involve *self*-ascriptions of knowledge. So they have the subject of the knowledge-ascription coincide with both the ascriber and an assessor. They therefore can suggest, at most, that we should be receptive to *some* sort of variabilism. They are powerless to motivate one rather than another of the variabilist views. Can we find some crucial experiments?

Here is a simple kind of case that has seemed to contextualists to favour their view over IRI:

Case 5 (*High stakes ascriber, low-stakes subject*): Ashley and Bobbie are situated as in Case 2. They ask Chris, another customer who is leaving the building, whether the bank will be open tomorrow.

Chris says “Yes, I happen to know it will – I was in here a couple of weeks ago on a Saturday.” Ashley says to Bobbie *sotto voce*, “Hmm. **That person doesn’t know any better than we do.** We had better join the queue.”

Invited intuition: Ashley speaks truly even though – as we may suppose – there is nothing at stake for Chris, the subject, in whether the bank will open on the Saturday or not. Here, it seems the interests that count are those of the ascriber, even when the subject is someone else whose interests are different (and less urgent).

The significance of this kind of case is *prima facie* countered, however, by the following simple case that may seem to point back towards IRI:

Case 6 (*Low-stakes ascriber, high-stakes subject*): Ashley, Bobbie and Chris are again situated as in Case 5. Chris is puzzled that Ashley and Bobbie have joined the queue again notwithstanding the advice they were just given about a Saturday opening and asks them about this. They explain their concern about the risk of foreclosure of their mortgage. Chris says, “OK, I understand now. I guess you guys had better not assume that the bank *will* be open tomorrow.”

Invited intuition: Chris speaks truly. But since “You know that P but had better not assume that P” is some kind of conceptual solecism, Chris’s remark is presumably a commitment to “**You do not know that the bank will be open tomorrow.**”¹¹

So, neither contextualism nor IRI does well in all the cases – in fact they do just as well and badly as each other: well enough in cases where subject and ascriber are identified, but badly in various kinds of cases where they are distinct – which are of course the crucial cases. This might encourage the thought that *both* have the location issue wrong, and one might therefore wonder whether knowledge-relativism promises an overall better ride. And indeed we can very simply modify Case 5 to get one that seems to favour knowledge-relativism over contextualism *and* IRI:

Case 5*: Ashley and Bobbie are dithering in the foyer and then merely overhear Chris (in a phone conversation) say, “Look, I don’t need to wait here now. My partner, Denny, was here a couple of weeks ago on a Saturday and can vouch that this bank will be open tomorrow.” Ashley remarks, *sotto voce*, “We can’t rely on that; **that Denny doesn’t know any better than we do.**”

Invited intuition: Ashley speaks truly.

However while knowledge-relativism may possibly best explain some intuitions in cases like this where subject, ascriber and assessor are all distinct, it faces the basic problem that it must coincide in its predictions with contextualism in any case where ascriber and assessor are one. So any two-agent problem cases for contextualism, like Case 6, are problems for relativism, too.

These conflicting intuitions present a potential paradox if we think that they do, near enough, show that there is *some* kind of relativity to pragmatic factors in the offing. How can that be so if the intuitions also suggest that each of the possible hypotheses about location is open to counterexample?

V Attempts to explain away the hard cases

Maybe (some of) the intuitions are misleading and should be explained away rather than accommodated. What might contextualism (and knowledge-relativism) say to explain away Case 6, where the intuition is that the correctness/incorrectness of the knowledge-ascription is determined by the subject’s relatively high standards, rather than the ascriber’s/assessor’s relatively low ones?

We should flag one tempting but futile tactic of explanation: that Ashley's and Bobbie's relatively high stakes and standards have an adverse effect on their confidence. It might be suggested that they do indeed not know that the bank will be open on that Saturday, as Chris's remark implies, but this is not because, as IRI would have it, the question is properly assessed by reference to their own high standards, but rather because they don't believe – or anyway *sufficiently confidently believe* – that the bank will then be open.

This suggestion has three problems. First, there is no general prohibition on the idea of a relatively diffident belief being knowledgeable. (Think of the schoolteacher's encouraging remark to a hesitant pupil: "Come on, Jonny: you *do* know the answer to this.") Second, there is anyway no need to make it a feature of the example that Ashley be in any significant degree of doubt that the bank will be open on that Saturday (and indeed I explicitly refrained from doing so, as the reader may care to check). Finally, Case 3 – involving low ascriber stakes, but *ignorant* high subject stakes – may be adapted to refurbish the objection as follows:

Case 7 Ashley and Bobbie are situated as in Case 3 – the risk of foreclosure if the cheques are not banked by Monday is real, but they are unaware of this. They ask Chris, another customer who is leaving the building, whether the bank will be open tomorrow. Chris says, "Yes, I happen to know it will – I was in here a couple of weeks ago on a Saturday." Ashley says to Bobbie, "Great. Let's get out of this and go get a coffee." Asked to explain why Ashley and Bobby have left the queue, Chris would doubtless say, **"Because they now know that the bank will be open tomorrow."**

Invited intuition: Chris speaks falsely. Given Ashley's and Bobbie's – the subjects' – actual high stakes, they are in no position to acquire knowledge by testimony from Chris, even though Chris's low-stakes self-ascription is unexceptionable.

There is, however, another response that at least one leading contextualist has offered to this kind of case that is potentially something of a game-changer. Keith DeRose observes¹² that in taking patterns of conversation like those illustrated by Cases 6 and 7 to constitute *prima facie* counterexamples to contextualism, we are implicitly taking it for granted that the mechanism whereby the context of a token knowledge-ascription contrives to set the standards for its truth is simply by identifying them with the standards of the ascriber: that "X knows that P" as uttered by Y is true just if X's relevant epistemic situation, replicated by Y but without change in the pragmatic aspects of Y's situation, would suffice for the truth of "Y knows that P". DeRose points out that there is absolutely no reason why that has to be the only kind of case. It is very familiar that in a wide range of examples – 'impure indexicals' like some personal pronouns, demonstratives and gradable adjectives – the semantic values of context-sensitive expressions featuring in particular utterances are settled as a function, in part, of the intentions of the utterer. It is therefore open to the contextualist to allow a similar role for the intentions of the author of a knowledge-ascription in determining the standard of epistemic achievement to be applied in fixing its truth-conditions. This can of course be the standard she would (take herself to) have to meet in order to satisfy the relevant ascription. But it need not be. In certain contexts – like those of Cases 6 and 7 – an ascriber may instead set a standard that defers to the needs, interest or saliences of the subject. In such a case, IRI and contextualism will coincide in their predictions of the truth-conditions of the knowledge-ascription.

I described this 'flexible contextualist' manoeuvre as a potential game-changer. It is, of course, merely *ad hoc* unless a principled and comprehensive account is provided of the conditions under which relevant variations in a speaker's intentions can be expected, enabling empirically testable predictions of variable truth-conditions. DeRose expends some effort in

that direction, to not implausible effect. His basic suggestion is that knowledge-ascriptions may be harnessed to two quite different kinds of project: whether X knows that P may be of interest because one wishes to rate X as a potential *source of information*; but it may also be of interest in the context of assessing X's performance as a *rational agent*. In the former type of case one will naturally impose standards on X's claim to knowledge appropriate to one's own needs and interests. (Just this is what seems to be happening in the high-stakes ascriber, low-stakes subject cases reviewed.) But in the latter type of case, when the focus shifts to what it is rational for X to do, it may well be (one's conception of) X's needs and interests that determine what level of cognitive achievement it is reasonable to demand if X is to be credited with the knowledge that P. And this seems to be the driver for the (invited) intuitions operative in the low-stakes ascriber, high-stakes subject cases like 6 and 7.

I have no space here to consider further whether the flexible contextualist manoeuvre can be developed so as to deliver fully satisfyingly on its initial promise. However two points are worth emphasis. The first is that an exactly analogous flexibility on the location question is, obviously, available to knowledge-relativism. Whatever potential shifts of interest are offered to explain variations in the location of standards from the point of view of a knowledge-ascriber, they will be available also to explain such variations from the perspective of a knowledge-ascription assessor. Flexibility thus offers no prospect of an advantage for contextualism over relativism. Second, there is no analogous move open to IRI, which is stuck with the idea that the standards for the truth of a knowledge-ascription are inflexibly set as a function of the needs, interests or saliences of its subject. If IRI is to restore dialectical parity after (and presuming the success of) the flexible contextualist manoeuvre, it must therefore explain away cases, like Case 5, where the location seems to go with an ascriber (or assessor), rather than the subject, as some kind of linguistic mistake. What are the prospects?

It is important to take the full measure of the challenge. Any presumed *knowledgeable* ascription of knowledge to a third party entails – by closure and factivity – an ascription of the same knowledge to oneself. And of course if IRI is right, and one's standards are relatively high, one may not have that knowledge. In that case, one won't be in position to ascribe it to a third party either, whatever their standards. There is therefore, in general, no difficulty for IRI in explaining our *reluctance to ascribe* knowledge in such cases. That, however, is not the relevant *explanandum*. What the defender of IRI has to explain – what the high-stakes ascriber, low-stakes subject examples are meant to illustrate – is a readiness of high-stakes ascribers to (falsely) *deny* knowledge that P to a relevant low-stakes subject. (Thus Ashley: “That person doesn't know any better than we do.”)

It would take us too far afield to pursue the details of all the responses that defenders of IRI have offered to this challenge.¹³ Suppose, however, that it proves that IRI has no good account to offer of the patterns of knowledge-ascription and denial that we seem to apply in certain high-stakes ascriber, low-stakes subject cases. How damaging is that? Jason Stanley¹⁴ contends that contextualism has an exactly matching set of problems as soon as we consider the relevant *self-ascription* by the low-stakes subject. Thus in Case 5 above, the subject – Chris – affirms that “I happen to know that it [the bank] will [be open tomorrow]” and, from a contextualist point of view, this self-ascription ought to be (absent any detail in the example suggesting the contrary) in perfect order: a true knowledge-ascription made by a low-stakes ascriber on adequate evidence. Yet is that not (a truth-conditional equivalent of) the very claim that Ashley, in a high-stakes context, correctly – by contextualist lights – contradicts?

Non-indexical contextualism will accept that consequence: Ashley and Chris are, in their respective contexts, perfectly correctly endorsing incompatible claims. But historically contextualists have shown no stomach for this near-enough¹⁵ relativistic stance. Rather here is a place where the putative context-sensitivity of ‘knows’ is made to do some serious theoretical work. When Chris affirms “I happen to know [that the bank will be open tomorrow]” and Ashley

affirms “That person [Chris] doesn’t know any better than we do [that the bank will be open tomorrow]”, the contradiction is finessed by the shift in the semantic value of ‘know’ engineered by the differing standards operative in their two contexts.

We’ll return to this shortly.

VI Ugly conjunctions

We have so far been concerned with the challenge to the different variabilist views to capture and explain not just some but all the pragmatically variable patterns of use of ‘knows’ and its cognates that, according to the ‘intuitions’, competent speakers seem to find acceptable. And at this point, provided they are prepared to go ‘flexible’, and thus steal the cases that otherwise favour IRI, contextualism and relativism seem to be tied in the lead. But there is also an obverse challenge: to avoid predicting uses to be acceptable which are apt to impress as anything but. How do the different theories fare on this?

IRI imposes a condition on knowledge-ascriptions as follows:

X knows that P at t is true only if X’s belief at t that P is based on cognitive accomplishments that meet standards appropriate to X’s practical interests (or whatever) at t,

and consequently appears to do very badly. Suppose X fails this condition – his practical interests are such that it is vitally important at t for him to be right about whether or not P, and he does at t truly believe that P, but does so on the basis of evidence that, though probative to a degree, impresses us as too slight to confer on him knowledge that P. Then IRI seems to treat as on an equal footing either of two remedies: X can either improve his evidence, or he can work on his practical interests in such a way that much less is at stake whether he is right about P or not. He can grow his evidence to meet the standards for knowledge imposed by his practical interests at t, or he can so modify his practical interests as to shrink, as it were, the standards of knowledge that P requires. Suppose he takes the latter course. Then a situation may arise at a later time, t^* , when we can truly affirm an ‘ugly conjunction’ like:

X didn’t (have enough evidence to) know P at t but does at t^* and has exactly the same body of P-relevant evidence at t^* as at t.

Such a remark seems drastically foreign to the concept of knowledge we actually have. It seems absurd to suppose that a thinker can acquire knowledge without further investigation simply because his practical interests happen to change so as to reduce the importance of the matter at hand. Another potential kind of ugly conjunction is the synchronic case for different subjects:

X knows that P but Y does not, and X and Y have exactly the same body of P-relevant evidence

when affirmed purely because X and Y have sufficiently different practical interests. IRI, as we noted earlier, must seemingly allow that instances of such a conjunction can be true.¹⁶

So far, so bad for IRI. But does contextualism escape any analogue of these problems for its competitor? Certainly, there can be no commitment to either form of ugly conjunction so long as we are concerned with cases where the relevant standards are set as those of an ascriber distinct from X and Y. In that case the same verdict must be returned about X at t and at t^* , or about X and Y, simply because some single set of standards is in play. But what if the context is one

where contextualism has gone *flexible*, availing itself of the licence to defer to standards set by the (changing) pragmatic characteristics of the subject(s)? In that case, *non-indexical* contextualism, at least, can offer no evident barrier to the assertibility in suitable circumstances of either type of ugly conjunction. So much is simply the price of the flexibility it appropriates to accommodate the cases that seemed to favour IRI.

Regular (indexical) flexible contextualism, by contrast, stands to suffer a commitment only to the metalinguistic counterparts:

“X doesn’t (have enough evidence to) know P” was true at *t* but “X does (have enough evidence to) know P” is true at *t*^{*} and X has exactly the same body of P-relevant evidence at *t*^{*} as at *t*;

“X knows that P” and “Y does not know that P” are both true and X and Y have exactly the same body of P-relevant evidence.

These are spared ‘ugliness’ by the postulated shifts in the semantic values of the occurrences of ‘know’ which are the trademark of the classical contextualist view and block disquotation. Nevertheless, they are unquestionably extremely strange to an English ear.

Does knowledge-relativism fare better with these potential snags? Again, the interesting question concerns a flexible relativism, one with the resources to handle cases where the pragmatic features of its subject determine the standards that a correct knowledge-ascription has to meet. And of course for the relativist, as for the non-indexical contextualist, there are no complications occasioned by shifts in the semantic value of ‘knows’. We know to expect that relativism will coincide in its predictions with non-indexical contextualism in all scenarios where knowledge is ascribed in the indicative mood and where there is no contrast between the ascriber and an assessor. It is therefore no more than the price paid for the flexibility to copy the verdicts of IRI in cases that reflect well on the latter that relativism, like non-indexical contextualism, will sanction certain cases, both synchronic and diachronic, of ugly conjunctions.

So, here is the scorecard.

IRI is, seemingly, encumbered by a commitment to the assertibility, in suitable circumstances, of both forms of ugly conjunction.

However, commitments of this kind are not, as is sometimes assumed, a distinctive problem for that particular form of variability.

Non-indexical contextualism and *relativism* both share that commitment provided they avail themselves of the option of ‘flexibility’. And, of course, if they do not so avail themselves, the IRI-favourable cases stand as counterexamples to their proposals.

Classical (flexible) contextualism is committed only to metalinguistic versions of ugly conjunctions. That is not as bad only provided (i) the metalinguistic versions are not as ugly and (ii) their disquotation is indeed blocked, i.e. provided ‘knows’ is indeed context-sensitive.

VII Is there any good reason to think that ‘knows’ is context-sensitive?

When utterances of the same type-sentence in different contexts appear to be able to take differing truth-values, context-sensitivity – that is, sensitivity of the content expressed to features of the utterance-context – is plausibly the most natural explanation. So, anyway, it must have seemed to the original authors of contextualism when first reflecting on the apparent variability of ‘knows’, but that was before the rival invariantist kinds of explanations considered here entered the scene. Can evidence be mustered to restore the presumption that context-sensitivity is at the root of the variability phenomena, and so give classical contextualism an edge?

The literature on the matter is complex, extensive and inconclusive; it is fair to say that there are no uncontroversial, or even generally agreed upon, criteria for (non-) context-sensitivity.¹⁷ Jason Stanley argues persuasively¹⁸ that the alleged context-sensitivity of 'knows' is not felicitously assimilated to that of any of gradable adjectives ('rich', 'tall'), pronouns ('I', 'you', 'this') or quantificational determiners ('all', 'many', 'some'). Schaffer and Szabo grant this but suggest instead a comparison with so called A-quantifiers ('always', 'somewhere').¹⁹ Still, there is no reason in any case why a bone fide context-sensitive expression should behave exactly like context-sensitive expressions of other kinds. Is there any *general* reason to think that 'knows' and its cognates are context-sensitive, whether or not their behaviour sustains close comparison with that of other, uncontroversially context-sensitive expressions?

Here is a natural litmus. If 'S' contains context-sensitives, then distinct tokens of 'S' in different mouths may have different truth-conditions. So distinct token questions, "S?" in the mouths of different questioners may have different conditions for affirmative answers. Hence, if 'knows' and its cognates are context-sensitive, it should be possible to design a pair of conversational contexts within which a pair of tokens of the question, "Does X know that P" presented simultaneously to a single agent – the *questionee* – can respectively properly deserve *prima facie* conflicting answers. Why simultaneously? Because 'knows' is of course context-sensitive, at least to the extent of admitting of significant tense and X, the subject, may know different things at different times. (We needn't require strict simultaneity though. It will be enough to ensure that both the questionee's and [if distinct] the subjects' information-states are relevantly unaltered throughout the interval when the two questions are put.) Why a single questionee? Because, again, we want to ensure that if different answers are appropriate to the distinct token questions, they are so because of variations in pragmatic factors determined by their respective conversational contexts, rather than variations in the information of the questionee.

Call this the *forked-tongue test*. It's pretty crude – for instance, it won't distinguish context-sensitivity from simple ambiguity. Still, its credentials as at least a necessary condition for context-sensitivity seem good. Let's construct a simple illustration. Suppose Ashley and Bobbie are wondering whether to duck out of the queues at the bank and go to get coffee and cake. Chris meanwhile, standing nearby, is on the phone to Denny. Bobbie overhears Chris say, "Yes, my dear, there is. There is a Caffè Nero just two minutes away where they serve excellent coffee and *torta di cioccolato*." Bobbie says, "Excuse me, but did you say that there is a nice coffee shop just two minutes away." Chris replies, "Ah. Actually, no. I mean: I did say that, but I was talking to my partner about a location downtown."

Thus, "just two minutes away" passes the test. It was the context of Denny's question, rather than Bobbie's, that set the reference of "just two minutes way" in Chris's original remark. When Bobbie puts a token of essentially the same type-question, the reference shifts and the correct answer changes.

Can we get a similar result with 'know'? Let's try to construct an analogously shaped case, but where the questioners' respective contexts differ in respect of the stakes they have in the truth of the answer.

In Case 8, Ashley and Bobbie are dithering in the foyer of the bank as before. They talk about the risk of foreclosure and Bobbie says, "Look, we had better ask someone." Chris and Denny, standing near the back of one of the queues, happen to overhear their conversation. Denny is also perturbed by the length of the queues and says to Chris, "Do you know if the bank will be open tomorrow? We could come back then if it will, but I'd rather not leave it till Monday since I have a hairdresser's appointment on Monday morning and am meeting Stacy for coffee and a chat in the afternoon." Chris, recalling the Saturday morning visit of two weeks earlier says, "It's OK. I happen to know the bank will be open tomorrow. I'll drive you over after breakfast." Ashley,

overhearing, says, “Excuse me, but did you say that you know the bank will be open tomorrow?” Chris, mindful of Ashley and Bobbie’s overheard priorities, replies, “Ah. Actually, no. I mean, I did use those words, but I was talking to Denny here, who has less at stake than you guys.”

Case-hardened contextualists may find this dialogue unexceptionable, but I would suggest that Denny, Ashley and Bobbie might reasonably be baffled by Chris’s last reply. It is also striking that, if the dialogue *is* regarded as unexceptionable, it should remain so if all play with ‘know’ is dropped and the operative question is rephrased as simply, “Will the bank be open tomorrow?” But in that case the explanation of the acceptability of Chris’s final remark will presumably have nothing to do with context-sensitivity in the operative question. So it looks as though the contextualist faces a choice between admitting that ‘know’ fails the forked-tongue test in this instance or insisting that it passes but that this fact has no significance for its putative context-sensitivity.

VIII Is there any good reason to think that ‘knows’ is *not* context-sensitive?

The consideration that has proved perhaps the most influential in this regard in the recent debates, and indeed has provided the prime motivation for knowledge relativism, is provided by ostensible patterns of *correction* and *retraction* that our knowledge-talk seems to exhibit. Here’s a toy example of the relevant kind. Chris and Denny have gone away for the weekend and have left Ashley and Bobbie the keys for the use of their car:

Ashley: Do you know where their car is parked?

Bobbie: Yes, I do – Chris texted me that they left it in the multi-storey lot as usual after badminton on Friday.

Ashley: But, as you very well know, there have been several car thefts in the neighbourhood recently. We should have gone to get it earlier. What if it’s been stolen?

Bobbie: I wasn’t reckoning with that. OK, I guess I don’t *know* that it is in the multi-storey lot – we had better go and check.

Here, the reader is intended to understand, Ashley’s second question doesn’t change Bobbie’s epistemic situation – it doesn’t give her any more evidence. But it does persuade her that it is appropriate to impose more demanding standards of evidence on her answer than she started out doing – and she now disavows the knowledge she originally claimed.

Now, the crucial point for the opponent of contextualism is the suggestion that this disavowal is to be understood as a *retraction*. Consider this continuation of the dialogue:

Ashley: Was your first answer, about knowing where the car is, true when you originally gave it, before I raised the possibility of the car being stolen?

and two possible responses:

Bobbie: Either (a) Sure, but I could not truly repeat the words I used, once I was reminded of the recent incidence of car-theft.

Or (b) No; as I just said, I wasn’t thinking about the possibility of the car being stolen. I shouldn’t have claimed to know that it is in the multi-storey lot.

The relativist’s idea is that contextualism ought to predict that answer (a) can be acceptable. For if the content of a knowledge-ascription is relative to standards set by the context of ascription, then suitable changes in that context may be expected to go along with a shift in content consistent

with tokens of a single type-ascription being respectively true in an original context but false in a later context. But in fact answer (a) is, on the face of it, simply bizarre, and the natural answer, in context, is answer (b), which notably not merely supplants but critiques and retracts the original. That is evidence, it is alleged, that the content of the knowledge claim has not shifted in response to the change of standards, but has remained invariant throughout.²⁰

Note that the contextualist can of course allow Bobbie to affirm not merely that she doesn't know now where the car is but that she *didn't know* when she made her first answer. That is because the referent of 'know', even as used in that past-tense claim, will – according to contextualism – have shifted to some high-standards knowledge relation in response to Ashley's invoking the possibility of theft, whereas Bobbie's original claim will have involved some different, low-standards relation. So contextualism can actually predict what *sounds like* a retraction: "I didn't know that P". What, the critic will charge, it cannot predict is agents' willingness to treat such remarks *as* retractions – their refusal to stand by the different thing that, according to contextualism, they originally said.

Challenged to explain that refusal, some contextualists²¹ have taken recourse to the idea – usually captioned (by their critics) as 'semantic blindness' – that ordinary speakers are ignorant of the context-sensitivity of 'know' and its cognates and so are prepared in certain respects to (mis)use this family of expressions to talk misguidedly as if they were not context-sensitive. It can be said in mitigation of such a move that a kind of semantic blindness was anyway part of the epistemic contextualist package from the start: after all, the contextualist has to allow that her thesis is controversial – that it is not just straight-off evident to us that 'knows' is a context-sensitive term.²² Still, it is one thing to maintain that we are ignorant about the gist of the correct semantic theory for an expression in common parlance, and another to hold that we systematically use that expression in ways that conflict with that theory; that is, that we systematically *misuse* it. It is, at least until more is said, rank bad methodology for proponents of a theory whose whole project is systematically to describe and explain aspects of our linguistic practice, to fall back too readily on the idea that aspects of that practice which fail to accord with the theory may be discounted as misuses.²³ At the least, contextualism needs to show how a prediction of the recalcitrant aspects of our practice may be elicited from its own theoretical resources. If no such account is forthcoming, the retraction data, *provided they are solid*, must constitute a serious strike against the view.

IX But are the retraction data solid?

It is, however, a further question whether our patterns of apparent retraction of knowledge claims really *do* provide the powerful argument for relativism that its supporters, notably MacFarlane, have urged. I'll canvass two doubts.

To begin with, there are issues about what exactly should count as the manifestation in practice of the relevant kind of retraction. Do we, in response to changes in pragmatic factors, really retract former ascriptions of knowledge in exactly the sense that relativism needs? A moment ago we already noted an important distinction in this connection. Consider this dialogue:

Ashley (on a fast moving train): Look, there is a cougar!

Bobbie: Where? I don't see it.

Ashley: Just there, crouching by those rocks.

Bobbie: I still don't see it.

Ashley: Oh, I am sorry. I see now that it was just a cat-shaped shadow on the rocks. There wasn't a cougar.

Here Ashley's last speech is a retraction in anyone's book: she is denying, using appropriately changed context-sensitive language, exactly the thing she originally said. But to accomplish this, it suffices merely to change the tense of the original and negate it. Whereas under the aegis of contextualism about 'knows' group ability, corresponding moves do *not* suffice for retraction of a knowledge-ascription, as we observed. Contextualism allows that Bobbie may perfectly properly admit, in response to Ashley's canvassing the possibility of car-theft, both that she does not know where the car is and *did not know when first asked*. The latter admission is not a retraction of the original claim, since – according to contextualism – it concerns a different, high-standards knowledge relation. Accordingly, the relativist needs to point to clear evidence in our linguistic practice that the disposition to retract knowledge claims when the stakes are raised goes deeper than the apparent denial involved in merely changing the tense and negating the result. Speakers will have to be reliably and regularly disposed to say things that distinguish what they are doing from such merely apparent retractions that contextualism can take in stride.

What kinds of sayings would manifest that distinction? Bobbie was presented above as doing something of the needed sort by saying, "I wasn't thinking about the possibility of the car being stolen. *I shouldn't have claimed to know that it is in the multi-storey lot.*" But that is exactly *not* what he should say on the assumption of knowledge-relativism. Relativism allows that the earlier claim, in the lower-standards context then current, can have been perfectly appropriate – indeed, from the standpoint of that context, true. So if that were the form that retractions of knowledge claims were generally to assume, the fact would be at odds with rather than advantageous to relativism. What is wanted, it seems, is a form of repudiation which is neither a simple denial, modulo any needed changes in tense, etc., nor a repudiation of the propriety of one's making the earlier claim in its original context.

The salient remaining possibility is something along the lines of "What I said before is false." So, let the relativist contention be that we are characteristically prone to retractions on this model of former knowledge-ascriptions when pragmatic factors suitably change. Unfortunately, even this pattern of retraction, should it be prevalent, is too coarse to be unpredictable by contextualism. The reason it is so is because in order to give what passes as an appropriate disquotational specification of what was said by some utterance in a previous context – "What he said before by S was that P" – it is not necessary, or indeed possible, to adjust *every* kind of context-sensitive expression that S may have contained. To be sure, if Ashley says, "Right now, I am going crazy waiting in this queue," then in order to specify what she said we'll need to shift pronouns and tenses and temporal adverbs in routine ways: what Ashley said was that, *at that time, she was* going crazy waiting in *that* queue. But this does not apply in general to, for instance, gradable adjectives nor, so the contextualist may contend, to 'knows' and its cognates. If an inexperienced hospital theatre orderly asserts, "This scalpel is very sharp," intending roughly that you could easily cut yourself if handling it carelessly, he may quite properly be reported to an expert surgeon as having said that *that particular scalpel is very sharp*, even when the context set by conversation with the surgeon is understood as one in which the notion of an instrument's sharpness is high standards – for instance, is tied to its suitability for refined neurosurgery. And in such a context, the orderly may have to accept a reprimand and allow that "What I said – *viz.* that that scalpel is very sharp – was false." In short, where some kinds of context-sensitive language are involved, admissible ways of specifying what was said are not guaranteed to deliver an actual content previously asserted rather than a counterpart spawned by differences between the original context of use and the context of the specification.

Of course it's usually easy enough to disambiguate in such cases if the conversational participants find it important to do so. The hospital orderly may (perhaps unwisely) protest that all he meant was that the scalpel had enough of a fine edge to be dangerous if handled carelessly.

Perhaps therefore the relativist argument should be that we don't go in for such disambiguation where knowledge claims are concerned but, as it were, *simply* retract. But is that true? With "sharp" now annexed to high (neurosurgical) standards, the orderly has to have recourse to other language to explain what he originally meant to say. If that is allowed to constitute sticking by his former claim, then we surely will want to say something similarly exculpatory about the credentials of our erstwhile epistemic situation and an associated knowledge claim even as we feel obliged to revoke the latter purely because of pressure of elevated standards.

It is, accordingly, open to question whether relativists have succeeded in tabling a notion of retraction with each of the needed features (a) that we do go in for retraction of knowledge claims under changes of pragmatic parameters of context, (b) that relativism predicts this and (c) that contextualism cannot predict as much.

The second doubt about the alleged pro-relativistic significance that our patterns of retraction of knowledge claims supposedly carry concerns the *extent* of the phenomenon. Relativism predicts that two contexts of assessment, c_1 and c_2 , differing only in the values of pragmatic parameters, may be such that one mandates an endorsement of a knowledge-ascription and another its repudiation. The examples so far considered have tended to focus on one direction: where a knowledge-ascription is made in a relatively low-standards context and then, apparently, retracted as the stakes rise, or certain error-possibilities become salient, or whatever the relevant kind of change is proposed to be. What about the converse direction? Does our practice pattern as relativism should expect?

Let's try an example:

Case 9 begins exactly as Case 2. It is Friday afternoon, and Ashley and Bobbie have arrived at the bank to deposit their salary cheques. However there are long queues at all the bank counters. Ashley recalls being at the bank on a Saturday morning two weeks ago and says, "Let's come back tomorrow. **I know the bank will be open tomorrow morning.**" Suppose that the bank will indeed be open on the Saturday morning. However the couple's mortgage lender has written to say the company will foreclose unless the cheques are in the account by Monday to service the monthly repayment, and Ashley and Bobbie are mindful of this. Bobbie says, "But what if the bank has changed its opening hours? Or what if the Saturday morning opening was some kind of one-off promotion?" Ashley says, "You're right. **I suppose I don't really know that the bank will be open tomorrow** (even though I am pretty confident that it will). We had better join the queue."

Invited intuition: Ashley correctly retracts her original claim. There is too much at stake to take the risk of e.g. a change in banking hours.

But now let's run the example on. Let it so happen that Eli, who is the manager of the local branch of Ashley and Bobbie's mortgage company, is also waiting in one of the queues and overhears their conversation. Remembering "that nice young couple" and taking pity on them, Eli comes across and says, "Don't worry, guys. Just between us, there is a degree of bluff about these 'final reminder' notices. We never actually foreclose without first making every effort to conduct an interview with the borrowers. It will be absolutely fine if this month's payment is serviced by the end of next week." Ashley and Bobbie are mightily relieved and Ashley says, "Aha. **So actually I did know that the bank will be open tomorrow!** Let's go and get a coffee and come back then."

Relativism predicts that Ashley's last emboldened remark is perfectly in order – indeed it expresses a commitment: the context after Eli's intervention is once again low stakes, so low standards, so Ashley's knowledge claim is now mandated by the original evidence, and the

intermediate knowledge denial should be retracted. But while relief and the decision to get a coffee are reasonable enough, Ashley's last remark is actually utterly bizarre.

This is a crucial issue for knowledge-relativism. I have no space here to pursue it in detail, but I conjecture that there are actually no clear cases where, moving from a high- to a low-standards context, and *mindful of the fact*, we are content, without acquiring any further relevant evidence, simply to retract a former knowledge-disclaimer and to affirm its contradictory. Where P was the proposition of which knowledge was denied, we may well say things like, "Well, I guess it's reasonable now if we take it that P" or "We can now probably safely assume that P". But the claim to now *know* that P will simply invite the challenge to re-confront the error-possibilities made salient in the previous high-standards context. And when the changes involved in the context shift are wholly pragmatic, we will tend to regard ourselves as, strictly, no better placed, epistemically, to discount those possibilities than we were before. For example, Ashley should not now after conversation with Eli, any more than earlier, want to claim *knowledge* that the Saturday opening of two weeks ago was not a one-off promotion.

The qualification "mindful of the fact" is crucial. No doubt it may happen that, forgetting altogether about a previous high-standards situation, we may in a new, relaxed context be prepared to make knowledge claims that contradict earlier disclaimers. But these claims will properly rank as *retractions* only if we recall the previous context and what we said then. And if we do that, recollection of the error possibilities that drove the early disclaimers is still likely to inhibit our outright claiming the relevant bits of knowledge even if it no longer seems urgent to reckon with those possibilities. Relativism, by contrast, predicts that there is now a mandate for such claims and that any such inhibitions about them conflict with the correct semantics for 'knows'.²⁴

X Conclusion

Variabilism, in all its stripes, is motivated by an *appearance*: that the language game of knowledge-ascriptions and denials incorporates a dependence of their truth-values on pragmatics – on interests, or saliences or stakes. Each of the four theoretical proposals here considered, albeit offering very different accounts of the nature of the dependence involved, takes this appearance to be veridical. If, as has been the general tendency of the foregoing discussion, none of these accounts is satisfactory – if each under-predicts (fails to predict some uses) or over-predicts (predicts uses with which we are uncomfortable) – the natural conclusion is that the appearance is *not* veridical: that our discourse involving 'knows' and its cognates is subject to no genuine pragmatics-sensitive variability of truth-conditions.

If we draw that conclusion, two possibilities remain. One, of course, is invariantism. But invariantism must come to a view about where the invariant threshold for knowledge falls, and, wherever it is placed, it will have to be acknowledged that a significant body of our knowledge claims, or knowledge disclaimers, are false, and an explanation will therefore be owing of why so much of our linguistic practice with 'knows' and its cognates falls into error. Invariantists have not been slow to respond to this challenge.²⁵ I here record the opinion, for which I have no space to argue, that to date their efforts have been unpromising.

The other possibility is a view concerning 'knows' and its cognates that stands in comparison with what deflationists about truth say about 'true'. For the deflationist about truth, very familiarly, it is a metaphysical mistake to ask after the character of the property that 'true' expresses. The proper use of the word is accountable, rather, not to the nature of an assumed referent in the realm of properties, but to the service of certain practical purposes – notably indirect endorsement and generalisation – that it enables us to accomplish. Correspondingly, a deflationism about knowledge will discharge the idea that there is any determinate epistemic relation or – in

deference to contextualism – family of relations that the proper use of 'knows' serves to record and whose character determines the truth-conditions of knowledge-ascriptions. Rather the use of the word needs to be understood by reference to the practical purposes – notably, for example, as DeRose observed, the accreditation of potential informants and the appraisal of agents' rational performance – that it enables us to accomplish. The variability phenomena surface as one or another of these purposes comes to the fore in a particular pragmatic context. But is a metaphysical mistake to project these phenomena onto the putative nature of an assumed referent, or referents, as IRI and contextualism attempt to do, and seek to explain them thereby.

Relativism doesn't make *that* mistake. Someone who holds that "X knows that P" is assessment-sensitive has already discharged the realism about the knowledge relation that deflationism would counsel us against. But if the suggestion of the preceding section about the asymmetries between our apparent retractions of knowledge-ascriptions and apparent retractions of knowledge-denials are correct, then the concept of knowledge we actually have betrays an (inflationary) invariantist tendency which relativism simply misdescribes. Of course it is open to a relativist to acknowledge this, and to present relativism as reformist. That proposal, however, stands in need of an argument that any purpose would be served by reform. The essence of the case for deflationism about 'knows' is twofold: negatively, that the combination of our tendency to allow the standards for its application to inflate indefinitely while unwilling to accept, with the skeptic, that it never applies, betrays a concept with certain inbuilt tensions and no determinate reference; positively, that the word nevertheless supplies the valuable resources that the variability phenomena reflect. The first part of that might suggest the desirability of reform, but that is compensated for by the second.

Such a general conception of knowledge – or better, of the function of 'knows' – is nothing new,²⁶ although the present suggestion, that its correctness is the principal lesson which the variability phenomena have to teach us, may be so. If it is correct, the idea that knowledge should come first in analytical epistemology could not be further from the truth. Knowledge – the presumed substantive referent of 'knows' – comes nowhere. But I must defer the further exploration of this form of deflationism to another time.²⁷

Notes

- 1 This second aspect – Williamson's "Knowledge First" programme – is of course strictly independent of and additional to the primitivism.
- 2 "The *intense* view of these manifold contradictions and imperfections in human reason has so wrought upon me, and heated my brain, that I am ready to reject all belief and reasoning, and can look upon no opinion even as more probable or likely than another. . . . Most fortunately it happens, that since reason is incapable of dispelling these clouds, nature herself suffices to that purpose, and cures me of this philosophical melancholy and delirium . . . I dine, I play a game of backgammon, I converse, and am merry with my friends; and when after three or four hours' amusement, I wou'd return to these speculations, they appear so cold, and strain'd, and ridiculous, that I cannot find in my heart to enter into them any farther" *Treatise* I, IV, 7 (Hume 1738, pp. 268–9).
- 3 Hawthorne (2004), Dretske (2005).
- 4 See Hawthorne (2005). See also ch. 12 of this volume.
- 5 Cohen (1986), DeRose (1992).
- 6 Hawthorne (2004), Stanley (2005); see also Fantl and McGrath (2007). See ch. 19.
- 7 Or sometimes: subject-sensitive invariantism.
- 8 MacFarlane (2005) and (2014). See Part V.
- 9 As MacFarlane terms it.
- 10 For experimentation with a version of this kind of view, see Cappelen (2008). Weatherston (2009) makes an interesting application of it to address certain puzzles with indicative conditionals.

- 11 This is different from – but perhaps not quite as clean-cut as – Stanley’s (2005) tactic which is to develop examples where a low-stakes ascriber does not know that the subject is high stakes. For instance, suppose Chris does not notice Ashley and Bobbie join the queue. But Denny, Chris’s partner, who has overheard the exchange, does and nudges Chris with a quizzical glance in their direction. Chris says, “Oh, I guess they must have remembered some reason why they can’t come back tomorrow – after all, **they now know that the bank will be open then.**” This time, we are supposed to have the intuition that the knowledge-ascription is false.
- 12 See DeRose (2009), ch. 7.
- 13 John Hawthorne (2004, ch. 4 at pages 162–6) attempts to enlist the help of what he calls the “psychological literature on heuristics and biases”. Hawthorne’s idea is that one lesson of this literature is that the becoming salient of a certain risk in a high-stakes situation (e.g. that of the bank’s changing its opening hours) characteristically leads us to overestimate its probability in general and hence to project our own ignorance onto subjects in low-stakes situations, too. DeRose (2009, ch. 7, section 3) counters that the phenomenon to be explained – high-stakes agents’ denial of knowledge to low-stakes subjects – extends to cases where the former take it that they *do* nevertheless know the proposition in question (because they take themselves to meet the elevated standards demanded by their high-stakes context). That seems right, but I do not see that Hawthorne needed the “projection of ignorance” component in his proposal in any case; a tendency to overestimation of the probabilities of salient sources of error would seem sufficient to do the work he wants on its own. The objection remains, however, that if an overestimation of the risk of a certain source of error underlies a high-stakes ascriber’s denial of knowledge to himself, the good standing of that denial is already compromised – whereas IRI requires precisely that the high-stakes context should validate it.
- 14 Stanley (2005), ch. 7.
- 15 Non-indexical contextualism allows speakers of the same proposition in distinct contexts to speak truly and falsely, respectively. That may seem relativistic enough for most people’s money, but it stops short of the contention, essential to MacFarlane’s understanding of relativism, that a single speaking of a proposition may take distinct truth-values as assessed in different contexts.
- 16 If evidence, too, were an interest-relative notion, then a possible direction of defence for IRI against these ugly-conjunctive commitments would be to try to make the case that variation in the interests of a subject sufficient to make the difference between her knowing that P and failing to do so must also affect what evidence she possesses, thus undercutting the assumption that evidence may remain constant for a subject at different times or for distinct subjects when their interests differ. Stanley canvasses this suggestion (2005, p. 181). It misses the nub of the difficulty, however, because there will presumably be cases where the relevant evidence is known with certainty and hence must be reckoned to be in common no matter what the practical interests of the subjects or subject at different times.
- 17 For discussion, see Cappelen and Hawthorne (2009), ch. 2.
- 18 Stanley (2005), ch. 3.
- 19 Schaffer and Szabo (2013). Their proposal deserves a properly detailed discussion. I believe the comparison is flawed but I have no space to enlarge on that here.
- 20 See ch. 20.
- 21 See, e.g. Cohen (2001).
- 22 See ch. 24.
- 23 Cf. Baker (2012).
- 24 This objection should be contrasted with another made by Montminy (2009). His contention is that when *in a high-standards context* we disclaim knowledge that P, we will also judge that we will be wrong to reclaim knowledge that P in a subsequent low-standards context, even though – he allows – that is what we will do once such a context is entered into, and relativism says we will be right to do so. I agree with the first part of that – namely, that we will take a dim view, while in the high-standards context, of the envisaged subsequent reclamation, and that since relativism says that there is nothing wrong with the subsequent reclamation, there is here a tension between something we are inclined to think and what relativism thinks we ought to think. But, unless I misread Montminy, I’m saying something different and stronger as well: namely, that we *won’t actually make a retraction* of the previous knowledge-denial when we get into the low-standards context.
Knowledge-relativism, in other words, mispredicts not just aspects of our attitudes to our practice with ‘knows’ but our practice itself. (MacFarlane responds to Montminy in section 8.6 of his 2014 book, see especially pp. 198 and following. His response does not engage the objection made here.)
- 25 See, e.g. Williamson (2005). See chs. 7, 16 and 17.

- 26 The germ is famously present in Austin (1946, pp. 97–103) where a view is outlined on which utterances of the form “I know that such-and-such” serve a *performative* rather than a descriptive function, and the function of “I know” is in effect to offer a *promise* of truth, on the basis of which others are entitled to act, form beliefs or claim to know, in turn. Austin’s ideas receive a thoroughgoing, sympathetic development in Lawlor (2013), though I do not know how far she would welcome the deflationism prefigured here.
- 27 I am grateful to Jonathan Jenkins Ichikawa for giving me the opportunity to write up and publish this material, the principal ideas in which were generated in graduate classes at New York University in 2005 and further refined at seminars at the Arché research center at St Andrews that took place as part of the AHRC-funded *Contextualism and Relativism* project (2006–9). Thanks to those involved on those occasions and to Filippo Ferrari, Patrick Greenough, Jonathan Jenkins Ichikawa, Carrie Ichikawa Jenkins and Giacomo Melis for more recent helpful discussion.

References

- Austin, J.L. (1946) “Other Minds”, *Proceedings of the Aristotelian Society*, Supplementary Volume 20: 148–187.
- Baker, C. (2012) “Indexical Contextualism and the Challenges from Disagreement”, *Philosophical Studies* 157(1): 107–123.
- Cappelen, H. (2008) “Content Relativism and Semantic Blindness”, in Garcia-Carpintero, M. and Kolbel, M. (eds.), *Relative Truth*, pp. 265–286. Oxford: Oxford University Press.
- Cappelen, H. and Hawthorne, J. (2009) *Relativism and Monadic Truth*. Oxford: Oxford University Press.
- Cohen, S. (1986) “Knowledge and Context”, *The Journal of Philosophy* 83: 574–583.
- . (2001) “Contextualism Defended: Comments on Richard Feldman’s ‘Skeptical Problems, Contextualist Solutions’”, *Philosophical Studies* 103: 87–98.
- DeRose, K. (1992) “Contextualism and Knowledge Attributions”, *Philosophy and Phenomenological Research* 52: 913–929.
- . (2009) *The Case for Contextualism: Knowledge, Skepticism, and Context: Vol. 1*. Oxford: The Clarendon Press.
- Dretske, F. (2005) “The Case Against Closure”, in Steup, M. and Sosa, E. (eds.), *Contemporary Debates in Epistemology*, pp. 27–40. Malden, MA: Wiley Blackwell.
- Fantl, J. and McGrath, M. (2007) “Knowledge and the Purely Epistemic: In Favor of Pragmatic Encroachment”, *Philosophy and Phenomenological Research* 75(3): 558–589.
- Hawthorne, J. (2004) *Knowledge and Lotteries*. Oxford: Oxford University Press.
- . (2005) “The Case for Closure,” in Steup, M. and Sosa, E. (eds.), *Contemporary Debates in Epistemology*. Malden, MA: Wiley Blackwell.
- Hume, D. (1738) *A Treatise of Human Nature* (ed. L. Selby Bigge 1888). Oxford: The Clarendon Press.
- Lawlor, K. (2013) *Assurance: An Austinian View of Knowledge and Knowledge Claims*. Oxford: Oxford University Press.
- MacFarlane, J. (2005) “The Assessment Sensitivity of Knowledge Attributions”, *Oxford Studies in Epistemology* 1: 197–233.
- . (2014) *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.
- Montminy, M. (2009) “Contextualism, Relativism and Ordinary Speakers’ Judgments”, *Philosophical Studies* 143(3): 341–356.
- Shaffer, J. and Szabo, Z. (2013) “Epistemic Comparativism: A Contextualist Semantics for Knowledge Ascriptions”, *Philosophical Studies* 168(2): 491–543.
- Stanley, J. (2005) *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Weatherson, B. (2009) “Conditionals and Indexical Relativism”, *Synthese* 166: 333–357.
- Williamson, T. (2005) “Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge”, *The Philosophical Quarterly* 55(2019): 213–235.