JONATHAN SCHAFFER

FROM CONTEXTUALISM TO CONTRASTIVISM

Contextualism treats 'knows' as an indexical that denotes different epistemic properties in different contexts. *Contrastivism* treats 'knows' as denoting a ternary relation with a slot for a contrast proposition. I will argue that contrastivism resolves the main philosophical problems of contextualism, by employing a better linguistic model. Contextualist insights are best understood by contrastivist theory.

1. CONTEXTUALISM

What is contextualism? According to Keith DeRose, "'[C]ontextualism' refers to the position that the truth-conditions of knowledge ascribing and knowledge denying sentences . . . vary in certain ways according to the context in which they are uttered" (1999, p. 187; see also DeRose, 1992, p. 914) DeRose is on the right track here, but what he means when he says that the truth-conditions 'vary in certain ways' needs clarification. After all, even the most rabid anti-contextualist should think that the truth-conditions of an utterance of 'Moore knows that he has hands' vary in some ways, since the denotation of 'Moore' is context-variable. The needed clarification is that, for the contextualist, at least some of the variability in truth-conditions is traceable to the occurrence of 'knows'.

So what is contextualism? Contextualism is the theory that 'knows' contributes semantical context-dependence to utterances in which it occurs (beyond any context-dependence due to other components of what is said). This raises three clarificatory questions: (i) How does 'knows' contribute semantical context-dependence? (ii) What context-dependence does 'knows' contribute? (iii) Why believe that 'knows' actually works this way?

So first, how does 'knows' contribute semantical context-dependence? That is, what is the semantic mechanism involved? Here

Philosophical Studies **119:** 73–103, 2004.

Stewart Cohen is explicit: "The theory ... construes 'knowledge' as an indexical. As such, one speaker may attribute knowledge to a subject while another speaker denies knowledge to that same subject, without contradiction" (1988, p. 97; see also DeRose, 1992, p. 920; Ram Neta, forthcoming b; John Hawthorne *ms*; *inter alia*) What is meant by calling 'knows' an indexical is that 'knows' fits David Kaplan's (1977) model for indexicals, in having a context-invariant *character*, which is a function from context to *content* (DeRose, 1992, p. 921; Hawthorne, *ms*). So just as 'I am in the Haymarket café' may be true said by Ann but false said by Ben, so 'Moore knows that he has hands' may be true said in the courtroom but false said in the classroom, because both 'I' and 'knows' are associated with semantical rules that generate context-variable content. Thus contextualists maintain:

(1) *Indexicality*: 'Knows' is an indexical.

Turning to the second question, what context-dependence does 'knows' contribute? That is, if 'knows' is an indexical, what is its character? Here David Lewis says: "*S knows* that *P* iff *S*'s evidence eliminates every possibility in which not-*P* – Psst! – except for those possibilities that we are properly ignoring" (1996, p. 425). Lewis's point is clear enough, though strictly speaking the phrasing should be meta-linguistic:[1] an utterance of the type '*s* knows that *p*' is true in context *c* iff (i) *R* is the set of worlds relevant in *c*, and (ii) *s*'s evidence eliminates every $\sim p$-world in *R*. Thus the question of content, for Lewis, becomes the question, "Which [possibilities] are the 'relevant alternatives'?" (p. 426). So just as the character of 'I' is a function from context to an individual by the rule *speaker of the utterance*, so the character of 'knows' is a function from context to an epistemic property by the rule *relevant alternatives*. Here Mark Heller is explicit: "Every selection of worlds defines a specific epistemic property. It is these various properties that are the possible referents for the term 'knowledge'. It is context which determines which of them is referred to on a particular occasion" (Heller, 1999, p. 117; see also Cohen, 1988, p. 96). Thus contextualists maintain:

(2) *Relevance*: The content of 'knows' is determined by the relevant alternatives.[2]

Moving on to the third and final question here, why believe that 'knows' actually works this way? That is, why should one believe that 'knows' is an indexical whose content is determined by the relevant alternatives? Contextualists offer three related arguments. The first (and most prominent) argument is that contextualism shields the immovable object of dogmatic belief from the unstoppable force of skeptical argument. As DeRose says:

> [O]ur ordinary claims to know can be safeguarded from the apparently powerful attack of the skeptic, while, at the same time, the persuasiveness of the skeptical argument is explained. For the fact that the skeptic can invoke very high standards that we don't live up to has no tendency to show that we don't satisfy the more relaxed standards that are in place in more ordinary conversations and debates. (1995, p. 5; see also Lewis, 1979, pp. 245–246; Cohen, 1999, pp. 66–67; Heller, 1999, p. 121; *inter alia*)

The idea here is that the dogmatist is right in ordinary context *c1* to say (i) 'Moore knows that he has hands', while the skeptic who brings up extraordinary possibilities such as being a brain-in-a-vat is right in such an extraordinary context *c2* to conclude (ii) 'Moore does not know that he has hands'. The skeptic is just wrong to think that these conflict. There is no more conflict between (i) and (ii) then there would be between Ann's saying (iii) 'I am at the Haymarket café', and Ben's saying (iv) 'I am not at the Haymarket café'. Any appearance of conflict is due to an equivocation on the reference of an indexical (be it 'knows' or 'I'). Thus contextualists maintain:

(3) *Equivocationism*: The contextually variable content of 'knows' reveals how skeptical arguments are equivocal.

There are two other (less prominent) arguments for contextualism that one finds in the literature. The second argument is that contextualism explains the shifty acceptability of knowledge ascriptions (DeRose, 1992, pp. 913–918; Cohen, 1999, pp. 57–60). It is intuitively evident that some knowledge ascriptions are shifty. For instance, a glance at the clock will ordinarily suffice to make it acceptable to say 'I know that it is 3:15'. But if the accuracy of the clock is in question, a glance at the clock will not suffice to make that ascription acceptable – further corroboration will be required. Contextualism provides a straightforward semantical explanation for such shiftiness: the question of the accuracy of the clock intro-

duces a further relevant alternative, thereby shifting the denotation of 'knows' to a more demanding epistemic property.[3]

The third argument for contextualism is that it suits 'knows' to the purposes of current inquiry (Neta, forthcoming a). Inquiries may be modeled as multiple-choice slates: inquiries are governed by questions,[4] and questions present alternatives.[5] Indexing 'knows' to the relevant alternatives suits 'knows' to multiple-choice purposes, because the relevant-alternatives set for 'knows' sways with the multiple-choice slate of inquiry. Suppose, for instance, that Mikey can discriminate Coke from Sprite but not from Pepsi. Then Mikey will (i) succeed if the inquiry concerns whether the drink is Coke or Sprite, but (ii) fail if the inquiry concerns whether the drink is Coke or Pepsi. So does Mikey know that the drink is Coke? Contextualism suits 'knows' to the purposes of current inquiry by allowing 'Mikey knows that the drink is Coke' to come out (i) true when the relevant alternative is that the drink is Sprite, but (ii) false when the relevant alternative is that the drink is Pepsi.

This concludes my characterization of contextualism. I should note two qualifications by way of conclusion. The first qualification is that the contextualist theses (1)–(3) are *neutral as to the analysis of knowledge*. Indeed there are substantive disagreements among contextualists on this point.[6] My aim is not to adjudicate between different versions of contextualism, but rather to compare the general contextualist framework of (1)–(3) to the general contrastivist framework, which will turn out equally neutral. So I will not discuss this issue here.

The second qualification to note is that my characterization of contextualism is *partly stipulative*. The literature hosts a family of theories that bear the name 'contextualism', which share little more than a pattern of familial resemblance.[7] My aim is not to engender verbal dispute, but rather to compare theses (1)–(3) to their contrastivist counterparts, which will turn out subtly different. So I would simply drop the term 'contextualism' should it prove contentious.

2. CONTRASTIVISM

What is contrastivism? *Negative characterization*: It is widely assumed that knowledge is a binary relation of the form K*sp*. The

project of understanding knowledge then becomes the project of completing the schema '*s* knows that *p* iff ...' The traditional epistemologist supposes that there is *one* K*sp* relation, while the contextualist suggests that there are *many*. The contrastivist says there is *none* – the assumption that knowledge is a binary relation is an error due to the seductive pull of the surface grammar of a special form of utterance.

Positive characterization: Contrastivism is the view that knowledge is a ternary relation of the form K*spq*, where *q* is a contrast proposition.[8] This raises three clarificatory questions: (i) How does *q* factor into knowledge ascriptions? (ii) What determines the value of *q*? (iii) Why believe that *q* actually exists?

So first, how does *q* factor into knowledge ascriptions? That is, what is the relation between K*spq* and 'knows'? The contrastivist has a range of options here, including: (a) treating 'knows' as lexically and thus syntactically ternary, (b) treating 'knows' as syntactically binary but requiring semantical supplementation in order to generate a proposition (this is akin to what Kent Bach (1994) calls 'completion'); (c) treating 'knows' as semantically binary but inviting pragmatic embellishment in order to make a decently specific claim (this is akin to Bach's notion of 'expansion' by what is implicit); and (d) treating 'knows' as fully binary but conceptually inferior, and deserving of replacement by a ternary 'knows*' (Sinnot-Armstrong, *ms*). In any case, the contrastivist maintains:

(4) *Ternicity*: 'Knows' denotes a three-place relation K*spq*.

For the sake of definiteness, I will pursue option (a), and explore the idea that 'knows' is lexically ternary. Since lexical structure projects onto syntactic structure (Chomsky, 1981; see also Hale and Keyser, 1993), this implies that all knowledge ascriptions contain a syntactically real contrast variable *q* in their logical forms. Some may dismiss the hypothesis of a hidden contrast variable out of hand (though still accepting one of the other options for contrastivism above). But I regard posits of syntactic variables as a subtle matter calling for subtle diagnostics.

Here are five arguments that *q* is syntactically real. First, there is the argument from *contrastive ascriptions*. One diagnostic for covert variables is the existence of overt counterparts that articulate

the same lexical items plus an extra argument place. For instance, one reason for thinking that 'prefers' denotes a ternary relation, despite the existence of binary preference ascriptions such as 'Ann prefers chocolate', is the existence of overtly contrastive preference ascriptions such as 'Ann prefers chocolate rather than vanilla'. Knowledge ascriptions have overtly contrastive forms too. There are binary ascriptions such as 'Holmes knows that Mary stole the bicycle', but there are also contrastive ascriptions such as 'Holmes knows that Mary stole the bicycle rather than the wagon', as well as interrogative ascriptions such as 'Holmes knows what Mary stole' that embed questions denoting sets of alternatives (§1). To consider only the binary ascriptions is to commit a sampling error.[9]

Second, there is the argument from *binding*. Another diagnostic for covert variables is the existence of binding effects given the right syntactic conditions (namely, co-indexing and c-command). For instance, in the case of 'Ann prefers chocolate', one can prefix a quantifier that binds the foil such as in: 'In every ice cream parlor, Ann prefers chocolate'. This has a bound reading on which it says, e.g., that in parlor1, Ann prefers chocolate rather than parlor1's other flavors; in parlor2, Ann prefers chocolate rather than parlor2's other flavors; etc. Knowledge ascriptions also generate contrast-bound readings. Suppose one says, 'On every test, Mikey knows that the drink is Coke'. This has a bound reading on which it says: on the first test, Mikey knows that the drink is Coke rather than Sprite; on the second test, Mikey knows that the drink is Coke rather than Pepsi; etc. Explanation: the quantifier is binding q.[10]

Third, there is the argument from *ellipsis*. A further diagnostic for covert variables is the existence of preservation effects into ellipsis sites. For instance, in the case of 'Ann prefers chocolate', one can suffix an ellipsis site that copies the foil such as in 'Ann prefers chocolate, and Ben does too'. Here if what Ann prefers is chocolate rather than vanilla, then the Ben conjunct is true iff Ben prefers chocolate to vanilla (Ben must prefer what Ann prefers). Knowledge ascriptions also feature contrast-preservation under ellipsis. Suppose that an inquiry is underway at to whether Mary stole the bicycle or the wagon, and one says 'Holmes knows that Mary stole the bicycle, and Watson does too'. Here what Holmes knows is that Mary stole the bicycle *rather than the wagon*, and the evidence for

the syntactic reality of the contrast is that the Watson conjunct is true iff Watson knows that Mary stole the bicycle rather the wagon. Explanation: the copying mechanism is reading q.[11]

Fourth, there is the argument from *focus*. Yet another diagnostic for covert variables is the existence of semantic sensitivity to focal differences. For instance, 'Ed prefers *drinking* tea' and 'Ed prefers drinking *tea*' differ in truth-conditions: if Ed's overall preference ranking is drinking coffee, then drinking tea, then bathing in tea, then the former is true but the latter false. Knowledge ascriptions also generate semantical focus-sensitivity, as noted by Fred Dretske:

> Someone claiming to know that Clyde *sold* his typewriter to Alex is not (necessarily) claiming the same thing as one who claims to know that Clyde sold his typewriter *to Alex* . . . A person who knows that Clyde *sold* his typewriter to Alex must be able to rule out the possibility that he *gave* it to him, or that he *loaned* it to him . . . But he needs only a nominal justification, if he needs any justification at all, for thinking it was Alex to whom he sold it. (1981, p. 373; see also Dretske, 1972)

Explanation: the focus differences are inducing differences in q. To know that Clyde *sold* his typewrite to Alex is to know that Clyde sold his typewriter to Alex rather than that he gave it to him or loaned it to him: Kspq_1; while to know that Clyde sold his typewriter *to Alex* is to know that Clyde sold his typewriter to Alex rather than Bonnie: Kspq_2.[12]

Fifth and finally, there is the argument from *surface paradoxes*. One last diagnostic for covert variables is the existence of surface paradoxes. Suppose that Ann's overall preference ranking is chocolate, then vanilla, then strawberry. This generates a surface paradox between 'Ann prefers vanilla' and 'Ann does not prefer vanilla', which may be resolved by noting that Ann does prefer vanilla – rather than strawberry, but Ann does not prefer vanilla – rather than chocolate. Knowledge ascriptions also resolve surface paradoxes. Suppose that Mikey can discriminate Coke from Sprite but not from Pepsi. This generates a surface paradox between 'Mikey knows that the drink is Coke' and 'Mikey does not know that the drink is Coke', which may be resolved by noting that Mikey knows that the drink is Coke – rather than Sprite, but Mikey does not know that the drink is Coke – rather than Pepsi. Explanation: the surface paradoxes are due to shifts in q.

Details aside, whatever story one prefers for 'prefers', tell that story for 'knows'.

Turning to the second question here, what determines the value of q? Contrastive, interrogative, and focused ascriptions feature explicit 'rather than'-clauses, embedded questions, and focus variables (respectively) that denote alternatives and thus explicitly determine q. What about binary ascriptions? In binary ascriptions q functions as a free variable – q may either find an antecedent or receive a pragmatically constructed value. Where there is a specific inquiry in the preceding discourse, then the alternatives under inquiry may serve as the antecedent for q. In general one might think of the dynamics of discourse in terms of what Robert Stalnaker (1998) calls the *context set*, which contains "all the situations among which the speakers intend to distinguish with their speech acts" (p. 99). The context set may serve as the antecedent for q. If that fails then the value of q must be pragmatically constructed.[13] What emerges from the range of knowledge ascriptions is a variety of linguistic mechanisms for saturating a contrast slot. What emerges is a ternary relation in various linguistic guises. Thus the contrastivist concludes:

(5) *Saturation*: The value of q is set by standard mechanisms for saturating argument places.

Moving on to the third and final question here, why believe that q actually exists? The contrastivist offers four related arguments. First, as explored above, the best linguistic treatment of knowledge ascriptions may require q as syntactically real.

The second argument for q is that q allows for the reconciliation of dogmatic knowledge with skeptical doubt. Does Moore know that he has hands? The dogmatist answers *yes*: Moore's hands are right before his eyes. The skeptic answers *no*: for all Moore knows he could be a brain-in-a-vat. The recognition of q reconciles these partly plausible answers by answering *yes and no*: *yes*, Moore knows that he has hands rather than stumps; but *no*, Moore does not know that he has hands rather than vat-images of hands. Thus the contrastivist maintains:

(6) *Compatibilism*: Contrastivism renders dogmatic knowledge and skeptical doubt compatible, as concerning different q-values.

The third argument for q is that q explains acceptability differences among knowledge ascriptions, including the shifty acceptability of binary knowledge ascriptions (§1). Acceptability differences are explained via shifts in the proposition expressed: Kspq_1 versus Kspq_2. Contrastivism thus provides a straightforward semantical explanation for why 'I know that it is 3:15' is ordinarily acceptable on the basis of a glance at the clock, but not when the accuracy of the clock is in question. The explanation is that in the ordinary case the contrast does not include the clock-error possibility, while the introduction of the question of accuracy shifts the implicit contrast to include such. Moreover, contrastivism provides a straightforward and uniform explanation for why differences in 'rather than'-arguments, in queried alternatives, and in focused constituents generate acceptability differences too. Suppose Holmes has found Mary's fingerprints at the crime scene, but has not bothered to verify that it was a bicycle that she stole. Then intuitively 'Holmes knows that Mary rather than Peter stole the bicycle' is true, but 'Holmes knows that Mary stole the bicycle rather than the wagon' is false; 'Holmes knows who stole the bicycle is true', but 'Holmes knows what Mary stole' is false; and 'Holmes knows that *Mary* stole the bicycle' is true, but 'Holmes knows that Mary stole *the bicycle*' is false. The explanation is that Kh⟨that Mary stole the bicycle, that Peter stole the bicycle⟩ is true, while Kh⟨that Mary stole the bicycle, that Mary stole the wagon⟩ is false.

The fourth and final argument for q is that q allows knowledge ascriptions to score overall progress in inquiry. Since each stage of inquiry may be modeled as a multiple-choice question (§1), overall progress in inquiry may be scored as a record of which multiple-choice questions one can successfully answer. Recognition of q allows such scoring because different values for q correspond to distinct stages of inquiry. Thus if Mikey can discriminate Coke from Sprite but not from Pepsi, then Mikey's partial overall progress in determining what he is drinking may be scored via: *success*, Mikey knows that the drink is Coke rather than Sprite; and *failure*, Mikey does not know that the drink is Coke rather than Pepsi.

This concludes my characterization of contrastivism. I should note two qualifications by way of conclusion. The first qualification is that the contrastivist theses (4)–(6) are *neutral as to the analysis*

of knowledge. Indeed there is room for a 'contrastivized' version of virtually any major account of knowledge. My aim is not to adjudicate between different versions of contrastivism, but rather to compare the contrastivist framework generally to the contextualist framework, which is equally neutral (§1). So I will not discuss this issue further here.[14]

The second qualification to note is that my characterization of contrastivism is *wholly stipulative*. By 'contrastivism' I mean the conjunction of (4)–(6). My aim is not mere stipulative definition of course, but rather to compare theses (4)–(6) to their contextualist counterparts (1)–(3). To this I now turn.

3. INDEXICALITY VERSUS TERNICITY

How if at all do contextualism and contrastivism differ?[15] There is a familial resemblance between contextualism and contrastivism: (i) both treat binary utterances as having context-dependent truth-conditions, (ii) both factor alternatives into the truth-conditions, and (iii) both shield ordinary knowledge from skeptical doubt. But there are also subtle differences between the shared features of contextualism and of contrastivism: (i) indexicality and ternicity are distinct linguistic models for generating context-dependent truth-conditions, (ii) relevance and saturation are distinct linguistic mechanisms for factoring alternatives into the truth-conditions, and (iii) equivocationism and compatibilism are distinct philosophical accounts of how ordinary knowledge is shielded from skeptical doubt.

So how if at all do contextualism and contrastivism differ? Starting with the difference between (1) indexicality and (4) ternicity, these are distinct linguistic models for generating context-dependent truth-conditions, which differ as follows. By indexicality, context-dependence is generated by a semantical rule triggered by the occurrence of 'knows', which outputs different epistemic properties in different contexts. The model for this mechanism is 'I', whose occurrence triggers a semantical rule that outputs different individuals in different contexts. Whereas by ternicity, context-dependence is generated by the absence of an explicit setting for q, which is then implicitly saturated by different alternatives in different contexts. The model for this mechanism is 'prefers',

which licenses reduced expression (such as 'Ann prefers chocolate') where the covert foil takes different values in different contexts. In short, by indexicality what is variable is the relation denoted by 'knows', whereas by ternicity what is variable is not the relation denoted by 'knows' but rather the value of the relatum *q* when left implicit.[16]

I offer three arguments that ternicity is preferable to indexicality. The first argument is the argument from *non-binary stability*: ternicity better predicts the relative non-shiftiness of non-binary ascriptions. First premise: I take it as intuitively clear that non-binary ascriptions are less shifty than their binary counterparts. Compare, for instance, the contrastive utterance type (i) 'Holmes knows that Mary stole the bicycle rather than the wagon', with the binary utterance type (ii) 'Holmes knows that Mary stole the bicycle'. It is intuitively evident that (i) is less shifty than (ii). Explanation: (i) can only encode the alternative that Mary stole the wagon, while (ii) can encode such diverse alternatives as: that Mary stole the skates, that Mary merely borrowed the bicycle, that it was Peter who stole the bicycle, and/or that the whole episode was a dream, etc.[17]

Second premise: indexicality predicts that non-binary ascriptions should be as shifty as their binary counterparts. This is because, with indexicality, the shiftiness is generated by a semantical rule triggered by the occurrence of 'knows'. Since 'knows' occurs in *all* the knowledge ascriptions, they should *all* generate this shiftiness. Thus indexicality predicts, for instance, that both (i) 'Holmes knows that Mary stole the bicycle rather than the wagon' and (ii) 'Holmes knows that Mary stole the bicycle' should share the same 'knows'-induced shiftiness.

Third premise: ternicity predicts that non-binary ascriptions should not be as shifty as their binary counterparts. This is because, with ternicity, what is context-variable is not the relation denoted by 'knows' but rather the value of the relatum *q* when left implicit. Since non-binary ascriptions are relatively explicit as to the value of *q*, they should have a lesser degree of shiftiness. Thus ternicity predicts, for instance, that (i) 'Holmes knows that Mary stole the bicycle rather than the wagon' explicitly fixes *q* and so should leave

little to shift; while (ii) 'Holmes knows that Mary stole the bicycle' does not explicitly fix q and so should be shiftier.

In summary, the formal difference between indexicality and ternicity yields a predictive difference for non-binary ascriptions, where q is made explicit. Putting the preceding pieces together, the argument from non-binary stability runs as follows:

(7) Non-binary ascriptions are less shifty than their binary counterparts;
(8) Indexicality predicts that non-binary ascriptions should be as shifty as their binary counterparts, because both contain an occurrence of 'knows'; and
(9) Ternicity predicts that non-binary ascriptions should be less shifty than their binary counterparts, because the non-binaries are more explicit as to the value of q.

Thus ternicity better predicts the relative non-shiftiness of non-binary ascriptions.[18]

The second argument for the preferability of ternicity is the argument from *scoring inquiry*: ternicity better suits 'knows' to its role in keeping score of the overall progress of inquiry. First premise: one of the roles of the knowledge ascription is to keep score of the overall progress of inquiry. We have an epistemic interest in truth. Inquiry is our method for seeking truth. So we have an epistemic interest in keeping score of the overall progress in inquiry, and we use our epistemic vocabulary to serve this interest – 'knows' is an honorific for successful inquirers.[19] By way of illustration, suppose Holmes is inquiring into who stole the bicycle. He finds Mary's fingerprints in the store, which enables him to conclude the inquiry successfully ('So, it was Mary!') I am maintaining that we attribute knowledge to Holmes when he has successfully answered the question – now we may say, 'He knows who stole the bicycle'.[20]

Second premise: indexicality precludes 'knows' from keeping score of the overall progress of inquiry. This is because, with indexicality, the denotation of 'knows' is always warped to the current context. As such 'knows' cannot keep consistent score *across contexts*. But scoring inquiry requires evaluating how a subject performs through a sequence of questions, and this requires a consistent score across contexts. (Imagine trying to score a baseball

game if the denotation of 'run' changed with every inning!) Suppose that Mikey can discriminate Coke from Sprite but not from Pepsi. 'Knows' as the indexicalist conceives it cannot a consistent keep score of Mikey's overall progress, since the introduction of the Pepsi alternative warps 'knows' in such a way that Mikey's success in discriminating Coke from Sprite is lost to our epistemic vocabulary: 'Mikey knows that the drink is Coke' just comes out false in such a context.

Third premise: ternicity allows 'knows' to keep score of the overall progress of inquiry (§2). This is because, with ternicity, all that is context-variable is the value of q. Thus success and failure can be consistently logged under different q-values. If Mikey succeeds in discriminating Coke from Sprite but fails in discriminating Coke from Pepsi, then Mikey's overall scorecard reads: Mikey knows that the drink is Coke rather than Sprite, but Mikey does not know that the drink is Coke rather than Pepsi: Kmcs & ∼Kmcp. Here Mikey's partial progress towards identifying what he is drinking is on context-invariant display.

In summary, the formal difference between indexicality and ternicity yields a serviceability difference for inquiry, where a consistent score is needed. Putting the preceding pieces together, the argument from scoring inquiry runs as follows:

(10) One of the roles of 'knows' is to keep score of the overall progress of inquiry;
(11) Indexicality precludes 'knows' from scoring the overall progress of inquiry, because indexicals cannot keep a consistent score across contexts; and
(12) Ternicity allows 'knows' to score the overall progress of inquiry, because the various stages of inquiry may be consistently logged under various values of q.

Thus ternicity better suits 'knows' to its role in keeping score of the overall progress of inquiry.

The third argument for the preferability of ternicity is the argument from *precedent*: ternicity is the more precedented phenomenon with respect to words like 'knows'. First premise: it is simply implausible that 'knows' is some sort of lexical freak, whose behavior is unprecedented. To say otherwise would be to engage in special pleading for 'knows'.

Second premise: indexicality is unprecedented for 'knows'. As Jason Stanley (2000, pp. 400, 431) has observed in this regard, the only other (non-controversial) examples of indexicals are the obvious indexicals, demonstratives, and pronouns. These differ from 'knows' in lexical kind, in the blatancy of their shiftiness, and in their resistance to being bound. Thus to lump 'knows' in with these terms is to treat 'knows' as a lexical freak, a lone verb among proforms, a lone subtle shifter among blatant shifters, and a lone bindable indexical among binding-resistant indexicals.

Third premise: ternicity groups 'knows' within a class of verbs including 'prefers', 'explains',[21] and 'selects', which denote ternary relations that license reduced binary expressions such as 'Ann prefers vanilla [rather than what?]', 'Refraction explains rainbows [rather than what?]', and 'Mikey can select Coke [from what?]'. These terms feature varying degrees of blatancy: 'selects' is blatantly ternary, 'prefers' less so, and 'explains' rather subtly so. And these terms pass the same tests for covert variables (as per §2). Thus to include 'knows' in with these terms is to treat 'knows' in a precedented way, amongst lexical kin, where subtlety already features, and where binding is understood.

In summary, the formal difference between indexicality and ternicity yields a conformity difference for 'knows' within the lexicon. Putting the preceding pieces together, the argument from precedent runs as follows:

(13) 'Knows' should not be treated as a lexical freak;
(14) Indexicality treats 'knows' as a lexical freak, because 'knows' differs from the other indexicals in its lexical kind, subtlety of shiftiness, and bindability; and
(15) Ternicity treats 'knows' in a precedented way, because 'knows' is similar to other ternaries in its lexical kind, subtlety of shiftiness, and bindability.

Thus ternicity is the more precedented phenomenon with respect to words like 'knows'.[22]

The three preceding arguments for ternicity may be reframed as a comparison between 'prefers' and 'I' as models for 'knows'. As to the argument from non-binary stability (7)–(9), 'prefers' loses shiftiness in ternary utterances such as 'Ann prefers vanilla to strawberry' because the foil is fixed at strawberry, while 'I' never loses

its shiftiness no matter how one might try.[23] As to the argument from scoring inquiry (10)–(12), 'prefers' may be used to keep score of overall rankings ('Ann does not prefer vanilla to chocolate, but she does prefer vanilla to strawberry'), but 'I' cannot be used to keep track of any individual, since it simply shifts denotation to whomever happens to be speaking. And as to the argument from precedent (13)–(15), 'prefers' makes a far more plausible precedent for 'knows' than 'I' does, since 'prefers' and 'knows' are lexically kindred, partially subtle, and bindable.

The argument from non-binary stability (7)–(9) reveals that, while both contextualists and contrastivists claim to explain shifty acceptability (§1, §2), it is the contrastivist who frames the better explanation. The argument from scoring inquiry (10)–(12) reveals that, while both contextualists and contrastivists claim to connect knowledge to inquiry (§1, §2), it is the contrastivist who forges the better connection. Thus I conclude that contrastivism is based on a more suitable linguistic model: 'knows' is a ternary relation, not an indexical.

4. RELEVANCE VERSUS SATURATION

I turn now to the second main difference between contextualism and contrastivism, namely that between relevance and saturation. These are distinct linguistic mechanisms for factoring alternatives into the truth-conditions, which differ as follows. By relevance, alternatives enter into the truth-conditions via the semantical rule of relevant alternatives (as triggered by the 'knows' indexical). Whereas by saturation, alternatives enter into the truth-conditions via such mechanisms as 'rather than'-arguments, interrogatives, focusing, and free variables (as serving to saturate the q slot). In short, by relevance there exists a 'knows'-specific relevance function, whereas by saturation there is no 'knows'-specific relevance function but rather just a variety of general linguistic mechanisms for saturating an argument slot.

I offer two arguments that saturation is preferable to relevance. The first argument is the argument from *the mysteries of relevance*: saturation is a less mysterious linguistic mechanism. First premise: I take it that an adequate approach to 'knows' should

not multiply linguistic mysteries beyond necessity. The extreme of inadequacy would be an approach according to which the alternatives are generated by 'magic', or just pulled out of the theorist's hat.

Second premise: relevance is a deeply mysterious mechanism for generating alternatives. No contextualist has ever offered anything near a precise account of relevance. Indeed the only contextualist who has even made a serious attempt in this direction is Lewis (1996), who proposes seven rules. According to Lewis's rules, the following possibilities are always relevant: (i) actuality, (ii) those the subject believes or ought to believe to be actual, (iii) those the attributor is attending to, and (iv) those that directly and saliently resemble those relevant by (i)–(iii); whereas the following possibilities are defeasibly irrelevant: (v) those concerning errors in reliable processes, (vi) those concerning errors in sampling and abduction, and (vii) those conventionally ignored. But with all due respect to Lewis, this is little more than a laundry list of rules of thumb, replete with unclear principles,[24] subject to a variety of counterexamples,[25] and open to skeptical usurpation as merely pragmatical.[26] (Compare the convolutions of Lewisian relevance to the simplicity of the speaker rule for 'I'.) As Ernest Sosa remarks, pending a precise account of relevance, contextualism "will remain unacceptably occult" (1986, p. 585). The mechanism of relevance remains as mysterious as magic.

Third premise: saturation is an antecedently recognized and decently understood mechanism. That is, the various mechanisms of 'rather than'-arguments, interrogatives, focusing, and free variables are all antecedently recognized and decently understood mechanisms for generating truth-conditional contrasts. Or at least, whatever mysteries these mechanisms harbor are mysteries we must face anyway, since these mechanisms are not 'knows'-specific – there is at least no multiplication of mysteries here. Saturation is old hat.

In summary, the difference in mechanism between relevance and saturation is the difference between magic and old hat. Putting the preceding pieces together, the argument from the mysteries of relevance runs as follows:

(16) Linguistic mechanisms should not multiply mysteries;

(17) The mechanism of relevance multiplies mysteries, because such a 'knows'-specific rule is completely obscure; and
(18) The mechanism of saturation does not multiply mysteries, because the general linguistic mechanisms involved are antecedently recognized and decently understood.

Thus saturation is a less mysterious mechanism for generating alternatives.[27]

The second argument for the preferability of saturation is the argument from *contrastive explicitness*: saturation better captures the alternatives encoded in contrastive ascriptions. First premise: I take it as intuitively clear that the alternatives of contrastive ascriptions are to be read off their 'rather than'-arguments. For instance, the alternative in 'Holmes knows that Mary stole the bicycle rather than the wagon' is: that Mary stole the wagon (§3).

Second premise: relevance is blind to the content of 'rather than'-arguments. Lewis's rules (which I will continue to work with for the sake of some degree of definiteness) deliver alternatives not listed in the 'rather than'-argument if those alternatives are actual, believed, attended to, or directly resembling such; and Lewis's rules fail to deliver alternatives listed in the 'rather than'-argument if those alternatives concern errors in reliable processes, sampling errors, or what is conventionally ignored.[28] For instance, if Peter actually stole the bicycle and no one seriously attends to the possibility that Mary stole the wagon, then the alternative which Lewisian relevance delivers for 'Holmes knows that Mary stole the bicycle rather than the wagon' is: that Peter stole the bicycle.

Third premise: saturation reads the alternatives off the 'rather than'-argument (§2). Here 'prefers' continues to serve as a useful precedent: in 'Ann's prefers vanilla rather than strawberry', strawberry saturates the foil. Likewise the alternative that saturation delivers in 'Holmes knows that Mary stole the bicycle rather than the wagon' is: that Mary stole the wagon.

In summary, the difference in mechanism between relevance and saturation yields an interpretive difference with respect to contrastive ascriptions. Piecing things together, the argument from contrastive explicitness runs as follows:

(19) The alternatives in contrastive ascriptions should be read off the 'rather than'-argument;

(20) Relevance is blind to the content of the 'rather than'-argument, because relevance is keyed into independent factors such as actuality and belief; and

(21) Saturation reads the alternatives off the 'rather than'-argument, because saturation is keyed directly into such an argument.

Thus saturation better captures the alternatives encoded in contrastive ascriptions.

The thought that there must exist a 'knows'-specific relevance function is an artifact of the contextualist's underlying indexical model. That is, the thought that 'knows' must be associated with some sort of rule of relevance is an artifact of the contextualist's assimilation of 'knows' to 'I', together with the association of 'I' with the specific *speaker of the utterance* rule. If 'knows' is not an indexical, then it need not be associated with any specific semantical rule. If 'knows' is a ternary relation, then it just requires saturation of its argument places in the standard ways. Thus I conclude again that contrastivism is based on a more suitable linguistic model: 'knows' is a ternary relation, not an indexical.

5. EQUIVOCATIONISM VERSUS COMPATIBILISM

I turn now to the third and final difference between contextualism and contrastivism, namely that between (3) equivocationism and (6) compatibilism. These are distinct philosophical accounts of how dogmatic knowledge is shielded from skeptical doubt, which differ as follows. By equivocationism, there is a dogmatic binary property on which dogmatic claims like 'Moore knows that he has hands' count as true, and a skeptical binary property on which such claims count as false. The skeptic is accused of equivocating on the 'knows' indexical. Whereas by compatibilism, it is true that Moore knows that he has hands rather than stumps, and false that Moore knows that he has hands rather than vat-images of hands. The skeptic is accused of illicitly shifting the contrast variable. These differences show up in the treatment of deductive closure (which must be reconceived for the K*spq* form), and in the diagnosis of skeptical anxiety.

I offer two arguments that compatibilism is preferable to equivocationism. The first argument is *the argument from deductive modesty*: only compatibilism can reconcile deductive closure with epistemic modesty. First premise: I take it that an account of knowledge must satisfactorily reconcile (i) the transmission of knowledge via deductive inference,[29] with (ii) the modesty of human knowledge. The problem is that deductive transmission *looks* to undermine modesty, for it looks as if Moore's modest knowledge that he has hands allows him to deductively infer immodest knowledge of not being envatted, or that Moore's modest ignorance as to envattedness allows him to deductively infer immodest ignorance as to handedness.[30]

Second premise: equivocationism does not reconcile deduction with modesty. The equivocationist insists that all the epistemic properties denotable by 'knows' are closed under deduction.[31] But that means that every epistemic property is either dogmatic or skeptical – *no epistemic properties are modest*. For the equivocationist, dogmatism reigns in the courtroom, so that there one can truly count as knowing that one is not a brain-in-a-vat (full stop) – but surely one can never know so much! And for the equivocationist, skepticism reigns in the classroom, so that there one can truly count as not even knowing that one has hands (or anything else about the external world) – but surely one can never know so little! Thus the equivocationist is forced to swing from a manic dogmatism to a depressive skepticism.

Third premise: compatibilism does reconcile deduction with modesty. For the contrastivist, modest knowledge is knowledge involving mundane contrasts, such as: Moore knows that he has hands rather than stumps (Kmhs); while modest ignorance is ignorance involving skeptical contrasts, such as: Moore does not know that he has hands rather than vat-images of hands (∼Kmhv).[32] Deductive transmission works as follows. Think of K*spq* as *knowledge relative to a question*, or, more picturesquely, *local knowledge within a partitioned region of logical space*:

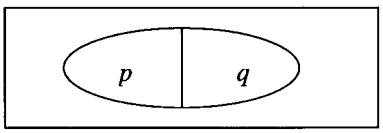

The following valid knowledge-transmitting inferences can be adduced from this picture of local knowledge: (i) Expand p: if $p_1 \rightarrow p_2$ then K$sp_1q \rightarrow$ Ksp_2q, and (ii) Contract q: if $q_2 \rightarrow q_1$ then K$spq_1 \rightarrow$ Kspq_2; whereas the following transmission principles will be invalid: (iii) *Contract p: *if $p_2 \rightarrow p_1$ then K$sp_1q \rightarrow$ Ksp_2q, (iv) *Expand q: *if $q_1 \rightarrow q_2$ then K$spq_1 \rightarrow$ Kspq_2, (v) *Replace p: *K$sp_1q \rightarrow$ Ksp_2q, and (vi) *Replace q: *K$spq_1 \rightarrow$ Kspq_2.[34] The validity of (i) and (ii) shows that the contrastivist respects deductive closure *suitably understood* (in particular (i) shows how mathematical proof still counts as knowledge-transmitting). But the invalidity of (vi) shows that Moore's knowledge that he has hands rather than stumps is compatible with his ignorance as to whether he has hands or vat-images of hands: Kmhs does not entail Kmhv.[34] Intuitively, 'hands or stumps?' and 'hands or vat-images?' are distinct multiple-choice questions; intuitively, knowledge relative to the former does not entail knowledge relative to the latter; intuitively, the former is an easy question answerable at a glance, while the latter is a trick question that cannot be answered at all. Alternative-relative knowledge is modest knowledge.

In summary, the philosophical difference between equivocationism and compatibilism is the difference between a handful of immodest properties, and one modest property. Piecing things together, the argument from deductive modesty runs as follows:

(22) An account of knowledge should reconcile deduction with modesty;
(23) Equivocationism fails to reconcile deduction with modesty, because it allows deduction to render every epistemic property immodestly dogmatic or skeptical; and
(24) Compatibilism reconciles deduction with modesty, because it relativizes knowledge in such a way that deduction preserves modesty.

Thus compatibilism better reconciles deductive closure with epistemic modesty.

The second argument for the preferability of compatibilism is *the argument from skeptical anxiety*: only compatibilism can explain why skeptical arguments feel so nightmarish. First premise: I take it that an account of knowledge must not just resolve the skeptical paradox, but it must also explain why the skeptical arguments feel so

nightmarish (DeRose, 1995, §1). Why do brain-in-a-vat nightmares panic us into fearing ignorance?

As the second premise, equivocationism cannot explain why skeptical arguments feel so nightmarish. As Stephen Schiffer (1996) points out, equivocationism diagnoses the error of skeptical arguments as an equivocation on an indexical, and so equivocationism requires the claim that we are panicked by such an equivocation. Yet it does not seem that equivocations on indexicals draw competent speakers into panic. If Ann says 'I am at the Haymarket café' and Ben says, 'I am not in the Haymarket café', then I trust that no competent speaker will announce a paradox, much less blink. In general, if one considers all the non-contentious examples of real indexicals, none seem likely to confuse a soul. So Schiffer concludes, "If that's the solution, what the hell was the *problem*?" (1996, p. 329).

As the third and final premise, the compatibilist can provide a decent explanation for why skeptical arguments feel so nightmarish. The compatibilist diagnoses the error of skeptical arguments as ignoring the covert contrast variable. And it does seem that covert variables can induce confusion among competent speakers. For instance, if Ann tells Ben 'I prefer chocolate' at one store (relatively blatant ternicity), it is relatively unlikely but not impossible that Ben will get confused and think the same must apply at a second store. While if Professor tells Gilligan 'Refraction explains rainbows' (subtle ternicity) it is relatively likely that Gilligan will get confused. The subtler the presence of the variable, the likelier competent speakers can be led into paradox. Skeptical arguments feel so nightmarish because the covert contrast variable in 'knows' is so easily missed. The wile of the skeptic is to shift the contrast.[35]

In summary, the philosophical difference between equivocationism and compatibilism is the difference between an indecent and a decent explanation of skeptical anxiety. Piecing things together, the argument from skeptical anxiety runs as follows:

(25) An account of knowledge must explain why skeptical arguments seem so nightmarish;

(26) Equivocationism cannot explain why skeptical arguments seem so nightmarish, because errors with indexicals do not occur; and

(27) Compatibilism can explain why skeptical arguments seem so nightmarish, because errors with suppressed argument places are common.

Thus compatibilism better explains why skeptical arguments seem so compelling.[36]

The thought that skeptical arguments are equivocal is another artifact of the contextualist's underlying indexical model. If 'knows' is a ternary relation, then the dogmatist and the skeptic need not equivocate at all. Rather the dogmatist and skeptic may be portrayed as pointing out compatible truths, where the appearance of incompatibility is due to the subtlety of the covert contrast variable. Thus I conclude once again that contrastivism is based on a more suitable linguistic model: 'knows' is a ternary relation, not an indexical.

Contextualism and contrastivism are sibling theories. On both theories, binary knowledge ascriptions have context-dependent truth-conditions, alternatives are factored into the truth-conditions, and dogmatic knowledge is shielded from skeptical doubt thereby. But I think that contextualism obscures these insights, and faces serious objections, by deploying an unsuitable indexical model. Contrastivism provides a more suitable ternary model, which preserves the core contextualist insights while resolving the main objections to contextualism. I conclude that contextualist insights are best understood by contrastivist theory.[37]

NOTES

[1] As Lewis acknowledges in his concluding paragraph (1996, p. 445).

[2] Some contextualists invoke *standards* (strength of epistemic position) instead of *alternatives*. Every standard corresponds to a set of alternatives: a given standard *s* determines a distance in logical space *d*, which determines a sphere of worlds (DeRose, 1995, pp. 34–35; Heller, 1999, p. 116). But not every set of alternatives corresponds to a standard: if a set of alternatives does not comprise a *sphere* in logical space, no standard will correspond to it. Thus alternatives are preferable to standards on grounds of generality (see Schaffer, forthcoming b for further discussion). Indeed, as will emerge below, the invocation of standards would preclude 'knows' from serving the purposes of inquiry, since the alternatives under inquiry need not comprise a sphere.

[3] For both the dogmatist and the skeptic, shifty acceptability should be explained pragmatically. For instance, the skeptic can invoke the pragmatics of *hyperbole* to

explain the shifty assertibility conditions for knowledge claims (Schaffer, forthcoming a). I suspect the dogmatist will have a much harder time here, but this raises issues I cannot discuss here.

[4] As John Dewey says "Inquiry and questioning, up to a certain point, are synonymous terms" (1938, p. 105). Jaakko Hintikka (1981) has developed a formally explicit interrogative model of inquiry, on which inquiry is a cooperative game played between Questioner and Answerer, represented by a sequence of question-and-answer pairs $\langle\langle Q_1, A_1\rangle, \langle Q_2, A_2\rangle, \ldots, \langle Q_n, A_n\rangle\rangle$.

[5] The association of questions with their alternative possible answers is known as *Hamblin's dictum* (after C.I. Hamblin, 1958). Hamblin's dictum connects the semantics of questions to answers in a way that explains a range of inferences, such as from 'A: Who stole the bicycle? B: Mary stole the bicycle' to 'B answered A's question', while generating straightforward treatments of such notions as *presupposition*. In this tradition, Nuel Belnap and Thomas Steel (1976) treat questions via the schema $?\rho\sigma$, where 'ρ' signifies the lexical request and 'σ' the lexical subject. The lexical subject presents a set of alternatives for selection in accord with the request. And James Higginbotham extends the notion of a set of alternatives into a partitioned region of logical space: "An *abstract question* [is] a nonempty *partition* Π of the possible states of nature into *cells* P_i for $i \in I$, having the property that no more than one cell corresponds to the true state of nature (i.e., the cells are mutually exclusive)" (1993, p. 196; see also Higginbotham, 1996). Note (to continue the discussion from note 2) that the queried alternatives need not comprise a sphere.

[6] There is virtually a one-one function from major analyses of knowledge to contextualists, by the rule 'has advocated a contextualized version'. This function maps the JTB analysis onto Cohen, the tracking analysis onto DeRose, the elimination analysis onto Lewis, and the reliability analysis onto Heller.

[7] David Annas (1978) and Michael Williams (1996), for instance, offer 'contextualist' theories that do not fit my definition. Roughly, both are social-foundationalists about the regress of reasons, maintaining that the regress can halt at propositions unchallenged by one's peers. Such a view about justification is compatible with virtually any view about the linguistic behavior of 'knows' (DeRose, 1995).

[8] Strictly speaking, a temporal relatum *t* is needed as well, since one may learn or forget. So the full 'binary' structure is really K*spt*, and the full 'ternary' structure is really K*spqt*. In the main text I suppress *t* for convenience.

[9] The binary theorist might rejoin that 'knows' is *ambiguous* (/variably polyadic) between a binary and contrastive relation, or that the additional 'rather than'-clause is merely *adjunctive*. The ambiguity option strikes me as very implausible, and I will not pursue it further here – ambiguities are not to be posited so lightly! The adjunctive reading is implausible as well. If the 'rather than'-clause were merely adjunctive in contrastive knowledge ascriptions, then (i) iterated 'rather than'-clauses would be licensed, (ii) the 'rather than'-clause would serve as a state modifier, and (iii) order would not matter were other adjunct phrases added. But (i) iteration is not licensed: ?'Holmes knows that Mary stole the bicycle rather

than the wagon rather than Peter'; (ii) state modification is obscure: it is clear what it is to be in the knowledge state *by looking*, but obscure what it could mean to be in the knowledge state *rather than Peter*; and (iii) order matters: 'Holmes knows that Mary stole the bicycle rather than the wagon by looking' is fine, but not: ?'Holmes knows that Mary stole the bicycle by looking rather than that Mary stole the wagon'. Further, the adjunctive reading would still not cover interrogative ascriptions.

[10] The binding test is due to Barbara Partee (1989), and is used extensively by Jason Stanley (2000). Caveats: (i) intuitions about bound readings can be very delicate, (ii) there may be a pragmatic explanation for some bound readings (Bach, 2000). Stanley (personal communication) grants that there is a bound reading here but suggests an alternative syntactic explanation, namely that what is being bound is the domain variable for the quantifier in 'the drink', not any contrast variable. But one still gets bound readings if one dubs the drink 'Sam' and says, "On every test, Mikey knows that Sam is Coke'. Thus the binding cannot be explained away in terms of a domain variable in 'the drink'." (Less artificially, one might start with a case where Mikey is being tested on height discrimination. He is presented with Tim the six-footer. He is first asked 'Is Tim six-feet or seven-feet?', and second asked 'Is Tim six-feet or six-foot-one?' Now the same bound reading arises for: 'On every question, Mikey knows that Tim is six-feet'.)

[11] Here I am assuming a *syntactic-copying* view of VP ellipsis, such as defended by Robert May (1985; see also Fiengo and May, 1994), on which the syntactic structure and lexical content of an antecedent gets copied into the ellipsis site. Though what I say is compatible with a pluralistic view on which different sorts of ellipsis receive different treatments, as per Shalom Lappin (1996). Caveat: it may be difficult to discern whether a variable is really being copied, or a presupposition is merely being carried.

[12] Focus differences are not always semantically effective. For instance, it seems that there is no truth-conditional difference between '*Mary* stole the bicycle' and 'Mary stole *the bicycle*'. Yet focus differences do yield truth-conditional differences in cases such as those involving adverbs of quantification, conditionals, modals, generics, superlatives, and factives like 'it is odd that' and 'I found out that' (Rooth, 1996, §1; see also Rooth, 1992; Hajicova, Partee and Sgall, 1998, §2.2.2). What unifies this list? That is, which focus differences generate truth-conditional differences? Partee (1991) hypothesizes that focal sensitivity arises from operators that apply to *multiple arguments* in an asymmetric manner. Thus focal sensitivity reveals a multiplicity of arguments where the surface displays just one.

[13] Thus when Moore mounts the podium and declares, "I know that I have hands" (a binary ascription in discourse-initial position), we the audience must pragmatically construct a contrast. We do so with charity. I suggest that this is why Moorean declarations of knowledge seem undeniable yet empty.

[14] My own view is that K*spq* iff (i) *p* is true, (ii) *s* has conclusive evidence that *p* rather than *q*, and (iii) *s* is certain that *p* rather than *q* on the basis of (ii) (see Schaffer *ms*. §4 for further discussion).

[15] Lewis (personal communication) maintains that contextualism and contrastivism differ not at all, quipping: "The only thing we disagree about is whether we disagree." Neta (Pacific APA, 2001) claims that contrastivism is a species of contextualism, by describing my contrastivist account of knowledge (*ms*.) as: ". . . one of the most explicit, comprehensive, and thoroughly defended *contextualist* theories of knowledge to date." What follows (§s 3–5) may be thought of as an extended reply to Lewis and Neta.

[16] The difference between indexicality and ternicity can be formally represented as the difference between an indexed binary relation $K_i sp$ (which may be given a natural language paraphrase as: *s* bears *that property* to *p*) and an unindexed ternary relation K*spq* (*s* knows that *p* rather than *q*). Given that the difference between indexicality and ternicity emerges at the level of formal semantics, many relevant alternatives theorists (Austin, Fred Dretske, Alvin Goldman, G.C. Stine) simply have not expressed their views in ways that decide. Lewis, Cohen, and DeRose are explicitly indexicalist. Indeed Cohen has recently become aware of this difference, asking:

> How from the point of view of formal semantics should we think of this context-sensitivity of knowledge ascriptions? We could think of it as a kind of indexicality . . . But we could instead . . . think of knowledge as a three-place relation between a person, a proposition, and a standard. (1999, p. 61)

Though he immediately dismisses this question as "irrelevant to the epistemological issues" (1999, p. 61).

[17] The contrastive ascription still retains some context-dependence, in that there is context-dependence concerning which set of worlds is denoted by a that-clause. Thus the alternative that Mary stole the wagon may or may not (depending on context) include worlds in which Holmes is a brain-in-a-vat veridically hallucinating Mary's thieving. But this remaining bit of shiftiness is still considerably less than that permitted the binary utterance: in no context may the alternative that Mary stole the wagon include worlds in which what she stole was the roller skates, or in which it was Peter who did the thieving, etc.

[18] One possible contextualist rejoinder is to say that non-binary ascriptions are less shifty in virtue of doing more to fix the relevance function. The idea is to explain the relative non-shiftiness of 'Holmes knows that Mary stole the bicycle rather than the wagon' by postulating that the 'rather than'-argument manages to fix relevance. I will rebut this rejoinder in §4: the relevant-alternatives rule is as blind to the contents of the 'rather than'-argument as the speaker-of-the-utterance rule for 'I' is to the rest of the utterance.

[19] Here I follow Christopher Hookway: "The central focus of epistemic evaluation is . . . the activity of inquiry . . . When we conduct an inquiry, or deliberate on some matter, we attempt to formulate questions and answer them correctly" (1996, p. 7).

[20] The argument of the main text is consistent with a plurality of roles for 'knows'. For instance, our epistemic vocabulary may also serve to explain action, and to stand as the norm for assertion, etc. (Williamson, 2000). That said, I suspect

that these other roles may be derivative from the role of 'knows' in inquiry. Actions are only explicable in light of which options the subject considers, where this may be modeled in terms of a multiple-choice inquiry concerning 'What should I do?' And assertions are only appropriate in contexts in which they advance the conversation, where this may be modeled in terms of winnowing the alternatives in the context set (Stalnaker, 1998, p. 99).

[21] The contrastive treatment of 'explains' is due to Bas van Fraassen (1980).

[22] One possible contextualist rejoinder is to adduce more controversial precedents such as Mark Richard's (1990) account of 'believes', on which the truth of belief ascriptions varies contextually with a 'faithful representation' parameter. If a critical mass of defensible precedents could be adduced I would accept this rejoinder. Though I would first ask whether the alleged precedents aren't themselves better understood via covert argument places!

[23] Attempts to fix 'I' on the speaker produce such redundancies as 'I, Jonathan . . .', while attempts to fix 'I' on anyone other than the speaker produce such catastrophes as 'I, Napoleon . . .'. Other indexicals allow a bit more flexibility: I can specify the boundaries of 'here' as 'Here in Amherst . . .' or as 'Here in Massachusetts . . .', but still the center cannot shift. Attempts to shift the center produce such absurdities as 'Here on Mars . . .'

[24] Some unclarities: What is it for two possibilities to 'saliently resemble'? What is it for a possibility to be one '*s* ought to believe'? What is it for a possibility to be 'conventionally ignored'?

[25] See Jonathan Vogel (1999) for some rather convincing counterexamples here.

[26] The skeptical usurper takes all knowledge ascriptions to be false, and reinterprets the Lewis rules as constraints on the *assertibility* of such falsehoods, on the model of assertible *hyperbole*. Hyperboles generally are assertible in contexts in which their true implications are presumptively informative and their false implications are presumptively ignorable (Sperber and Wilson, 1986; similar conclusions will follow from a neo-Gricean treatment of hyperbole in terms of flouting the maxim of Quality, as per Stephen Levinson (1983)). For instance, 'the airplane is a mile long' is assertible in contexts in which (i) the true implications, such as that the airplane is greater than 100 feet long, are presumptively informative, and (ii) the false implications, such as that the airplane is greater than 1000 feet long, are presumptively ignorable. From this perspective the Lewis rules just look like decent generalizations about typical audiences: one's audience will presumptively take seriously the actual, believed, and attended possibilities, as well as those that saliently resemble them; and one's audience will presumptively ignore the reliable-process breakdown, methodological error, and conventionally ignored possibilities. Thus relevance invites skeptical plucking (Schaffer, forthcoming a).

[27] The contextualist might rejoin by withdrawing any 'knows'-specific account of relevance and instead invoking general linguistic mechanisms only. Indeed Lewis starts off by sounding as if he is invoking general linguistic mechanisms only, both in his 1979 where he assimilates 'knows' to *accommodation*, and in his 1996 where he assimilates 'knows' to *quantifier domain restriction* (though in

his 1996 the specifics of 'knows' soon take center-stage). And Neta (personal communication) disavows any 'knows'-specific view of relevance, preferring instead to invoke the speaker's intentions. But this rejoinder faces three problems. The first problem is that it is inconsistent with indexicality. That is, indexicality (in the intended sense) is the view that 'knows' has a context-invariant character, which is a specific semantical rule that determines a function from context to content. – No specific semantical rule, no indexical. The second problem with this rejoinder is that it vitiates most applications of the theory. For instance, if the contextualist deploys anything like Stalnaker's notion of a context set, then the contextualist must forego such Lewisian rules as Actuality, Belief, and Resemblance, since the context set need not contain actuality, need not correspond to anyone's beliefs, and is not closed under resemblance (Stalnaker, 1998, §2). As such, contextualism would no longer underwrite, e.g., Lewis's solutions to skepticism, Gettier cases, and the lottery paradox, since these require Actuality, Belief, and Resemblance. The third problem with this rejoinder is that it opens the door even wider to skeptical usurpation, since the skeptical usurper can already generate pragmatic counterparts of Lewis's rules from the antecedently recognized pragmatic of hyperbole, thus completely obviating the need for any further shiftiness in the semantics.

[28] The mere mention of the alternatives in the 'rather than'-argument is not sufficient to defeat presumptive irrelevance, since mentioning is not sufficient for attending (Lewis, 1996, pp. 435–436). Nor can the gap between mentioning and attending be closed without generating such implausible results as that the skeptical defense attorney should win every time: 'Your honor, that witness knows nothing!'

[29] Knowledge is transmitted by deduction. That is, successful deduction from previous knowledge allows for further knowledge. How could it not, given that our epistemic interest is truth, and deduction is truth preserving? How could it not, given that mathematical proof is deductive, and mathematical proof transmits knowledge? (Williamson, 2000).

[30] Some variations: one's modest knowledge that the beast is a zebra allows one to deductively infer immodest knowledge that the beast is not a painted mule, one's modest knowledge that one's car is parked on Elm allows one to deductively infer immodest knowledge that it has not been stolen away, and one's modest knowledge that one will visit Montreal next year allows one to deductively infer immodest knowledge that one won't die this year, etc. (see Hawthorne *ms*. for a number of further examples).

[31] Dretske (1970) takes contextualism to license a denial of closure because of the shiftiness of 'knows'. The point that each epistemic property should still satisfy closure is due to G.C. Stine (1976), and has been accepted by virtually all subsequent contextualists (see especially Cohen, 1988, §3; DeRose, 1995, §10).

[32] Clarification: the propositions h and s are intended to be restricted here, not just to any old world in which Moore has hands/stumps, but to those in which he is awake and perceiving accurately, etc. That is, there is a subclass of Moore-has-hands worlds in which Moore's handedness is apparent, and a subclass of

Moore-has-stumps worlds in which Moore's stumpedness is apparent, and these are the worlds I intend h and s to denote. It is this subclass of stump-worlds that Moore's experience of apparent hands allows him to eliminate, this that his modest knowledge concerns, and this that I intend s to denote. Moreover, h, s, and v are intended to constitute a restricted partition on logical space, one naturally corresponding to the question 'h, s, or v?' While Moore could be a brain-with-hands-stapled-to-his-medulla in a vat, I do not intend this possibility to be in h, s, or v.

[33] See Schaffer (*ms*., §5) for further discussion. The basic idea is that the valid principles preserve the elimination of all-but-*p*.

[34] Point of clarification: Expand *p* does allow Moore to know that he is not a brain-in-a-vat rather than someone with stumps. But this is actually a very modest bit of knowledge, secured not by Moore's possessing any sort of positive evidence for not being a brain-in-a-vat, but rather by Moore's possessing conclusive negative evidence against having stumps (reminder: the stumps-possibility is restricted to worlds in which stumpedness would be apparent). The associated question 'Are you (i) not a brain-in-a-vat, or (ii) a creature with stumps?' is an easy question since option (ii) is so easily eliminated. And the invalidity of Expand *q* precludes this modest bit of knowledge from licensing the immodest conclusion that Moore knows that he is not a brain-in-a-vat rather than a brain-in-a-vat.

[35] Schiffer himself considers both the indexical and the ternary theories (as well as a vagueness-based theory, all of which he lumps under 'contextualism'). But when he comes to objecting to the diagnosis of the error of skepticism, he only considers the indexical model, and does not consider whether the diagnosis of error might be more plausible given ternicity (no doubt the fact that Schiffer does not investigate ternicity further is a result of his having lumped indexicality and ternicity together under 'contextualism').

[36] Further application: Hawthorne (*ms.*) has raised the objection that (i) knowledge ascriptions appear amenable to disquotation, while (ii) indexicals are not amenable to disquotation. The same error theory that the contrastivist uses to explain away skeptical anxiety also explains away the appearance of disquotability. Binary utterances appear disquotable only because the suppressed argument place is relatively subtle. In general, contrastivism predicts that 'Moore knows that he has hands' should have a similar disquotational profile to 'Ann prefers vanilla': neither is disquotable *salva veritate* (unless the contrasts are fixed), but both might appear disquotable to the extent that the suppressed contrast is subtle.

[37] Thanks to Kent Bach, Stew Cohen, Fred Dretske, Ed Gettier, John Hawthorne, Mark Heller, David Hunter, Ernie LePore, David Lewis, Ram Neta, Walter Sinnot-Armstrong, Rob Stainton, Jason Stanley, and especially to Robert Stalnaker.

REFERENCES

Annas, D. (1978): 'A Contextualist Theory of Epistemic Justification', *American Philosophical Quarterly* 15, 213–219.

Austin, J.L. (1946): 'Other Minds', *Proceedings of the Aristotelian Society* 20(suppl.), 149–187.
Bach, K. (1994): 'Conversational Impliciture', in *Mind and Language* 9, 124–162.
Bach, K. (2000): 'Quantification, Qualification and Context: A Reply to Stanley and Szabó', *Mind and Language* 15, 262–283.
Belnap, N. and Steel, T. (1976): *The Logic of Questions and Answers*, New Haven: Yale University Press.
Chomsky, N. (1981): *Lectures on Government and Binding*, Dordrecht: Foris.
Cohen, S. (1988): 'How to be a Fallibilist', in *Philosophical Perspectives* 2, 91–123.
Cohen, S. (1999): 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* 13, 57–89.
DeRose, K. (1992): 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* 52, 913–929.
DeRose, K. (1995): 'Solving the Skeptical Problem', *The Philosophical Review* 104, 1–52.
DeRose, K. (1999): 'Contextualism: An Explanation and Defense', in Greco and Sosa (eds.), *The Blackwell Guide to Epistemology* (pp. 187–205), Oxford: Basil Blackwell.
Dewey, J. (1938): *Logic: The Theory of Inquiry*, New York: Holt, Rinehart, and Winston.
Dretske, F. (1970): 'Epistemic Operators', *The Journal of Philosophy* 67, 1007–1023.
Dretske, F. (1972): 'Contrastive Statements', *Philosophical Review* 81, 411–437.
Dretske, F. (1981): 'The Pragmatic Dimension of Knowledge', *Philosophical Studies* 40, 363–378.
Fiengo, R. and May, R. (1994): *Indices and Identity*, Cambridge, MA: MIT Press.
Goldman, A. (1976): 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73, 771–791.
Hajicova, E., Partee, B. and Sgall, P. (1998): *Topic-Focus Articulation, Tripartite Structures, and Semantic Content*, Dordrecht: Kluwer.
Hale, K. and Keyser, S.J. (1993): 'On Argument Structure and the Lexical Expression of Syntactic Relations', in Hale and Keyser (eds.), *The View from Building 20: Essays in Honor of Sylvain Bromberger* (pp. 53–109), Cambridge, MA: MIT Press.
Hamblin, C.L. (1958): 'Questions', *The Australasian Journal of Philosophy* 36, 159–168.
Hawthorne, J. (2003): 'Knowledge and Lotteries', Oxford: Oxford University Press..
Heller, M. (1999): 'Contextualism and Anti-Luck Epistemology', *Philosophical Perspectives* 13, 115–129.
Higginbotham, J. (1993): 'Interrogatives', in Hale and Keyser (eds.), *The View from Building 20: Essays in Honor of Sylvain Bromberger* (pp. 195–228), Cambridge, MA: MIT Press.

Higginbotham, J. (1996): 'The Semantics of Questions', in Lappin (ed.), *The Handbook of Contemporary Semantic Theory* (pp. 361–383), Oxford: Basil Blackwell.
Hintikka, J. (1981): 'On the Logic of an Interrogative Model of Scientific Inquiry', *Synthese* 47, 69–84.
Hookway, C. (1996): 'Questions of Context', *Proceedings of the Aristotelian Society* 96, 1–16.
Kaplan, D. (1977): 'Demonstratives', in Almog, Perry and Wettstein (eds.), *Themes from Kaplan* (pp. 481–563), Oxford: Oxford University Press, 1989.
Lappin, S. (1996): 'The Interpretation of Ellipsis', in Lappin (ed.), *The Handbook of Contemporary Semantic Theory* (pp. 145–175), Oxford: Basil Blackwell.
Levinson, S. (1983): *Pragmatics*, Cambridge: Cambridge University Press.
Lewis, D. (1979): 'Scorekeeping in a Language Game', *Journal of Philosophical Logic* 8, 339–359.
Lewis, D. (1996): 'Elusive Knowledge', *The Australasian Journal of Philosophy* 74, 549–567.
May, R. (1985): *Logical Form: Its Structure and Derivation*, Cambridge, MA: MIT Press.
Neta, R. (forthcoming a): '*S* knows that *p*', *Nous*.
Neta, R. (forthcoming b): 'Skepticism, Contextualism, and Semantic Self-Knowledge', *Philosophy and Phenomenological Research*.
Partee, B. (1989): 'Binding Implicit Variables in Quantified Contexts', in *Proceedings of the Chicago Linguistics Society*, vol. 25 (pp. 342–365), Chicago: University of Chicago Press.
Partee, B. (1991): 'Topic, Focus, and Quantification', in Moore and Wyner (eds.), *Proceedings from SALT I* (pp. 257–280), Ithaca: Cornell University Press.
Richard, M. (1990): *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them*, Cambridge: Cambridge University Press.
Rooth, M. (1992): 'A Theory of Focus Interpretation', *Natural Language Semantics* 1, 75–116.
Rooth, M. (1996): 'Focus', in Lappin (ed.), *The Handbook of Contemporary Semantic Theory* (pp. 271–297), Oxford: Basil Blackwell.
Schaffer, J. (forthcoming a): 'Skepticism, Contextualism, and Discrimination', *Philosophy and Phenomenological Research*.
Schaffer, J. (forthcoming b): 'What Shifts? Thresholds, Standards, or Alternatives?', in Preyer and Peter (eds.), *Contextualism in Philosophy*, Oxford: Oxford University Press.
Schaffer, J. (*ms.*): 'Contrastive Knowledge'.
Schiffer, S. (1996): 'Contextualist Solutions to Scepticism', *Proceedings of the Aristotelian Society* 96, 317–333.
Sinnot-Armstrong, W. (*ms.*): 'Classy Pyrrhonism'.
Sosa, E. (1986): 'On Knowledge and Context', *The Journal of Philosophy* 83, 584–585.
Sperber, D. and Wilson, D. (1986): *Relevance: Communication and Cognition*, Cambridge, MA: Harvard University Press.

Stalnaker, R. (1998): 'On the Representation of Context', *Context and Content* (1999) (pp. 96–113), Oxford: Oxford University Press.
Stalnaker, R. (this volume): 'Comments on "From Contextualism to Contrastivism in Epistemology" '.
Stanley, J. (2000): 'Context and Logical Form', *Linguistics and Philosophy* 23, 391–434.
Stine, G.C. (1976): 'Skepticism, Relevant Alternatives, and Deductive Closure', *Philosophical Studies* 29, 249–261.
Van Fraassen, B. (1980): *The Scientific Image*, Oxford: Oxford University Press.
Vogel, J. (1999): 'The New Relevant Alternatives Theory', *Philosophical Perspectives* 13, 155–180.
Williams, M. (1996): *Unnatural Doubts*, Princeton: Princeton University Press.
Williamson, T. (2000): *Knowledge and its Limits*, Oxford: Oxford University Press.

University of Massachusetts-Amherst
157 Fayette Street
Watertown, MA 02472
USA
E-mail: schaffer@philos.umass.edu