Does *knowledge* function like a quantifier?
A critique of Stanley

Giovanni Mion

Abstract: In "Elusive Knowledge" (1996), David Lewis advocates epistemic contextualism on the basis of an analysis of the nature of knowledge. For Lewis, the context-sensitivity of knowledge depends on the fact that 'knowledge that p' implies the elimination of all the possibilities in which '~p'. But since all is context-sensitive, knowledge is also context-sensitive. In contrast to Lewis, in Knowledge and Practical Interests (2005), Jason Stanley argues that since all context-sensitive expressions can have different interpretations within the same discourse, contextualists cannot consistently embrace the following two claims: (i) knowledge functions like a quantifier and (ii) distinct occurrences of knowledge within the same discourse must be associated with the same standard. In response to Stanley, I argue that (i) and (ii) are both true. More specifically, I argue that with the help of global domains, we can overcome Stanley's objections to Lewis and, accordingly, provide the linguistic basis that epistemic contextualism needs.

Keywords: epistemic contextualism; global domains; David Lewis; Jason Stanley; Christopher Gauker

In order to substantiate the claim that *knowledge* is contextual, contextualists invoke different classes of context-sensitive expressions. Notoriously,

context-sensitive expressions come in many distinct classes. They include (*a*) 'core' indexicals such as 'I', 'here', and 'now', (*b*) demonstratives such as 'this' and 'that', (*c*) gradable adjectives such as 'tall', 'flat', and 'large', (*d*) determiners such as 'many',[1] (*e*) modal auxiliaries such as 'can', (*f*) relational expressions such as 'local' and 'enemy', (*g*) noun phrases such as 'every student' and 'most professors', and (*h*) adverbial quantifiers such as 'usually' and 'always' (Stanley 2005: 32).

Given this large collection of context-sensitive expressions, the contextualists' hope to find the right class of expressions for *knowledge* seems to be justified. However, in order to validate the context-sensitivity of *knowledge*, contextualists need much more than an analogy between two linguistic expressions.

---

[1]   'Many' is a quantifier. 'The' *is* a determiner.

For example, in "Contextualism and Knowledge Attributions" (1992), Keith DeRose first compares 'knowledge' to 'this': a demonstrative (1992: 920-921), and then to 'here': a core indexical (925). But, since there is no *analytic* connection between 'here' and 'knowledge', or 'this' and 'knowledge', DeRose's claims amount only to mere analogies between linguist expressions. In fact, DeRose's ideas that knowledge varies according to the strength of epistemic position (DeRose 1995) and that the context-sensitivity of *knowledge* resemblances the context-sensitivity of demonstratives are wholly unrelated. As Jonathan Shaffer points out,

no linguistically general parameter has been identified [in connection with our *strength of epistemic position*],[2] much less one associated with expressions of which 'know' is an instance. There seems to be no precedent for this form of parameter in the language. It seems a pure invention. (Shaffer 2005: 123)

In short, on the basis of DeRose's analysis of knowledge, we cannot conclude that *knowledge* is something like an indexical or a demonstrative, and therefore a context-sensitive expression.

In contrast to DeRose's, Lewis and Cohen's semantic claims are stronger because they do not derive exclusively from contentious linguistic analogies, but from a prior conceptual analysis of what knowledge entails. According to Cohen, since knowledge implies justification and justification comes in degrees, and is thus context-sensitive, knowledge also comes in degree, and is thus also context-sensitive (see Cohen 1999: 60). In other words, according to Cohen, since knowledge relations are ordered according to justificatory strength, as tallness is ordered according to a scale, *knowledge* belongs to the class of context-sensitive expressions commonly labeled 'gradable adjectives'.

Subsuming *knowledge* under a more general class of context-sensitive expressions is precisely what the contextualist needs in order to substantiate his or her semantic thesis about knowledge. Nevertheless, Cohen's argument must be either invalid or unsound because as Stanley persuasively argues in *Knowledge and Practical Interests* (2005), any attempt to treat *knowledge* as a gradable expression fails.[3] But, while Stanley's arguments against Cohen are persuasive, Stanley's arguments against Lewis's idea that *knowledge* functions like a quantifier are not.

According to Lewis, the context-sensitivity of knowledge depends on the fact that 'knowledge that *p*' implies the elimination of all the possibilities in

---

[2]   My addition.
[3]   In particular, while gradable expressions allows for modifiers (for example: 'very', 'really', 'very much', 'very well') and for comparatives structures (for example: 'flatter than', 'taller than', 'smaller than'), 'knowledge' does not.

which '~*p*'. But since *all* is context-sensitive, *knowledge* is also context-sensitive:

> What does it mean to say that every possibility in which not-*P* is eliminated? An idiom of quantification, like 'every', is normally restricted to some limited domain. If I say that every glass is empty, so it's time for another round, doubtless I and my audience are ignoring most of the glasses there are in the whole wide world throughout all of time. They are outside the domain. They are irrelevant to the truth of what was said.
>
> Likewise, if I say that every uneliminated possibility is one in which *P*, or words to that effect, I am doubtless ignoring some of all the uneliminated alternative possibilities that there are. They are outside the domain, they are irrelevant to the truth of what was said. (Lewis 1996: 553)

In other words, for Lewis, since *knowledge* involves quantification over possibilities and in the natural language quantification is generally restricted, *knowledge* belongs to the class of context-sensitive expressions generally labeled 'quantifiers'. But according to Stanley, if *knowledge* were context-sensitive in the way envisioned by Lewis, then we should be able to associate different occurrences of 'know' within the same discourse with different sets of possibilities because:

> it is a well-established fact that different occurrences of the same quantified expression within a discourse can be associated with different domains. (Stanley 2005: 61)

Nonetheless, according to Stanley, in order to disallow the possibility for a discourse such as the following (see DeRose 1992: 924), "contextualists have exploited the view that distinct occurrences of 'know' within a discourse must be associated with the same standard" (Stanley 2005: 63):

> A: Is it a zebra?
> B: Yes, it is a zebra.
> A: But can you rule out its being merely a cleverly painted mule?
> B: No, I can't.
> A: So you admit you didn't know it was a zebra?
> B: No, I did know then that it was a zebra. But after your question, I no longer know.

How shall a contextualist respond to this dialogue? According to DeRose,

> The objection […] is based upon a mistake. The contextualist believes that certain aspects of the context of an attribution or denial of knowledge attribution affects its content […] If in the context of the conversation the possibility of painted mules has being mentioned, and if the mere mention of this possibility has an effect on the conditions under which someone can be truly said to 'know', then any use of 'know' (or

its past tense) is so affected, even a use in which one describes one's past condition. (DeRose 1992: 925)

In other words, according to DeRose, the final remark ('No, I did know then that it was a zebra. But after your question, I no longer know') is infelicitous because the first part of the remark ('No, I did know then that it was a zebra') is to be evaluated by the same standard for knowledge as the second part ('I no longer know'). So, in short, according to Stanley, contextualists embrace the following incompatible claims:

(i)   'Knowledge' functions like a quantifier.
(ii)   Distinct occurrences of 'knowledge' within the same discourse must be associated with the same standard.

Obviously, if Stanley were right, we should abandon the idea that *knowledge* is context-sensitive or we should discard Lewis's idea that *knowledge* functions like a quantifier (or both, as Stanley does in his book). But, it is not obvious that different occurrences of the same quantified expression (at the same syntactic level)[4] within a discourse could be associated with different domains. Therefore it is not obvious that in order to function as a quantifier, different occurrences of 'knowledge' (at the same syntactic level) within a discourse could be associated with different domains.

In order to show that 'every' can vary within a single discourse, Stanley offers the following dialogue:

A: Every Van Gogh painting is in the Dutch National Museum.
B: That's a change. When I visited last year, I saw every Van Gogh painting, and some were definitely missing. (Stanley 2005: 65)

According to Stanley,

The domain for the first occurrence of the quantifier phrase 'every Van Gogh painting' is (if you like) maximally large. But we can, with no difficulty at all, understand the domain for the second occurrence to be a subset of this domain, restricted to last year's collection in the Dutch National Museum. (Stanley 2005: 65)

Under a charitable interpretation, we can make sense of the previous conversation. But Stanley's dialogue cannot be taken as the default interpretation for sentences containing quantified expressions because Stanley is bending

----

[4]   I am prepared to grant that the two occurrences of 'child' in 'Every child waved to every child' are associated with different domains and that the occurrences of 'one' in 'Someone angered everyone' are associated with different domains. But in these cases the one noun is syntactically subordinated to the other.

the rules that govern quantification. For instance, the sentence 'I saw every Van Gogh painting, and some were definitely missing' is a plain contradiction. Given the awkwardness of the dialogue in question, no theory about quantification should be designed according to it. Therefore, the felicity or infelicity of the dialogue should be judged, and maybe accommodated, only in light of a preexisting theory. In contrast, Stanley takes this dialogue as uncontroversial evidence for his view that multiple occurrences of the same quantifier phrase can take on differing values in the same discourse.

It is certainly true that some context-sensitive expressions can have different interpretation within the same discourse ('This is larger than this'), but Stanley believes that *all* context-sensitive expressions can have different interpretations within the same discourse. In fact, he explicitly offers the following generalization:

Since semantic context-sensitivity is traceable to an individual element [the expression 'she', in the case of 'She is tired'; a comparison class variable associated with the adjective 'tall', in the case of 'John is tall'…],[5] multiple occurrences of that element in a discourse should be able to take on different values. In the case of an utterance such as 'This is larger than this', where two different objects are pointed to by the person uttering the sentence, this feature is obviously confirmed. But it is present in a broader range of constructions. (Stanley 2005: 57)

Following this generalization, Stanley assimilates quantifiers to demonstratives and argues in a series of papers[6] that since in some constructions, the semantic context-sensitivity is traceable to an unpronounced element in the syntactic structure of a sentence (for example, in the case of 'John is tall' is traceable to a comparison class variable associated with 'tall'), a semantic theory for a natural language may associate different domains with different quantifiers phrases, even within a single sentence.

In short, according to Stanley, since in some structures the semantic context-sensitivity is traceable to an unpronounced element in the syntactic structure of the sentence in question, the syntax of a sentence containing a quantified noun phrase is best understood by associating each noun with a *hidden indexical* (see Gauker 2010, pages 246-247) of the form '$f(i)$': where 'f' denotes a constant function that, given an object denoted by 'i' as input, yields a domain for the noun. In this way, different occurrences of the same quantified expression at the same syntactic level can be associated with different domains, even within the same sentence.

---

[5] My addition.
[6] See for example Stanley and Szabó (2000).

But Stanley's theory must be wrong because if we can assign different domains with different quantifier phrases within a single sentence, then plain contradictions would turn out to be true. Consider the following sentence:

(I)   Every student is happy and some student is not happy.

On Stanley's theory, (I) can be interpreted as follow:

(I*)   For every j in ‹student, $f(i)$›, j is happy, and for some h in ‹student, $f(k)$›, h is not happy.

(I*) will be true in a context that assigns to '$f(i)$' a group of students who are happy and assign to '$f(k)$' a different group of students that includes a student who is not happy. But as Stanley himself admits (see Stanley and Williamson 1995: 291), (I) is clearly a contradiction, and therefore it is false in *all* contexts.[7]

As Christopher Gauker (2010) suggests, in order to vindicate the fact that (I) is a contradiction, we need to postulate the existence of *a single global, albeit contextually-determined, domain* relative to which sentences containing different occurrences of the same quantified expression at the same syntactic level has to be evaluated.

Global domains can be readily applied to knowledge attributions. Multiple occurrences of 'this' in a discourse or in a sentence can take on different values because speakers by pointing out an object as they speak can easily exploit the environment in which the conversation takes place in order to differentiate the referent of one occurrence from the referent of another. In contrast, in the case of knowledge, a speaker has no reliable way of indicating which set of possibilities he or she takes the context to assign to the word 'knowledge'.

So, given that speakers taking part in a conversation have no reliable way to indicate what they take to be the content of *knowledge*, it is safer to assume that, in any given context, there will be a unique set of possibilities associated with each occurrence of the word 'knowledge'.[8] Therefore, for any given conversation, there will be a unique set of epistemic standards *governing* it. Accordingly, if the possibility of it being a cleverly painted mule belongs to the context that governs the conversation between A and B (for example, suppose that there are several cleverly painted mules in the surroundings), then, unless B were in the position to rule it out, B does not know that it is a zebra even in the case in which the possibility in question remains unmentioned. Whereas, if the possibility of it being a cleverly painted mule does not belong to the

----

[7]   For a longer discussion of (I), see Gauker (2010).

[8]   These are the possibilities that the putative knower must be able to rule out in order to have knowledge.

context that governs the conversation between A and B, then B *does* know it is a zebra even in the case in which the possibility in question is raised and B cannot rule it out.[9]

Finally, Lewis's idea that *knowledge* functions like a quantifier plus global domains helps to understand what is wrong with the following sentence:

> (II)   If there is an external world, many normal non-philosophers know that there is, but, by contrast, no epistemologists know that there is.

According to Stanley (see page 62), (II) is a challenge for the contextualist because if *knowledge* functions like a quantifier (and therefore 'know that there is an external world' can be associated with different standards within the same sentence, and consequently can have different contents when predicated of non-epistemologists and philosophers), then (II) would be "felicitous and true". But (II) is not felicitous and true, therefore – pace Lewis – knowledge cannot function like a quantifier.

Stanley's argument is valid, but its first premise is false. With the help of global domains, the contextualist is in the position to offer a perfectly legitimate account of why (II) is not "felicitous and true": if *knowledge* calls for global domains, then there does not exist a conversational context in which (II) is true.

In conclusion, global domains plus Lewis's semantics for *knowledge* offer the linguistic basis that contextualists need.

References

Cohen, Steward, 1999, "Contextualism, Skepticism, and the Structure of Reason", in Toberlin, James ed. *Epistemology, Philosophical Perspective 13, Epistemology 1999*, Blackwell, Malden, MA: 57-90.

DeRose, Keith, 1992, "Contextualism and Knowledge Attributions", in *Philosophy and Phenomenological Research*, 52, 4: 913-929.

DeRose, Keith, 1995, "Solving the Skeptical Problem", in *Philosophical Review*, 104, 1: 1-52.

Gauker, Christopher, 2010, "Global Domains versus Hidden Indexicals", in *Journal of Semantics*, 27, 2: 243-270.

Lewis, David, 1996, "Elusive Knowledge", in *Australasian Journal of Philosophy*, 74, 4:

---

[9]   In Mion (2013a and 2013b), I argue that the distinction between those possibilities that the putative knower must be able to rule out in order to have knowledge and those that are irrelevant depend on the goals of the conversation in question.

549-567.

Mion, Giovanni, 2013a, "Epistemic Disagreements: A Solution for Contextualists", in *Studia Philosophica Estonica*, 6, 1: 15-23.

Mion, Giovanni, 2013b, "Skepticism and Objective Contexts: A Critique of DeRose", in *International Journal for the Study of Skepticism*, 3, 2: 119-129.

Schaffer, Jonathan, 2005, "What Shifts? Thresholds, Standards, or Alternatives", in Preyer Gerhard, Peter Georg eds. *Contextualism in Philosophy. Knowledge, Meaning, and Truth*, Oxford University Press, Oxford: 115-130.

Stanley, Jason, 2005, *Knowledge and Practical Interests*. Oxford University Press, Oxford.

Stanley, Jason, Zoltán G. Szabó, 2000, "On Quantifier Domain Restriction", in *Mind and Language*, 15, 2-3: 219-261.

Stanley, Jason, Timothy Williamson, 1995, "Quantifiers and Context-dependence", in *Analysis*, 55, 4: 291-295.

Giovanni Mion
Istanbul Technical University
gmion22@gmail.com