

Knowledge and Conversation

In a paper from 2014, Jonathan Schaffer and Zoltán Gendler Szabó develop a contextualist account of propositional knowledge ascriptions that exploits an analogy with adverbial quantifiers. Contextualism is often presented as an anti-skeptical strategy, but I side with Schaffer (2004) that the real issue of skepticism lies elsewhere and that we should choose the account of knowledge ascriptions that fits the linguistic data best. I argue that the analogy with adverbial quantifiers is good, but that the context-sensitivity is best understood as a merely pragmatic aspect. The appeal to the question under discussion is more naturally a pragmatic resource, and we should use it that way for adequate flexibility. Also, Schaffer and Szabó offer no satisfactory explanation of the “stickiness” of skeptical standards. I therefore advocate an infallibilist semantic account of knowledge ascriptions combined with a pragmatic account drawing from Schaffer and Szabó.

In 2004, Jonathan Schaffer suggested an infallibilist account of propositional knowledge ascription, i.e. an account according to which knowing requires absolute certainty. More recently, Schaffer has given up this account in favor of a contextualist account – according to which the truth conditions of knowledge ascriptions vary with their context – he has defended together with Zoltán Gendler Szabó (2014). Schaffer and Szabó provide an answer to what they call “the semantic problem” for contextualism: contextualism should provide an actual semantic of knowledge ascriptions that clarifies *how* their truth conditions vary. They do so by claiming that knowledge ascriptions involve a domain restriction roughly analogous to adverbial quantifiers and accordingly giving a semantics similar to the standard linguistic interpretation of them. While this analogy seems fair and the proposed semantics give the intuitively right results, I will challenge the claim that this is best construed as a *semantic*

theory. I will argue that infallibilism is the better semantic account of knowledge ascriptions, and that it can usefully be combined with Schaffer and Szabó's theory understood *pragmatically*. To this end, I will first briefly discuss the scope of the enterprise of analyzing knowledge ascriptions (I). I will then review the linguistic data (II), consider some suggested analogies to knowledge ascriptions and Schaffer and Szabó's suggestions (III) and argue that this suggestion is better understood as a pragmatic explanation given an infallibilist semantic account (IV).

I

What does it *mean* to say that S knows that P? This is first and foremost a question about the correct linguistic analysis of knowledge ascriptions. It is related to another question, namely whether we know anything about the external world. Skepticism has often been characterized as the claim that we do not *know* anything about the external world. Given this, the meaning of knowledge ascriptions becomes crucial to deciding whether skepticism is correct: if to say that S knows that P means that S believes P and can rule out any possible error with respect to her belief, then the skeptic is right because S cannot rule out that she is being deceived by an evil demon. If, on the other hand, knowledge ascriptions can be true even if the subject cannot rule out some types of error, then the skeptic will often be wrong. This line of reasoning makes a *fallibilist* account of knowledge ascriptions a desideratum because it allows us to “defeat skepticism”.

Schaffer (2004) is not too impressed by this: he argues that infallibilism only results in a “shallow skepticism” characterized by the claim that we do not *know* things about the external world. Meanwhile, the fact that we have the capacity to *discriminate* between certain possibilities, and the acknowledgment that we can improve this capacity, has to be assumed even by the skeptic. This, he argues, is a much more meaningful admission.

I may add that this characterization of skepticism should not be placed too much weight upon. Why should we characterize a position in terms of concepts we disagree about? It would be easier to distinguish between the claim that we do not know things about the external world with absolute certainty, the claim that we can never rule out all contextually relevant possibilities of error with respect to some empirical proposition, the claim that we never have justified true beliefs, and so on. If we want to defend the possibility of empirical knowledge with absolute certainty, our prospects are dim; but we are not forced to defend that claim. Much more important is the fact that our *practice* of knowledge ascriptions is justified as long as they (perhaps pragmatically) *convey* something true.

II

Let us look at the linguistic data regarding knowledge ascriptions. Let me run through a number of observations familiar from the literature.

1. *Variance*. Which knowledge ascriptions appear to be correct depends (at least) on the context of attribution. In a *skeptical context* in which we are concerned with the possibility of a deception by an evil demon, few knowledge attributions appear valid. In quotidian contexts, however, many of these attributions are felicitous.

2. *Homophonic Reportability*. An initial characterization (e.g. Cohen 1999, 61) of the contextualist perspective is to compare its view of knowledge ascriptions to indexicals. But Herman Cappelen and Ernest Lepore (2005, 86-98) have pointed out a disanalogy: indexicals cannot always be embedded in indirect speech without threatening to alter their reference. For example, if Anne said

(1) “I like strawberries”,

I cannot report (1) by saying

(2) “Anne said that I like strawberries.”

Meanwhile, knowledge ascriptions can be reported homophonically even if the context is shifted.

3. *No Shifting*. Jason Stanley (2004) points out that the reading of many candidates for context-sensitivity can change within a conversational setting, even within a sentence. This also applies to gradable adjective such as “large”. Consider (Stanley 2004, 135):

(3) “That butterfly is large, but that elephant isn’t large.”

Here, both “that” and “large” change their reference or associated standard. However, no such shift of standards within a conversational setting has been observed for knowledge ascriptions.

4. *Non-Gradability*. Stanley (2004) observes another disanalogy with gradable adjectives: these always come with comparatives such as “x is larger than y.” Comparatives for knowledge attributions, on the other hand, seem artificial. In addition, gradable adjectives can usually be modified by “really” or “very” (“x is very large”), by which the standard of largeness seems to be raised. If knowledge ascriptions are combined with such a modifier, this merely seems to add emphasis, such as in “I really know that P”.

5. *Closure*. A common principle in epistemology is that knowledge is closed under known entailment: if (a) I know that P, and (b) I know that P entails Q, then I must (c) know that Q. On the linguistic level, it seems that we can make a case for this principle at least holding within a given context, for we can object to my claim (a) on the grounds of (b) and (c). However, such an objection may only be possible within a fixed conversational setting.

6. *Non-Concessiveness*. A motivation for infallibilism that has received particular attention is the apparent infelicity of so-called *concessive knowledge attributions* (e.g. Lewis 1996, 550). Saying something like

(4) “I know that P, but I cannot rule out all possible errors”

appears to be, if not a contradiction, at least a statement that carries a tension between its two parts.

7. *Stickiness*. Given the point of Variance from above, there are a few more noteworthy details. From the perspective of a skeptical context, empirical knowledge ascriptions made in a previous quotidian

context seem *false*. But conversely, if I move from a skeptical context to a quotidian context, I will *not* be drawn to say that my denials of knowledge ascriptions in the skeptical context were false. We would at least admit that they were true in their context. Thus, there seems to be a certain “stickiness” of the skeptical contexts. Lewis (1979) describes the move to a skeptical context as a shift in the “conversational score” that is not reversible.

The linguistic project of providing an analysis of knowledge ascriptions should respect these observations and provide an explanation of each of them. However, such an explanation may be pragmatic. In fact, it is hard to see a way of incorporating all these observations in a semantic theory of knowledge ascriptions while avoiding inherent contradiction. How can Variance be concealed with Non-Concessiveness and Stickiness without making many ascriptions of empirical knowledge both true and false?

III

There is one further restraint on our theorizing: while we could postulate a meaning of knowledge ascriptions just based on the linguistic evidence, such a postulation would be *ad hoc* if it fails to integrate with our other linguistic theories. It would be implausible to claim that knowledge ascriptions exhibit an insular pattern not to be found in any other elements of our linguistic life. In the face of this, many have exploited analogies to other expressions to motivate their suggestions.

As mentioned, Homophonic Reportability points to an important difference between indexicals and knowledge ascriptions, and Non-Gradability and No Shifting mark a distinction between knowledge ascriptions and gradable adjectives. The latter point is noteworthy, as epistemologists of both the infallibilist (Unger 1975) and fallibilist (e.g. Cohen 1999; Lewis 1996, 554) camp have tried to exploit gradable adjectives as an analogy. Lewis’s (1996, 553) central analogy are quantifiers, specifically “every”. The known proposition is supposed to be true in every possibility not eliminated by the

knower's evidence, but the domain of "every" is restricted because some possibilities can be properly ignored. While the analogy is of great explanatory help, Schaffer and Szabó (505-7) have pointed out that "every" also fails the criterion of Homophonic Reportability: the relevant domain restriction may get lost when we report a quantified statement in a different context. Such quantifiers also fail the No Shifting criterion, as an example by Stanley and Timothy Williamson (1995) illustrates:

(5) "Every sailor waived to every sailor"

may mean that every sailor on ship A waived to every sailor on ship B (and conversely), so that the two quantifiers would have different domains.

But there are other candidates for analogies. Schaffer and Szabó (507-15) suggest adverbial quantifiers, specifically "always".¹ For example, the sentence

(6) "Claire always steals the diamonds"

is standardly interpreted as quantifying over (actual) situations (or "cases", cf. Lewis 1975). But the relevant domain of situations will depend on the context: we might be saying, for example, that in all of her burglaries, Claire steals the diamonds (rather than the money), or we might be saying that whenever some diamonds (in a certain area, maybe) are stolen, Claire is always the thief.

According to Kai von Fintel (2004), we can best think of the domain of adverbial quantifiers as restricted by a contextual variable. Schaffer and Szabó (522-4) argue that this variable is provided by the *question under discussion* (QUD), e.g. "Who stole the diamonds?" or "What did Claire steal?" This QUD is thought of as a set of propositions (*alternatives*), where these alternatives are sets of possible worlds. We can define a *partial answer* as a statement that entails an evaluation of at least one of the alternatives, whereas a *complete answer* would evaluate all of them (Roberts 2012). "Always" here only quantifies over situations which satisfy the *presupposition* of the QUD (e.g. "someone stole the

¹ Another interesting analogy are counterfactuals (Ichikawa 2011, Lewis 2017). If we evaluate them by reference to possible worlds, there appears to be a context-dependent restriction of the domain of possible worlds in which the material conditional is supposed to hold. This domain restriction strikingly exhibits the features mentioned above.

diamonds” or “Claire stole something”). Schaffer and Szabó suggest that knowledge ascriptions may be similarly understood, namely as truth in all possibilities that satisfy the presupposition of the QUD.²

The analogy with “always” works well in many respects. In particular, we find Closure intuitive and the mechanism of Stickiness works similarly. Also, “Claire always steals the diamonds, but in some situations she does not.” has a contradictory feel to it, corresponding to Non-Concessiveness. Furthermore, these adverbs can at least usually be reported homophonically and are not gradable. A difference to knowledge ascriptions appears to be the fact that the domain of situations can shift within a context, e.g. in:

(7) “Alice always promotes employees if they always show up to work on time.”³

Despite this, the analogy is close enough to dispel the worry that knowledge ascriptions are an insular issue.

So the analogy with “always” is acceptable. Schaffer and Szabó’s explanation of knowledge ascriptions furthermore can explain the relevant data: the different QUD’s in different contexts can account for Variance. Meanwhile it is also easy to see that pointing to the existence of uneliminated alternatives, as happens in concessive knowledge attributions, will often alter the QUD. A similar idea seems to apply to the issue of Stickiness: when we are in a skeptical context, our QUD accommodates worries about Cartesian deception, and we are thus inclined to deny not only that we have empirical knowledge, but also that knowledge ascriptions we have made in the past that implied the contrary were false. When looking back at this skeptical context, these worries are re-introduced, or at least it is recognized that

2 I am omitting here the optional explicit domain restriction, which is possible both for quantificational adverbs (“always except on Saturdays”) as well as for knowledge ascriptions (“knows that P if she assumes Q”, cf. Schaffer and Szabó 2014, 529).

3 Schaffer and Szabó (535) point out that there are *some* constraints on the shifting of the domain of “always”, for example “Claire never steals anything, and she always steals the diamonds” can never be felicitous even though both parts of it can be acceptable in the right contexts (if she never steals from accomplices). But that much is true of regular quantifiers, too. For example, “I took everything from the fridge and I left something in the fridge” could not be felicitous, even if I can use the parts of it to express that I took all the food from the fridge and that I left the light bulbs in the fridge. The observation about knowledge ascriptions we make is that there are *no* such shifts.

these worries were on the table in the skeptical context, so we are not tempted to retreat our denial of empirical knowledge made in that context.

IV

So Schaffer and Szabó can explain the data, but *how*? I wish to argue that what they give is essentially a good *pragmatic* explanation. My claim is that it would be most suitable to accept a semantic meaning of knowledge that fully accommodates the infallibilist intuitions and view the domain restriction introduced by the QUD as a pragmatic alteration of that meaning. This claim might ultimately entail that the linguistic view of domain restrictions for quantificational adverbs ought to be construed in a similar fashion, although that is beyond my scope here. The distinction between pragmatics and semantics has become blurred, so I should add a qualification: my point here is not to insist on the “true” meaning of semantics, but rather that there is an important level of meaning which should feature in our theory of knowledge ascriptions, and which lies *before* a restriction of the domain of possibilities.

I would like to make three points in favor of this view. First, the resource Schaffer and Szabó use to explain the linguistic data are more easily understood as pragmatic. Von Stechow (2004) characterizes the contextual variable of adverbial quantifiers as a “[hole] in the semantic structures which will be filled by the pragmatics”. And Craige Roberts (2012, 36-46), who introduced the concept of QUDs, is aiming at an integrated theory of pragmatics, and she even uses them to argue against the claim that domain restrictions are introduced semantically by aspects like focus (partly contrary to what Schaffer and Szabó (524-7) claim). The idea is this: given a context provides a QUD, *how* does this restrict quantifier domains? In virtue of the fact that participants of the conversation mutually accept the QUD and make an effort to provide an answer to it, at least a partly one; and, crucially, that they mutually assume their conversational partners are also making such an effort. This assumption is just another

way of stating Grice's Principle of Cooperation. From this we can derive that speakers are trying to say something relevant. But of course saying something about situations or possibilities that contradict the presupposition of the QUD must always be *irrelevant*; even more so, when it would clearly be false. On the other hand, it is not clear why we should assume that the QUD can modify the domain of possibilities if there is no semantic meaning associated with the utterances in question yet. To be sure, I see that there is a mechanism of arriving at results, but this mechanism does not provide an *explanation*.⁴

The idea of relevance to the QUD leads to a second point. Roberts (2012, 20; cf. Schaffer & Szabó 2014, 523) goes further and says that to be relevant, an assertion must either introduce a partial answer or be part of a strategy of arriving at such an answer. This gives us better resources for explaining the precise restriction of the domain than merely appealing to the presupposition of the QUD. Just appealing to the (logical) presupposition alone will not always give us the intuitively right results. On the one hand, we cannot accommodate strategic assumptions that may be common ground; on the other hand we might be forced to accept attributions of knowledge to people who do not believe the presupposition of the QUD even though they intuitively only possess conditional knowledge. The pragmatic account I am suggesting is more flexible and allows us to exploit strategic assumptions; it also allows us to exempt subjects from the domain restriction who do not believe the presupposition of the QUD.

Thirdly, I think Schaffer and Szabó *are*, despite their claiming otherwise (534-5), actually still vulnerable to a meaningful objection from Stickiness and "semantic blindness" (i.e. the claim that ordinary speakers misattribute truth values due to Stickiness, which contextualists are typically

4 The explanation I am hinting towards here escapes Schaffer and Knobe's (2012) objections because (a) I assume the semantic meaning to be infallibilistic, so most knowledge ascriptions are actually false (which explains why we can be driven to retract them on closure-based arguments), but (b) the pragmatic weakening does occur on grounds of assumed cooperation, not based on incomplete processing (which plausibly also plays a role in some cases).

committed to). Let me give an illustration: suppose we are in a skeptical context c_1 , where the QUD is concerned with the possibility of absolute certainty. I say:

(8) “We don’t know that smoking causes cancer.”

I am right to say so, here, because of scenarios in which smoking does not cause cancer, but all scientists claiming the opposite are either mistaken due to an incredible coincidence in their choice of samples, or have manipulated their studies, or do not even exist.

Now, say I move on to a quotidian context c_2 where we specifically discuss the effects of smoking and what we have learned from science about them. I say:

(9) “We know that smoking causes cancer.”

Again, this seems felicitous: ignoring the possibilities mentioned above is fine here, for they would preclude us from ever saying anything about the effects of smoking (a strategical assumption). But what is more noteworthy, even in c_2 I will still be inclined to say that (8) was correct, at least when uttered in c_1 . However, had we moved from c_2 to c_1 , I would be inclined to deny the truth of (9). This is just an instance of Stickiness described above; but it is surprisingly hard for Schaffer and Szabó to deal with it.

The main problem here is that a semantic theory has to be understood as some form of *attributer contextualism*. Attributer contextualists claim that the truth conditions are sensitive to the context in which the knowledge ascription was made.⁵ This predicts the right result for the move from c_1 to c_2 , but the intuitively wrong result for the move in the other direction. Schaffer and Szabó (534) adapt an “error theory” about our evaluation of (9) from c_1 . But I think they lack a good explanation of this error. The error, one could argue, arises because we evaluate knowledge ascriptions *as if* they were made in the current context. But this raises the question why we accept (8) as true even from the perspective of

⁵ We might wish to say that the relevant QUD is to be found in the context of *evaluation*. But (semantic) *evaluator contextualism* would assign multiple truth conditions to an utterance, one for each context of evaluation. This is not feasible for a semantic theory. My suggestion is to adopt a pragmatic strategy that can allow for the evaluation of knowledge ascriptions to be also guided by the QUD in the context of evaluation.

c2. One could argue that looking back at c1 re-introduces an epistemological QUD. But this does not always work; specifically not in this case. The QUD in c2 is an important one, and this forbids such maneuvers of distraction. If we need to make an important decision, we will insist that we *do* know that smoking causes cancer even when someone reminds us of a possibility of error. But even then we would still say that (8) was true, at least in its context. If we posit an infallibilist semantic meaning of knowledge ascriptions, we have a straightforward explanation of this: we still accept that (8) was true because *it is* semantically true. But (9) communicates something true and important given our QUD, so we also insist on its correctness in c2.

These three points suggest that assuming an infallibilist semantic meaning gives us a better explanation of the linguistic data. We can then appeal to the QUD like Schaffer and Szabó, but we have an explanation why and how the QUD interferes with the communicated content. We also gain a more elegant way to account for the “stickiness” of skepticism. As I have mentioned at the beginning, this does not constitute a triumph for skepticism, for admitting an infallibilist semantic meaning does not threaten our practices of ascribing knowledge. A worrisome kind of skepticism would be one that rejects even discussing questions that presuppose or otherwise require that there exists an external world. The semantics of knowledge ascriptions do not justify anything like this.

References

- Cappelen, H. and E. Lepore (2005). *Insensitive Semantics*. Oxford University Press.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Noûs* 33, 57-89.
- von Stechow, K. (2004). A minimal theory of adverbial quantification. In H. Kamp and B. Partee (eds.), *Context-dependence in the analysis of linguistic meaning* (137–175). Elsevier.
- Ichikawa, J. (2011). Quantifiers and epistemic contextualism. *Philosophical Studies* 155, 383–398.

- Lewis, D. (1975). Adverbs of Quantification. In E. Keenan (ed.), *Formal Semantics of Natural Language* (3-15). Cambridge University Press.
- Lewis, D. (1979). Scorekeeping in a Language Game. *Journal of Philosophical Logic* 8, 339-359.
- Lewis, D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy* 74, 549-567.
- Lewis, K. (2017). Counterfactuals and Knowledge. In J. Ichikawa (ed.), *The Routledge Handbook of Epistemic Contextualism* (411-424). Routledge.
- Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics* 5 (article 6), 1–69.
- Schaffer, J. (2004). Skepticism, Contextualism and Discrimination. *Philosophy and Phenomenological Research* 69, 138-155.
- Schaffer, J. (2015). Lewis on Knowledge Ascriptions. In id. and B. Loewer: *A Companion to David Lewis* (473-490). Wiley Blackwell.
- Schaffer, J., and J. Knobe (2012). Contrastive knowledge surveyed. *Noûs* 46, 675–708.
- Schaffer, J. and Z. Gendler Szábo (2014). Epistemic comparativism: a contextualist semantics for knowledge ascriptions. *Philosophical Studies* 168, 491–543.
- Stanley, J. (2004). On the Linguistic Basis for Contextualism. *Philosophical Studies* 119, 119-146.
- Stanley, J. and Williamson, T. (1995). Quantifiers and Context-Dependence. *Analysis* 55, 291–295.
- Unger, P. (1975). *Ignorance. A Case for Scepticism*. Oxford University Press.