Synthese (2023) 201:13
https://doi.org/10.1007/s11229-022-04010-4

ORIGINAL RESEARCH

From contrastivism back to contextualism

Da Fan[1]

Received: 1 April 2022 / Accepted: 14 December 2022 / Published online: 31 December 2022

Abstract
Contrastivism is the view that knowledge is a ternary relation between an agent, a content proposition, and a contrast, and it explains that a binary knowledge ascription sentence appears to be context-sensitive because different contexts can implicitly fill the contrast with different values. This view is purportedly supported by certain linguistic evidence. An objective of this paper is to argue that contrastivism is not empirically adequate, as there are examples that favor its contextualist cousin. Thereafter, I shall develop a contextualist account for the relevant linguistic data. The account consists of a contextualist semantics and some rules of pragmatics. The two parts combined show that contrastivism is neither sufficient nor necessary as a satisfactory theory of knowledge ascriptions.

Keywords Knowledge ascription · Contrastivism · Contextualism · Pragmatics

1 Introduction

Knowledge-*that* has long been thought of as a binary relation between an agent and a content proposition. Accordingly, in the search for an appropriate linguistic model for knowledge ascription sentences, it is natural to take "to know" as a binary verb, which is to combine a subject noun phrase and a *that*-clause to form a complete knowledge ascription sentence. But this view is challenged by *epistemic contrastivism*, which claims that a binary ascription, such as "Moore knows that he has hands," is incomplete in itself, in the same way that "I prefer ice cream" is incomplete: whether I prefer ice cream depends on what alternatives are under consideration; I prefer ice cream to pretzels but not to cinnamon rolls. According to the contrastivist, knowledge ascriptions are similar: it is true that Moore knows that he has hands rather than stumps, but that doesn't entail that he knows he has hands rather than, say, brain-in-a-vat hand-images. In other words, knowledge ascriptions appear to be sensitive to

✉ Da Fan
phi.fanda@gmail.com

[1] School of Philosophy, Wuhan University, Wuhan 430072, Hubei Province, China

contrasts: a binary knowledge ascription is true with respect to some contrasts but not to others, just as a binary ascription of preference is true with respect to some contrasts but not to others. Hence, the contrastivist claims that knowledge is a *ternary relation* between an agent, a content proposition, and a contrast proposition, in the same way as preference is a ternary relation between an agent, an object that is preferred, and an object that is preferred to. It is natural, according to the contrastivist, to assume that the best linguistic model for knowledge ascription sentences should treat "to know" as a ternary verb, and the complete form of knowledge ascriptions is thus $Kspq$, or "s knows that p rather than q."

In this paper, I will argue against epistemic contrastivism and defend the received wisdom that an appropriate linguistic model for knowledge ascriptions is binary. As illustrated above, our truth-value judgments on knowledge ascription sentences shift according to different contrasts provided, and this apparent contrast-sensitivity is supported by empirical evidence (Schaffer & Knobe, 2012) . As epistemic contrastivism predicts this empirically attested shiftiness, this lends some strength to the position on grounds of empirical adequacy. However, I will argue that empirical evidence is actually against epistemic contrastivism, as various examples show that there are important dissimilarities between knowledge ascriptions and ascriptions of paradigmatically ternary relations. Thus, epistemic contrastivism is not empirically adequate.

Although I argue against the empirical adequacy of the view, the data the contrastivist alludes to are still illuminating: at least, they show that knowledge ascriptions appear to be sensitive to contrasts. Thus, with contrastivism rejected, we still need an alternative account of why knowledge ascription sentences behave as such. I will propose a binary treatment within the contextualist framework. The account aims at explaining the apparent contrast-sensitivity by appeal to a contextualist semantics and certain rules of pragmatics. First, it adopts a contextualist semantic framework for knowledge ascription sentences, according to which the content of such a sentence depends on a contextual parameter that can be resolved by some pragmatic processes. Second, I will propose some pragmatic rules governing how contextual parameters should relate to utterances of knowledge ascription. By appeal to these pragmatic rules, I will explain how contrast-triggering expressions can affect the contextual parameter to which knowledge ascription sentences are sensitive. Hence, according to the account, the apparent contrast-sensitivity is explained through a pragmatic detour: contrastive constructions are able to influence some aspects of the context, which in turn affect the content expressed by knowledge ascription utterances.

The remainder of this paper is organized as follows. In Sect. 2, I will briefly lay out the basic tenets of epistemic contrastivism, together with its purported empirical evidence that knowledge ascriptions are sensitive to contrasts. In Sect. 3, I present three arguments that are intended to show that empirical evidence is in fact against the contrastivism. In Sect. 4, I provide a contextualist account of the apparent contrast-sensitivity of knowledge ascriptions. Specifically, after summarizing some success conditions, I will formulate the semantic and pragmatic constituents of the account and apply them to the relevant data.

2 Contrastivism

Since epistemic contrastivism is formulated and motivated in various ways in the literature, I will only focus on two prominent aspects of the view, without surveying every variant of it.[1] First, epistemic contrastivism is a view on *the concept of knowledge*, according to which knowledge is a ternary relation. Call this thesis TERNICITY.

> **Ternicity.** Knowledge is a ternary relation of the form $Kspq$, which holds between an agent s, a content proposition p, and a contrast proposition q.[2]

Ternicity is only a claim about the form of the knowledge relation, and a complete contrastive theory of knowledge ultimately needs a plausible analysis of this ternary relation K. But roughly speaking, Ternicity is inspired by the idea that knowledge is relative to a class of contrast propositions, or "relevant alternatives", in the sense that one's knowing something amounts to a certain class of alternative propositions being ruled out by her epistemic state (cf. Dretske, 1972). The epistemic state of an agent may include everything she has as evidence, such as her sensory experiences, memory, etc., and that epistemic state's ruling out a proposition means that in every possible world where the agent has exactly that epistemic state as she actually has, the proposition is false.

According to contrastivism, to know is always to know some p by ruling out some propositions that are incompatible with p, the disjunction of which is the contrast proposition q. One's ruling out one class of contrasts does not mean that she can rule out others. When ascribing knowledge, we are ascribing a ternary relation $Kspq$, meaning that *s knows that p rather than q*; this, however, by no means implies that s also knows p rather than q'. For example, Moore knows that he has hands rather than stumps, but he doesn't know that he has hands rather than BIV hand-images, because his epistemic state does rule out the stump-possibilities but in no way rules out the BIV possibilities.

However, knowledge ascription sentences, in ordinary language use, are often binary in the surface form *s knows that p*, which does not appear to involve a contrast proposition. Contrastivists, such as Schaffer (2004), take the binary surface form as incomplete and claim that the contrast proposition is implicitly fixed by the context in such cases. In this sense, epistemic contrastivism is similar to contextualism, because they both allow the content of a (binary) knowledge ascription sentence to shift with contexts. Moreover, if the contrast can only be determined contextually, then

[1] Major advocates of epistemic contrastivism include Karjalainen and Morton (2003), Sinnott-Armstrong (2004), Sinnott-Armstrong (2008), Schaffer (2004), Schaffer (2005), Schaffer (2008), and Morton (2012).

[2] Contrastivists may not have a consensus on what "knowledge is a ternary relation" means. For example, Schaffer (2004) defends the view that this ternary relation is the denotation of the verb "to know" in our ordinary language, according to which contrastivism amounts to a descriptive theory of our knowledge ascription language and the concept of knowledge that is denoted by it. On the other hand, Sinnott-Armstrong argues for a revisionist version of contrastivism, which indicates that the ternary relation is what we should make use of if the goal is to "describe a person's epistemic position as precisely as possible" (2008, p. 268). This paper is focused on the descriptive version of contrastivism that is put forth in a series of papers by Schaffer.

contrastivism is nothing more than a particular version of contextualism.[3] However, Schaffer (2004) distinguishes the two views in terms of how contrasts are fixed:[4]

> [C]ontext-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. (Schaffer, 2004, pp. 82–83)

According to this view, the linguistic model of the verb "to know" directly reflects the ternicity of the concept of knowledge, in the same way as "to prefer" directly reflects the ternary relation of preference. Thus, there are two ways that the contrast proposition is determined. First, if a knowledge ascription sentence is ternary in its surface form, the contrast proposition is directly determined by the contrast clause. For example, the *rather-than* clause in (1) determines the contrast proposition to be that Moore has stumps.

(1) Moore knows that he has hands rather than (he has) stumps.
(2) Moore prefers ice cream to pretzels.

There is no mystery in this case: the contrast proposition is explicitly determined, just as the third argument in the preference relation expressed in (2) is determined by the explicit complement phrase.

On the other hand, for an apparently binary knowledge ascription sentence, the contrast is contextually determined. For example, when asked if Moore can tell whether hands or stumps are attached to the end of his arms, we can leave out the contrast and just say (3) to expresses the same content as (1). This is comparable to apparently binary preference ascription sentences, such as (4). When the context is clear—e.g. when asked which of ice cream and pretzels Moore prefers, (4) is felicitous to utter and in effect expresses the same content as (2).

(3) Moore knows that he has hands.
(4) Moore prefers ice cream.

Summarizing the two ways in which the contrast is determined, the contrastivist is committed to the following linguistic claim:

[3] Ternicity alone cannot distinguish contrastivism from contextualism. In the literature, the contextualist can treat "to know" as either an indexical or a predicate that has a third argument (cf. Bach, 2005, Montminy, 2008, and Baumann, 2016). The latter kind of contextualism is compatible with Ternicity. Thanks to an anonymous reviewer for this point. However, I do not think these contextualist views can agree with the other contrastivist thesis introduced below, i.e. Saturation. For the contextualist, even if she stipulates a third argument position of "to know", she treats it as an argument whose value has to be determined by the context. On the other hand, if the contrastivist is right, the third argument can be saturated by contrast expressions alone.

[4] Sometimes the contrastivist takes the view as a variant of contextualism in the sense that they both allow the shiftiness of binary knowledge ascription: "This is because the contrastive view allows that one ascriber could truly say 's knows that p,' while a second ascriber in a second context (with a different range of relevant alternatives) could truly deny 's knows that p.' " (Schaffer & Knobe, 2012, p. 687) However, this is just a terminological difference. I will use the term "contextualism" in a stronger sense that excludes contrastivism, and the distinction is made by the thesis of SATURATION below.

Saturation. The verb "to know" denotes the ternary relation $Kspq$. In a knowledge ascription sentence, if there is a contrast expression (e.g. *rather-than*), the value of q is explicitly saturated by it; otherwise, the value is implicitly saturated by the context.

With Saturation, contrastivism can be separated from contextualism, even though both allow knowledge ascriptions to be context-dependent. In short, binary ascription is the paradigm case for the contextualist, and the shiftiness of binary ascription sentences is explained by the context-dependency of the verb "to know". For a ternary ascription sentence, insofar as the contextualist allows explicit contrast expressions to affect the context, the truth-conditional effect of these expressions can be explained by the context-dependency of "to know". On the other hand, the contrastivist views ternary ascriptions as the paradigm and explains binary ascriptions, and the context-dependency thereof, in terms of the context playing the role of completing an incomplete sentence.

With the view of contrastivism laid out and separated from contextualism, the next step is to examine some arguments for it. In the literature, contrastivism is motivated and argued for from various perspectives. I will focus on its empirical adequacy.

In a survey done by Schaffer and Knobe (2012), participants are asked to read the following [JEWEL THIEF] vignette and then judge to what degree they agree with various knowledge ascription conditions. The results show that the judgments are sensitive to contrasts triggered in several different ways. In what follows, I use '#' to mark conditions the participants tend to disagree with, according to the data.

> [Jewel Thief] Last night, Peter robbed the jewelry store. He smashed the window, forced open the locked safe, and stole the rubies inside. But Peter forgot to wear gloves. He also forgot about the security camera. Today, Mary the detective has been called to the scene to investigate. So far she has the following evidence. She has been told that there was a theft, she has found and identified Peter's fingerprints on the safe, and she has seen and recognized Peter on the security video, filmed in the act of forcing open the safe. She has no further information.

(5a) Mary knows that Peter rather than anyone else stole the rubies.[5]
(5b) # Mary knows that Peter stole the rubies rather than anything else.
(6a) Mary knows who stole the rubies.
(6b) # Mary knows what Peter stole.

The above two pairs are cases where contrasts are signaled explicitly by the *rather-than* construction and *wh*-clauses. Limiting the scope of this paper to knowledge-*that* ascription, I simply list (6a) and (6b) here without discussing them.[6]

[5] Originally in Schaffer and Knobe (2012), these sentences read 'Mary now knows...' But here the word 'now' is dropped. Given that Gerken and Beebe (2016, pp. 139–142) are able to replicate the experimental results, and given that the conditions they use are exactly like those used by Schaffer and Knobe except that 'now' is dropped, it seems fair to present Schaffer and Knobe's data with this minor change.

[6] Schaffer and Knobe (2012) do think that contrastivism successfully explains the contrastive effect shown by the pair in (6a). One reason why I set them aside is that it seems not clear to me how the contrastivist theses, i.e. Ternicity and Saturation, could be extended to cover knowledge-*wh* ascriptions, especially when the aim is a compositional semantics that preserves the uniformity of "to know". For example, Schaffer (2009)

Also, (7a) and (7b), uttered in [C1] and [C2], display a similar effect without explicit contrast expressions.

[C1] Everyone is now asking the big question: Who stole the rubies? The news reporter is about to write a story about Mary. He is wondering if Mary now knows who stole the rubies. He writes:

(7a) Mary knows that Peter stole the rubies.

[C2] Everyone is now asking the big question: What did Peter steal? The news reporter is about to write a story about Mary. He is wondering if Mary now knows what Peter stole. He writes:

(7b) # Mary knows that Peter stole the rubies.

In addition, Schaffer (2008) indicates that *it*-cleft and focus can also elicit similar contrastive effects. This claim, unfortunately, is not tested in the above survey, but I nevertheless assume its truth and thus include the following pairs as part of the data alongside (5)–(7).

(8a) Mary knows that it was Peter that stole the rubies.
(8b) # Mary knows that it was the rubies that Peter stole.
(9a) Mary knows that PETER stole the rubies.
(9b) # Mary knows that Peter stole THE RUBIES.

Call these results collectively *the contrastive data*. Contrastivism provides a straightforward account of the difference between these a-sentences and b-sentences, assuming that *rather-than*, *it*-cleft, and focus serve to explicitly signal contrasts and that topical questions in the context can do so implicitly. Although the content propositions of all the above sentences are the same, due to the difference in aspects such as *rather-than*, the contrast propositions of the a-sentences are different from those of the b-sentences. In particular, all the a-sentences have the contrast proposition that someone other than Peter stole the rubies, which is ruled out by Mary's epistemic state, but all the b-sentences have the contrast that Peter stole something other than the rubies, which isn't ruled out. This explains why all the a-sentences are true while the b-sentences are false: the ternary knowledge relation holds between Mary, the content proposition that Peter stole the rubies, and the contrast proposition that someone else stole the rubies, but it doesn't hold if the contrast proposition is that Peter stole something else.

If that is the case, the intuitive difference between the a-sentences and b-sentences is explained by different propositions' filling into the contrast position: the logical form of the a-sentences is $Kspq$, whereas that of the b-sentences is $Kspq'$. Insofar as the contrastive data are accounted for, one may think that contrastivism is an empirically

Footnote 6 continued
proposes that a knowledge-*wh* ascription, with Q being the question that corresponds to its embedded *wh*-clause, is true iff there is a proposition p such that $KspQ$ and p is the true answer to Q. But I do not see how this existentially quantified truth condition could be unified with the version of contrastivism for knowledge-*that* ascription, so it seems better, for my purposes, to focus on only knowledge-*that* ascriptions at this point. For Schaffer's view on knowledge-*wh* ascriptions and relevant discussions, see Schaffer (2007b), Schaffer (2009), Brogaard (2009), Kallestrup (2009), Aloni and Égré (2010), and Steglich-Petersen (2014).

adequate view with regard to the language use of knowledge ascriptions. However, I argue against the view from two perspectives. First, as will see in the next section, I show, by appeal to various examples of knowledge ascription, that contrastivism is empirically *in*adequate. Second, in Sect. 4, I develop an alternative explanation within the good old contextualist framework, which better explains the examples considered earlier.

3 Arguments against contrastivism

3.1 Defensible negations

According to the contrastivist, the ascription (5a) has a complete ternary form, where the contrast argument position is saturated by the *rather-than* expression. Moreover, it is true because the ternary relation holds between the subject, the content, and the contrast. If that is correct, then (5a) is true in the same way as true ascriptions of other ternary relations, e.g. (10).

(10) Alice introduced Beth to Carl.

Suppose Alice did introduce Beth to Carl, (10) is true because it ascribes the introduction relation to Alice, Beth, and Carl. Consequently, we expect its negation, the denial of such a relation holding between the three, to be unarguably false. By the same token, the negation of (5a) should be unarguably false as well.

(11) Alice didn't introduce Beth to Carl.
(12) Mary doesn't know that Peter rather than anyone else stole the rubies.

However, there is a disanalogy between the two negations. While (11) is a flat-out falsehood, (12) is defensible. Suppose someone challenges the speaker of (12) by asking: "But Mary saw the footage of Peter force opening the safe and identified his fingerprints on the safe, how could she not know that Peter rather than anyone else stole the rubies?" The speaker would be able to respond: "Mary doesn't know that it was the rubies that were stolen, so she doesn't know that Peter rather than anyone else stole the rubies" and thus stand by what she said, without retracting it.[7] The case for (11), however, is very different. When someone points to the facts, the speaker of (11) has to retract what she said, as she just cannot defend it.

This dissimilarity indicates that while (11) invariably expresses a false proposition, what is said in (12) does not have to be false, and thus its opposite, (5a), does not have to be true. Thus, for a knowledge ascription that has a ternary form, such as (5a), there

[7] An anonymous reviewer raises the concern that this response presupposes the denial of contrastivism. In particular, the speaker's defense, that Mary doesn't know that Peter rather than anyone else stole the rubies, is directly against contrastivism. According to contrastivism, as Mary *can* rule out the contrast, i.e. that someone else stole the rubies, the knowledge denial should be false. However, the argument above is meant to rely on our intuitive judgment about the speaker's defense. As it appears to me, the speaker's defense is somewhat acceptable, while it is hard to imagine how (11) can be defended in any sensible way. If there is this difference between (11) and (12), I think it suffices for the main point here: unlike ternary ascriptions of paradigmatically ternary relations (e.g. the introduction relation), a ternary knowledge ascriptions can be true in a context and be false in another.

can be two distinct interpretations. The contrastive data show that (5a) can have an interpretation on which it is true, whereas the above example shows that there is an interpretation on which it is false. Moreover, the speaker's defense, in effect, makes the latter interpretation more natural. Thus, even for knowledge ascriptions that have a ternary form, they can still be shifty. This is utterly differnt from paradigmatically ternary predicates such as "introduce": binary ascriptions of a ternary relation can be shifty because their incomplete surface form does not determine a unique logical form, but ternary ascriptions of them are not shifty, as the logical form is completely determined—the three argument positions are all saturated.

As the contrastivist uses ternary verbs as linguistic model for knowledge ascriptions, the above example is a challenge to the view. According to Saturation, (5a) has its contrast proposition saturated by the *rather-than* expression, and thus it should invariably express a true proposition, which is, as illustrated, incorrect. This suggests that the shiftiness of knowledge ascriptions may not be understood as comparable to the shiftiness of, say, introduction ascriptions. For the latter, they are shifty only when they have a binary, incomplete surface form. However, knowledge ascriptions appear to be shifty even when the surface form is ternary. Thus, using ternary verbs as a linguistic model of "to know" may be wronghead in explaining the shiftiness of knowledge ascriptions.

3.2 Odd sequence

This point is to be further attested by more examples that display the variation of the perceived content of the a-sentences in different contexts. First, note that speakers judged (5a) as true but (5b) false, according to the contrastive data. This happens when we judge them individually. However, if the negation of (5b) is uttered first and followed by (5a), the previous judgment no longer holds.

(13) Mary doesn't know that Peter stole the rubies rather than anything else. But she knows that Peter rather than anyone else stole the rubies.

Something strange is easily detected, and we are much less willing to say that (13) is true than we used to when saying (5b) is false and (5a) true.[8] The contrastivist takes the logical forms of (5a) and (5b) as $Kspq$ and $Kspq'$ respectively and deems the former to be true and the latter false. It follows that (13) is just saying $\neg Kspq' \wedge Kspq$, which should be true. However, this prediction is incorrect.

Furthermore, we can generate counterexamples like (13) easily in a systematic way. Let $Kspq$ be a knowledge ascription that the contrastivist deems as true. According to contrastivism, for $Kspq$ to be true, the person s doesn't have to rule out all the $\neg p$-possibilities; rather, ruling out all the q-possibilities is enough. Then, here is a way to turn $Kspq$ false by manipulating the context: add something to the context

[8] Theoretically, it is possible for the contrastivist to give a pragmatic account of the infelicity of (13). She can insist that the sequence expresses two true propositions, while explaining the infelicity by pragmatics. Thanks to an anonymous reviewer for this point. However, it is not clear to me what pragmatic account is available for the contrastivist, as similar sequences of ternary relation ascriptions does not seem to violate any pragmatic rules. For example, (18) below sounds felicitous. In the end, I think it is the burden is on the contrastivist to show that there are pragmatic accounts that work in her favor.

indicating that s can't rule out q', where (i) q' is inconsistent with p and (ii) doesn't entail q. By (ii), we avoid contradicting the truth of $Kspq$, as by the contrastivist's lights $Kspq$ entails that s can rule out q and thus everything that entails q. By (i), we push the speakers' intuition towards the rejection of $Kspq$: since s can't rule out q' and q' is inconsistent with the content p of the purported knowledge, then p can't be ascribed to s as knowledge.

To illustrate, suppose Goldman is seeing a (real) barn and he can rule out the possibility that the object is a house. Then (14) sounds true, in a normal, non-philosophical context.

(14) Goldman knows that it's a barn rather than a house.

Applying the above method, we need to find a proposition that is inconsistent with that the object is a barn, and that doesn't entail that it is a house. Here is an unsurprising one: the proposition that the object is a barn facade. Adding that to the context, we get:

(15) Goldman doesn't know that it's a barn rather than a barn facade. But he knows that it's a barn rather than a house.

This, like (13), doesn't sound good, or at least sounds much less acceptable than (14) uttered alone.

Compared to the paradigm ternary predicate "to introduce" again, it is clear that "to know" deserves a different treatment. Suppose Alice introduced Beth to Carl but didn't introduce Beth to Catherine. Then (16) is true and (17) is false.

(16) Alice introduced Beth to Carl.
(17) Alice introduced Beth to Catherine.
(18) Alice didn't introduce Beth to Catherine. But Alice introduced Beth to Carl.

Obviously, uttering the negation of (17) and then (16) in one breath, as in (18), comes out true. This result is well expected, given the ternicity of "to introduce": In both (16) and (17), all the argument places of the predicate are explicitly filled by appropriate objects, so when uttered in the sequence (18), the truth values do not change. However, "to know" does not behave the same way.

3.3 Odd content expression

The above cases are intended to show that a knowledge ascription remains shifty even when an explicit contrast expression is present. Similar examples are found in the literature on the issue of epistemic closure. Famously, the traditional, binary version of closure apparently leads to either skepticism, i.e. the denial of (P1), or dogmatism, i.e. the acceptance of (C).

(P1) Moore knows that he has hands.

(P2) If one knows that p, and if p entails q, then one knows that q.[9]
(C) Moore knows that he is not a brain in a vat (BIV).

Schaffer (2004, 2007a) offers a contrastivist solution. As knowledge ascriptions are treated as ternary, Schaffer (2007a) advocates the following ternary version of closure.[10]

> **Expand**-p. If one knows that p rather than q, and if p entails p' and p' is incompatible with q, then she knows that p' rather than q.[11]

The contrastivist argues that if we accept the ternicity of knowledge relation and substitute this contrastive version of closure for the binary one, the undesirable inferences do not go through. First, Moore doesn't know that he isn't a BIV. But by the contrastivist light, the complete form of this is that Moore doesn't know that he isn't a BIV *rather than he is a BIV*, which is true. But this does not mean that Moore does not have ordinary knowledge. By Expand-p, what follows from Moore's non-BIV ignorance is only that he doesn't know that he has hands *rather than he is a BIV*. Indeed, Moore *is* ignorant of this, as he can't rule out the BIV possibility. However, this is not what we ascribe to Moore as ordinary knowledge. Rather, Moore's ordinary knowledge must have an ordinary contrast that he is able to rule out, e.g. he knows that he has hands rather than stumps. On the other hand, this ordinary knowledge ascription does not entail dogmatism: from this ascription, Expand-p does not allow us to infer that Moore knows that he is not a BIV rather than he has stumps. That is because Expand-p only applies when the entailed proposition is incompatible with the contrast. Since not being a BIV is compatible with having stumps, Expand-p doesn't lead to the implausible consequence that Moore knows that he is not a BIV rather than he has stumps. Indeed, the contrastivist requires every true knowledge ascription to have incompatible content and contrast propositions, as the latter is supposed to be a genuine *contrast* to the content.

[9] This is an over-simplified formulation of epistemic closure. For this conditional to hold, the relevant subject arguably has to know that p entails q, deduce q from p, and thereby come to believe that q in virtue of that deduction. Conditions of this sort are omitted here for two reasons. First, it is still controversial which of them should be added to closure. Second, we can assume that the examples considered in this paper all satisfy these conditions, no matter what they may include. For example, we can assume that Moore believes that having hands entails not being a BIV, he deduces that he is not a BIV from the premise that he has hands, and he comes to believe that he is not a BIV in virtue of the deduction.

[10] Expand-p is only part of the contrastivist version of closure in Schaffer (2007a). There is a parallel principle concerning the contrast position and two other principles covering multi-premise cases:

> CONTRACT-q: If one knows that p rather than q, and if q is entailed by q' and q' is not necessarily false, then she knows that p rather than q'.
> INTERSECT-p: If one knows that p_1 rather than q, and if she also knows that p_2 rather than q, then, *with certain other conditions satisfied*, she knows that $p_1 \wedge p_2$ rather than q. UNION-q: If one knows that p rather than q_1, and if she also knows that p rather than q_2, then, *with certain other conditions satisfied*, she knows that p rather than $q_1 \vee q_2$.

For our purposes, we only have to focus on Expand-p, as one of its counterexamples illustrates that the shiftiness of knowledge ascriptions cannot be eliminated by explicit contrast expressions.

[11] Similar to the binary version of epistemic closure (see fn. 9), Expand-p also requires provisos about how the subject in question comes to believe that p', or, if belief is also contrastive (cf. Blaauw, 2012), how the subject comes to believe that p' rather than q.

For our purposes, we can set aside the issue whether contrastive closure successfully avoids both skepticism and dogmatism. However, a counterexample to Expand-p found in Hughes (2013) well illustrates the main point here: the shiftiness of knowledge ascriptoins does not go away even when explicit contrast expressions are present. To see the problem, suppose one has some ordinary knowledge where the content is p and the contrast is q, where p and q are incompatible. Accordingly, we can design a skeptical proposition: one is deceived by an evil demon into believing that p while in fact it is the case that q. For Moore's ordinary knowledge, (19) is its corresponding skeptical proposition:

(19) Moore is deceived by an evil demon into believing that he has hands while in fact he has stumps.

Let $S(p, q)$ denote the skeptical proposition generated by p and q. As Hughes (2013) shows, we can make use of a proposition like this to form the following counterexample to Expand-p, where (21) (brackets are just for readability), by Expand-p, follows from ordinary ascription (20).

(20) $Kmpq$: Moore knows that he has hands rather than he has stumps.
(21) $Km(p \wedge \neg S(p, q))q$: Moore knows that [he has hands and is not deceived by an evil demon into believing so while he in fact has stumps] rather than he has stumps.[12]

As the content proposition p of $Kmpq$ is incompatible with the skeptical proposition $S(p, q)$, the conjunction $p \wedge \neg S(p, q)$ is equivalent to p. Consequently, Expand-p applies, and (21) follows from (20). But this is problematic, as (21) sounds much less acceptable than (20). [13]

Given examples of this sort, a natural diagnosis is that the truth-conditional interpretation of a knowledge ascription does not merely depend on the propositions expressed by the embedded clauses. It is sensitive to not only what the content proposition is but also the way in which it is expressed. For instance, although p and $p \wedge \neg S(p, q)$ express the same proposition, the two express it in very different ways: while p plainly states that Moore has hands and thus picks out all the hand-possibilities, $p \wedge \neg S(p, q)$ picks out the same set of possibilities by additionally mentioning, and excluding, some skeptical possibilities. Apparently, if mentioning something brings them up to salience, it is natural to understand the examples as showing, again, that knowledge ascriptions, even with explicity contrast expressions, are sensitive to contextual factors, such as what is salient. Namely, the truth-conditional interpretation of a knowledge ascription is shifty, and it is not completely determined by the semantic values of its constituents,

[12] In Hughes (2013, pp. 586–589), what functions as the skeptical proposition is that Moore (or any other subject) is a BIV. I think both counterexamples work well, but the skeptical proposition I choose here better reflects the role it plays: $S(p, q)$ is a proposition that the subject in question cannot differentiate by evidence from q while being incompatible with p.

[13] It is easy to find similar examples which do not rely on skeptical possibilities. Adapting an example from Dretske (1970), starting from the ordinary knowledge **(a)**, allows us to infer **(b)**.

(a) Moore knows that it is a zebra rather than a mule.
(b) Moore knows that it is a zebra and not a cleverly disguised mule rather than a mule.

However, **(b)** is intuitively false (assuming that Moore is no expert in zoology).

even if a third argument is added. Again, this is unexpected from the contrastivist's perspective: if the shiftiness is due to the incompleteness of knowledge ascription sentences, it should no longer exist once all the argument positions, including the contrast, are saturated.

Before leaving this argument, I shall consider a potential response from the contrastivist. She might want to treat the *and-not* construction in (21) as a linguistic device that signals contrast, in the same way that *rather-than* does. If so, the logical form of (21) can be treated along the following line:

(22) $Kmp(S(p,q)\vee q)$: Moore knows that he has hands rather than [(he is deceived by an evil demon into believing so while he in fact has stumps) or (he has stumps)].

If that is the case, then Expand-p does not allow the inference from the ordinary ascription (20) to (22).[14] I have no qualm with this solution itself: If *rather-than* can signal contrasts, why can't *and-not* play the same role? But this move seems to be a misdiagnosis of the problem. The problem is actually not due to the conjunct $\neg S(p,q)$ being added to the p-position, but rather because of, as mentioned above, a skeptical possibility's being brought up in some way to salience. For example, compare the ordinary ascription (20) with the following.

(23) Moore knows that he has hands, which BIVs can't have, rather than stumps.

Unlike the ordinary ascription (20), (23) sounds false, while the only difference between them is the non-restrictive relative clause "which BIVs can't have" present in (23). If the contrastivist responds to (21) by treating *and-not* as a contrast-signaling device, she should also treat the non-restrictive relative clause in (23) in the same way. Namely, she should take the logical form of (23) as, roughly, *Moore knows that he has hands rather than [he is a BIV or he has stumps]*, where the relative clause fills in the contrast position of *know*.

But this cannot be right. A non-restrictive relative clause (NRC) does not compositionally contribute to the semantic value of its main clause. It is a well-known linguistic property that NRCs take "the widest scope". That is, the content of an NRC is not a constituent of the content of its main clause; rather, it is not embedded, as if it were expressed alongside with the content of the main clause. Consequently, when the sentence is uttered, the content of the NRC is not taken into the scope of any embedding environment but rather constitutes part of the speaker's commitment. For example, NRCs escape the scope of intensional operators in the main clause.

(24) Sheila believes that the agency interviewed Chuck, a confirmed psychopath, just after his release from prison. (Potts, 2005, p. 115)

What (24) says does not entail that Sheila believes that Chuck is a confirmed psychopath. By uttering (24), it is the speaker who is committed to that Chuck is a confirmed psychopath. In this sense, NRCs always "scope out", as if it were an independent sentence, making no contribution to the main clause embedding it.

[14] However, as the new contrast $S(p,q)\vee q$ is equivalent to the original contrast q, (22) still follows from (20) by Contract-q (see fn. 10), which is another principle in Schaffer's (2007a). Therefore, this proposal does not work in the contrastivist's way, if she is to defend Schaffer's whole package of closure. Thanks to an anonymous reviewer for this point.

Syntactic and semantic theories of NRCs abound. Based on the widest-scope property of NRCs, most theories share the common view that NRCs are, in some way, separated from the main clause, and thus their semantic values are calculated independently.[15] For example, Potts (2005) develops a multi-dimensional semantics. According to him, the content of a main clause is the at-issue meaning (also known as *what is said*) of the sentence, and what an NRC contributes is a conventional implicature. Both of them are semantic values of the sentence in question, but they are separately calculated and kept as separate dimensions of the sentence meaning. For (23), the two semantic values are:[16]

- At-issue meaning: that Moore knows that he has hands rather than stumps.
- Conventional implicature: that BIVs can't have hands.

Again, this separation is motivated by the widest-scope property of the NRC: that BIVs can't have hands is a commitment of the speaker, and it is not embedded in the scope of the knowledge operator. If this or some similar theory is correct, then the contrastivist's response is linguistically implausible. Recall that the response is that NRCs may fill an argument position of *know*. This directly contradicts the current view that NRCs have the widest scope.

3.4 A context-sensitive contrastivism?

So far we have seen that explicit contrast expressions cannot eliminate the shiftiness of knowledge ascriptions. If this is correct, Saturation must go. In particular, on the contrastivist's view, a binary knowledge acription is shifty because it is incomplete: as the logical form of knowledge ascriptions is ternary, the surface form of the sentence, which is binary, does not determine a unique logical form, and thus the sentence is shifty because the missing argument has to be determined by the context. Saturation captures this view: when the surface form is binary, the contrast has to come from the context; otherwise, when there is an explicit contrast expression, the sentence is complete and thus the contrast is determined by that expression. This is indeed what genuinely ternary relations are like, such as *introduce* and *prefer*. When the surface form is a binary ascription, they are shifty, but when the surface form is ternary, the shiftiness disappears. However, as the examples in this section illustrate, knowledge ascriptions are not like this, as they remain shifty even when they have a ternary surface form.

At this point, the contrastivist might protest that Saturation is too strict. She may insist that *know* is ternary, while claiming that the contrast argument position cannot be saturated by contrast expressions in the surface form but rather is always affected by the context. So, the proposal is that, for a binary ascription, the contrast is determined by the context alone, and for a ternary ascription, the contrast is determined by *both* an

[15] To be clear, there are exceptions where NRCs seem to take a narrow scope, but it is safe to say that they tend to take the widest scope. See Fabricius-Hansen (2020) for a recent survey.

[16] Not all theories of NRCs are multi-dimensional. For example, Schlenker (forthcoming) develops a unidimensional semantics. But in his theory, it is the syntax that serves to separate an NRC and its main clause. The semantic values of the two are calculated independently first, though they are combine to form the unique semantic value of the sentence at the last step.

explicit expression *and* the context. In other words, explicit contrast expressions, such as *rather-than*, never completely settle what the contrasts are—they never saturate the contrast argument. Call this view context-sensitive contrastivism.

This would solve all the above three problems at once. First, in Sect. 3.1, the speaker of (12) can stand by what she said: "Mary doesn't know that it was the rubies that were stolen, so she doesn't know that Peter rather than anyone else stole the rubies." The context-sensitive contrastivist would explain that the first sentence changes the context by making *something else was stolen* a salient contrast. Hence, in interpreting the second sentence, this contrast is taken into account, and as Mary cannot rule it out, the knowledge denial come out true. It is similar for the conjunction case: the first sentence in (13) makes that contrast salient in the context and the interpretation of the second sentence, being sensitive to salient contrasts in the context, is different from when it is uttered alone. Finally, the NRC in (23) may play the same role: it makes salient the possibility that Moore is a BIV, and this contrast is taken into account when interpreting the main clause.

But is this context-sensitive contrastivism a plausible view? I argue that it is not. First, part of the reason why contrastivism appears to be attractive is that it assimilates the shiftiness of binary knowledge ascriptions to a relatively well-understood linguistic model. Specifically, the shiftiness of elliptical expressions is completely unsurprising. For example, a binary ascription of preference is shifty, and that is because the incompleteness of its surface form requires the context to fill in with the missing argument. If the original contrastivism is right, the shiftiness of knowledge ascriptions can be explained by this familiar linguistic model. Whatever account correctly explains ellipses, it automatically applies to knowledge ascriptions. However, context-sensitive contrastivism loses this appeal. According to this view, there is no parallel between knowledge ascriptions and, say, preference ascriptions: the former is context-sensitive even when it has a complete, ternary surface form, whereas the latter is context-sensitive only if its surface form is incomplete.

Given that the ellipsis model is not applicable, the context-sensitive contrastivist needs a new theory to explain why knowledge ascriptions, whether binary or ternary, are shifty. According to context-sensitive contrastivism, the context always plays a role in determining what the contrast is for a knowledge ascription, even if there is an explicit contrast expression. This view, however, is strange. On the one hand, the view would use some contextualist apparatus to explain how knowledge ascriptions are context-sensitive and how the relevant contextual parameters can be affected by surrounding utterances, as in (13), and NRCs, as in (23). On the other hand, when such a contextualist apparatus is employed, there is a natural explanation of the contrastive data that is does not appeal to a third argument place of "to know". I argue that this undermines the motivation of stipulating a ternary logical form of knowledge ascriptions.

To be more specific, the context-sensitive contrastivist will appeal to a contextualist apparatus, roughly like the following, to explain (13) and (23), repeated in (25) and (26).

(25) Mary doesn't know that Peter stole the rubies rather than anything else. But she knows that Peter rather than anyone else stole the rubies.

Explanation: "rather than anything else" in the first sentence, in some way, brings about certain changes in the context, which in turn, in some way, affects the interpretation of the second sentence, making it come out false.

(26) Moore knows that he has hands, which BIVs can't have, rather than stumps.
Explanation: "which BIVs can't have", in some way, brings about certain changes in the context, which in turn, in some way, affects the interpretation of this very sentence, making it come out false.

If such a contextualist apparatus is available, I see no reason why it cannot apply to the contrastive data and explain the truth-conditional effect of contrast expressions. For example, for (5b), repeated in (27), the explanation can go:

(27) Mary knows that Peter stole the rubies rather than anything else.
Explanation: "rather than anything else", in some way, brings about certain changes in the context, which in turn, in some way, affects the interpretation of this very sentence, making it come out false.

Notice that according to this explanation, the *rather-than* expression is not part of the logical form the knowledge ascription. Instead, the role it plays is just to affect the context. In other words, if such an explanation is available, the logical form of (27) can be treated as binary—i.e. *Mary knows that Peter stole the rubies*—and this binary form is complete. What the *rather-than* expression does is not to fill an argument place in the logical form of the ascription, as there is no place to fill. Rather, it is only an adjunct which is able to change the context in some way, much like what an NRC does.

To summarize, context-sensitive contrastivism is an unstable view. On the one hand, in order to explain examples raised in this section, it has to abandon Saturation and appeal to some contextualist apparatus. On the other hand, once equipped with such a contextualist apparatus, ternicity in the logical form of knowledge acriptions becomes explanatorily redundant, as the core empirical evidence that putatively supports contrastivism, i.e. the contrastive data, can be easily explained by the contextualist apparatus within a binary framework. Thus, it seems that context-sensitive contrastivism slips to good old contextualism, according to which (i) the logical form of knowledge ascriptions is binary and (ii) the interpretation of a knowledge ascription is sensitive to the context, which can be affected by expressions like *rather-than* and NRCs. In the next section, I will develop such a contextualist theory, which explains both the contrastive data and the other examples raised in this section.

4 A contextualist account: semantics and pragmatics

Before expositing the contextualist account, let me make clear some success conditions for it. First, according to such a theory, a knowledge ascription of a binary surface form is taken as syntactically complete. Thus, to explain the shiftiness of binary ascriptions is to give an account of what contextual parameter is responsible for affecting the truth-condition of binary ascription, and how the value of this contextual parameter is determined. Second, such an account should accommodate the contrastive data that are alleged to favor the contrastivist view. In other words, an adequate binary account has

to explain why explicit contrast-signaling expressions can have some truth-conditional effect on knowledge ascription. Third, as demonstrated in Sect. 3, a correct account should allow ternary ascriptions to be shifty.

Turning to the promised contextualist account, let's start with a familiar idea: an epistemic agent has a body of evidence—including sensory experience, memory, etc.—that rules out some possibilities, i.e. the possibilities in which she doesn't have the body of evidence that she actually has. Let each possibility be represented by a possible world, and call the set of all the possible worlds that are not ruled out by an agent's evidence her *epistemic state*, and let E_a denote the epistemic state of agent a. The core doctrine of contextualism is that the truth of ascribing p as knowledge to an agent does not require the agent's epistemic state to rule out *all* the not-p possibilities. Rather, what possibilities she has to rule out in order for the truth of the knowledge ascription depends on the context. Thus, in a general contextualist framework, an utterance of a knowledge ascription sentence is interpreted with respect to a contextually determined *range of relevant (in some sense) possibilities*, and a knowledge ascription of p can be truthfully uttered in the context just in case all the not-p possibilities that are *within* the range of relevant possibilities are ruled out by the agent in question. Let a set of possible worlds R_c represent the set of *r*elevant possibilities determined by context c. Hence the truth condition:

> CONTEXTUALIST SEMANTICS: In a context c, an utterance of "a knows that p" is true if and only if p is true at all the possible worlds in $R_c \cap E_a$.

Gloss: "a knows that p" can be truthfully uttered in a context just in case a's epistemic state rules out all the relevant not-p possibilities.[17]

But what possibilities are relevant in a context? In general, it is governed by some pragmatic constraints; but among different contextualists, there is no agreement on what these constraints should be. To mention some, for Lewis (1996), R_c is governed by multiple rules, one of which is *Rule of Attention*: if relevant speakers in a conversation attend to some possibilities, then those possibilities must be included in R_c. For Cohen (1999), possibilities in R_c are those which are salient in the context c. And for Blome-Tillmann (2009), they are the possibilities that are compatible with the relevant speakers' presuppositions.

We don't have to decide between these particular contextualist theories. My aim is to point to an aspect of context that at least partially constrains what R_c should be like. It by no means constitutes a complete account of how R_c is determined, and the aspect of context that I am suggesting might well be reducible to some contextual parameter mentioned above.

This aspect of context is a question that speakers intend to see if an agent, based on her epistemic state, could answer and if so, what the answer would be. When a knowledge ascription sentence is uttered, the utterer in so doing typically intends to inform the hearer something about the subject. But what exactly is that piece of

[17] A serious contextualist semantics must be much more complex than this. For example, the current semantics doesn't even entail the factivity of knowledge, nor the belief condition. These are, of course, a shortcoming, but it does not matter for our purposes here. Indeed, we can cast these conditions as semantic constraints on the epistemic state: the actual world must be a member of the epistemic state, and the epistemic state must be a subset of the agent's doxastic state.

information? The proposal is this: in saying that a subject knows that p, the utterer is informing the hearer that if the subject is asked a certain question, and if she sincerely responds to it, she will answer, truthfully and based on her body of epistemic state, that p. The contention here is that such a question is an aspect in the context when knowledge is being ascribed. To illustrate, consider the scenario of [Jewel Thief] again. The relevant speakers may be interested in finding out how Mary's epistemic state fares with the question *Who stole the rubies?*. In particular, they are interested in whether Mary's epistemic state provides a true answer to the question, and if so, what the answer is. For convenience, call such a question *Epistemic Question*. Given its contextually determined nature, let's assume that it is an aspect of context and note such an epistemic question in a context c as Q_c.

Following the semantics of questions originated in Hamblin (1973), an epistemic question Q_c is formally represented as a set of (at least two) propositions, i.e. a set of sets of possible worlds. Roughly speaking, a question is defined by its possible direct answers, each of which is a proposition. For example, if we consider the epistemic question *Who stole the rubies?*, it corresponds to the set of all the propositions of the form "x stole the rubies," where x is any individual. For a polar question such as *Did Peter steal the rubies?*, it is formally treated as a set of two members: the proposition that Peter stole the rubies, and its negation.

By formulating a primitive contextualist semantics and two contextually determined factors, i.e. the set R_c of relevant possibilities and the epistemic question Q_c, the ground has been prepared for a pragmatic theory that explains the context-sensitivity of knowledge ascription. In what follows, I will offer some pragmatic rules governing the interaction between these factors and utterances of knowledge ascription sentences.

First, if p is ascribed as knowledge to a subject in a context, the speaker is attributing p as the subject's truthful and evidence-based answer to the epistemic question in the context. Given this question-answer relation, the embedded *that*-clause and the epistemic question should be subject to the general pragmatic constraint of question-answer congruence. For a question, a sentence or clause may be suitable for answering it while another may not, even if they are equivalent in terms of their propositional contents. Congruence is a relation about this kind of suitability: if a sentence or clause is suitable for answering a quesiton, then it is congruent with the question; otherwise it isn't. As will be clear later, this notion of question-answer congruence allows the interpretation of a knowledge ascription to shift according to contrast expressions. Very roughly, a contrast expression in a knowledge ascription signals what epistemic question it is congruent with, and the interpreter can assume that the ascription is suitably addressing such an epistemic question. This process of resolving the epistemic question will in turn affect the interpretation of the knowledge ascription.

Congruence is regulated by certain linguistic devices, such as focus. For instance, if in a context the epistemic question is *Who stole the rubies?*, then the embedded clause in (8b)—repeated as (28)—isn't congruent to it, given the focus on "the rubies" it involves. Because of this incongruence, (28) cannot be felicitously uttered in such a context. In contrast, an utterance of (29) is felicitous in the context, as the embedded clause in it is a congruent answer to the epistemic question.

(28) Mary knows that Peter stole THE RUBIES.

(29) Mary knows that PETER stole the rubies.

In order to capture this congruence relation, I follow Roberts (2012) to associate each sentential clause with a set of alternative propositions, which is called the CONGRUENT QUESTION of the clause.[18] Then, a sentence or clause is congruent with a question just in case the latter is the Congruence Question of the former.

> CONGRUENT QUESTION: For a sentential clause that involves a focused component, its congruent question is a set of propositions that are derived by substituting semantic entities of an appropriate type for the focused constituent.[19]

For example, the congruent question of the clause *PETER stole the rubies* is the set of propositions of the form *x stole the rubies*, where x is any individual. In contrast, the congruent question of *Peter stole THE RUBIES* is the set of propositions of the form *Peter stole x*, where x is any object or plurality of objects.

As the embedded clause in a knowledge ascription is supposed to be congruent with the epistemic question in the context, here comes a pragmatic constraint on how they should relate to each other:

> RULE OF CONGRUENCE: A knowledge ascription "*a* knows that *p*" is uttered in a context *c* felicitously only if the congruent question of the clause *p* is identical to the epistemic question Q_c.

In other words, this constraint requires an utterance of a knowledge ascription to be addressing, with its embedded *that*-clause, the epistemic question in context. Applying it to the example above, we can see that the constraint is violated by (28) in the intended context, while (29) complies with it.

I see no reason why focus is the only linguistic device that governs congruent questions. Indeed, many others can play the same role in constraining what question can be answered by a clause or sentence, including *it*-cleft, *rather-than* construction, etc.. For example, *It was Peter who stole the rubies* is congruent with *Who stole the rubies?* and *It was the rubies that Peter stole* is congruent with *What did Peter steal?*. Moreover, *rather-than* constructions are congruent with questions with two possible answers: For S_1 *rather than* S_2, its congruent question is a set of two propositions expressed by S_1 and S_2, i.e. *Which of* S_1 *and* S_2 *is true?*.[20] In addition, if a clause does not involve any of those devices governing the congruent question, I assume

[18] This notion is in line with a well-established view on the semantic interpretation induced by focus—focus triggers a set of alternative propositions. In this sense, the notion of congruent question is the same as "focus semantic value" in Rooth (1985, 1992), the set of "focal alternatives" in Roberts (2012), and "current question" in Simons et al. (2017), etc..

[19] This notion of Congruent Question is slightly different from Roberts's view. For Roberts, congruent questions are only derived at the whole sentence level and do not apply to embedded clauses. But the move I'm making here is a natural extension of her approach, given that the embedded clause in a knowledge ascription, as illustrated above, should be congruent with the epistemic question in the context in the same way as whole sentences should be congruent with the contextually determined question under discussion.

[20] A compositional derivation of sets of alternative propositions (i.e. congruent questions) can be found in Abusch (2010). He treats the *and-not* construction (e.g. *John is in Boston and not New York*) as triggering a set of alternative propositions (i.e. a congruent question) and doesn't mention *rather-than*, but it strikes me that the two are nothing different.

that by default its congruent question is a polar question, i.e. the set consisting of the proposition expressed by the clause, and its negation.

So much for this rule of Congruence. Now we turn to how, in a context c, the epistemic question Q_c should be related to the set R_c of relevant possibilities, possibilities that matter for knowledge ascriptions. The general idea is that, for a given set of relevant possibilities, some epistemic question can be asked while others are pragmatically deviant. To illustrate, a basic constraint of this sort is NON- TRIVIALITY, which requires that the epistemic question Q_c not be settled by R_c in a context c, because otherwise, it would be implausible for the speakers to *wonder* what the subject's answer to Q_c would be:

> RULE OF NON- TRIVIALITY: In a context c, R_c is compatible with at least two independent answers to Q_c.[21]

This pragmatic rule, together with Congruence, gives rise to the accommodation effect noticed in Lewis (1979): even in an ordinary context where no skeptical possibilities are in R_c, once a knowledge ascription like (30) is uttered, some skeptical possibilities become relevant and thus are accommodated into R_c.

(30) Moore knows that he is not a BIV.

Skeptical possibilities are accommodated because otherwise either Congruence or Non-Triviality would be violated. The congruent question of the embedded clause in (30) is a polar question: *Is Moore a BIV?*, and according to Congruence, this question must be the epistemic question in the context. But if there were no BIV-possibilities in R_c, then R_c itself would settle the question, with the answer that he is not a BIV. Then, some BIV-possibilities have to be accommodated into R_c, which makes the utterance of (30) false, since Moore's epistemic state can't rule out those skeptical possibilities.

The above rule requires R_c not to be too strong: it must not settle Q_c so that Q_c becomes trivial. However, for the purpose of explaining the contrastive data, what is important is another pragmatic rule that is intended to ban the opposite extreme: R_c must not be too weak so that Q_c becomes too substantive as a question. To illustrate, consider a rather philosophical, or skeptic-friendly, context c, in which speakers consider some skeptical possibilities as relevant, as well as some ordinary possibilities, intending to assess Moore's epistemic state with respect to *all* these possibilities. That means that R_c in this case involves both ordinary and skeptical possibilities. Let's also suppose, for *reductio*, that the speakers are currently interested in finding out what truthful and evidence-based answer Moore would give, if any, to the question *Does Moore have hands or stumps?*. In this case, Q_c is a question that has only two possible answers, that Moore has hands and that Moore has stumps. But now the context, set up as such, is strange. If some skeptical possibilities are considered relevant, how can the speakers put forth this epistemic question, which by itself rules out all skeptical possibilities? In other words, all the possible answers to Q_c will rule out the skeptical possibilities, which makes it implausible to entertain any skeptical possibility as relevant in the first place. Hence the constraint: the possible answers to Q_c must exhaust all the possibilities in R_c; given that the possible answers and R_c are sets of possible worlds, this means that all these answers must jointly "cover" R_c.

[21] Formally: $\exists p_1, p_2 \in Q_c \exists w_1, w_2 \in R_c[w_1 \neq w_2 \wedge (w_1 \in p_1 \wedge w_1 \notin p_2) \wedge (w_2 \in p_2 \wedge w_2 \notin p_1)]$.

RULE OF COVERING: In a context c, every possibility in R_c is committed to the truth of some possible answers to Q_c.[22]

As we have seen, the Rule of Non-Triviality enables pragmatic accommodation that makes R_c expand. Now, this Covering may achieve the opposite effect: when in a context c there are some possibilities in R_c in which every possible answer to Q_c is false, such possibilities are to be excluded from R_c, so that Covering won't be violated. This rule, then, enables a kind of accommodation that in effect makes R_c shrink.

By now I have introduced a general contextualist semantics and some pragmatic constraints, and these are all we need to form an account that satisfies the desiderata mentioned at the beginning of this section. First, according to Contextualist Semantics, the context-sensitivity of knowledge ascriptions is due to the contextual parameter R_c. Although we don't have a full analysis of how R_c is determined, the current theory gives a partial characterization of how it is constrained: it has to relate to Q_c in a way that complies with Non-triviality and Covering.

As to the second desideratum, the contextualist semantics combined with the pragmatic rules offers an account of the contrastive data. For example, we have to explain why the pair of sentences in (17), repeated as follows, appear to have different truth-values in the given context.

(31a) Mary knows that Peter rather than anyone else stole the rubies.
(31b) # Mary knows that Peter stole the rubies rather than anything else.

By Congruence, for (31a) to be felicitously uttered, Q_c has to be *Was it Peter or anyone else who stole the rubies?*. But for (31b), Q_c has to be *Did Peter steal the rubies or anything else?*. Hence, by Covering, the two utterances then impose different requirements on the contextual parameter R_c. On the one hand, R_c in the context for (31a) must not involve any possibilities where the rubies weren't stolen. In other words, in a context in which (31a) can be felicitously uttered, all the relevant possibilities in its R_c are someone-stole-rubies possibilities. On the other hand, the R_c parameter in the context for (31b) isn't constrained in such a way: it can involve some possibilities where the rubies weren't stolen.

This explains why the pair appear to have different truth-values. Upon hearing (31a), the hearer tends to accommodate the felicity of the utterance by making the pragmatic rules satisfied, including Congruence and Covering. The result is that all the relevant possibilities in the context are someone-stole-rubies possibilities. Among those, Mary's epistemic state is only compatible with the Peter-stole-rubies possibilities. Therefore, according to Contextualist Semantics, (31a) comes out true. In contrast, the felicity of (31b) doesn't require all the relevant possibilities in the context to be someone-stole-rubies possibilities. Rather, there might well be possibilities where what was stolen was something else. As Mary's epistemic states cannot rule out these possibilities, (31b) comes out false.

Now it is clear why *rather-than* can have a truth-conditional effect on knowledge ascriptions. Other contrast-signaling devices, e.g. focus and *it*-cleft, can be explained in the same way, so I leave the details aside. What is remaining in the contrastive data is the case where different topical questions in the context appear to affect the truth-condition

[22] Formally: $\forall w \in R_c \exists p \in Q_c (w \in p)$.

of knowledge ascription. In the contexts [C1] and [C2], the two questions are *Who stole the rubies?* and *What did Peter steal?*, respectively. I submit that these questions are the epistemic questions in the two respective contexts, because the relevant speakers are wondering whether and how Mary could truthfully answer the questions based on the evidence she has.[23] Hence, Covering, in the same manner presented above, explains how the R_c parameter has to be adjusted according to these two different epistemic questions, which in turn results in the truth-value difference of "Mary knows that Peter stole the rubies" when uttered in the respective contexts.

The current account also satisfies the third desideratum: it explains the examples in Section 3, those which demonstrate that ternary ascriptions can be shifty. To illustrate, consider (12), repeated as (32).

(32) Mary doesn't know that Peter rather than anyone else stole the rubies.

It is the negation of a knowledge ascription which comes out true in the contrastive data. However, given the contextualist semantics, the contents of the ascription and its denial both crucially depend on how the contextual parameter R_c is resolved. We have seen one resolution that makes it exclude all the possibilities where the stolen jewel was not the rubies. Based on that resolution of R_c, (32) is false. But note that this specific resolution is derived by the interpreter's inference. In other words, in order to accommodate felicity, the interpreter infers that the R_c should be restricted as such, so (32) turns out false. Such an inference determines a content of (32) based on the assumption that the utterance of (32) is complying with pragmatic rules such as Covering. However, other interpretations are possible as well, when there are pragmatic pressures conflicting with Covering. In this case, consider the speaker's defense: "Mary doesn't know that it was the rubies that were stolen, so she doesn't know that Peter rather than anyone else stole the rubies." What is going on here is that uttering the first sentence strongly suggests that some possibilities where something else was stole should be relevant. Given that, (32) is interpreted as true. This explains why, when challenged, the utterer of (32) can stand by what she said: what she does by refusing to retract the utterance is to make salient that alternative resolutions of R_c on which (32) is true. On the other hand, an implication of this account is that after the speaker stands by and doesn't retract what she said, the felicity of the original utterance is sacrificed, as Covering is violated. That does seem to be the case. Although the speaker can stand by what she said, the interlocutor would then get an uncomfortable feeling: "Now I see what she meant, and that's indeed not false. But it's a really strange way of saying it!"

[23] Schaffer and Knobe (2012) take the two questions as questions under discussion. I beg to differ: the knowledge ascription "Mary knows that Peter stole the rubies" is not a congruent answer to either question, in the sense of Roberts (2012, pp. 31–32). Although before the knowledge ascription is uttered, they seem to be the questions under discussion, because of the incongruence of the ascription to it, some other questions must be accommodated as the immediate questions under discussion, say, *Who was the person that stole the rubies, according to what Mary knows?* and *What did Peter steal, according to what Mary knows?*. But given these newly accommodated questions under discussion, the previous two questions now become the epistemic questions respectively in the two contexts.

5 Conclusion

With the three desiderata satisfied, I conclude that we should resist the contrastivist semantic revolution, which adds a contrast argument to the verb "to know". On the one hand, the contextualist semantics, combined with the above pragmatic rules, can successfully satisfy the theoretical and empirical desiderata. Thus, it is unnecessary to embrace a contrastivist approach. On the other hand, as I argued in Sections 3, the contrastivist semantics itself lacks empirical adequacy. Thus, contrastivism is insufficient regarding what initially motivates it. Given that it is neither necessary nor sufficient, a contrastivist revolution, touted as a better alternative to contextualism, does not succeed.

Acknowledgements I am grateful to two anonymous referees of this journal for their helpful comments on earlier versions of this paper. I also wish to thank Jordan Bell, I-Sen Chen, Rohan French, Natasha Haddal, Hanti Lin, Patrick Skeels, S. Kaan Tabakci, and other members of UC Davis LLEMMMa reading group. Special thanks to Adam Sennet for all the valuable guidance and support.

Funding Not funded.

Data availability Not applicable.

Code availability Not applicable.

Declarations

Conflict of interest The author declares that he has no conflict of interest.

References

Abusch, D. (2010). Presupposition triggering from alternatives. *Journal of Semantics, 27*(1), 37–80.
Aloni, M., & Égré, P. (2010). Alternative questions and knowledge attributions. *Philosophical Quarterly, 60*(238), 1–27.
Bach, K. (2005). The emperor's new 'knows'. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth* (pp. 51–89). Oxford University Press.
Baumann, P. (2016). *Epistemic contextualism: A defense*. Oxford University Press.
Blaauw, M. (2012). Contrastive belief. In M. Blaauw (Ed.), *Contrastivism in philosophy: New perspectives.* Routledge.
Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind, 118*(470), 241–294.
Brogaard, B. (2009). What Mary did yesterday: Reflections on knowledge-WH. *Philosophy and Phenomenological Research, 78*(2), 439–467.
Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives, 33*, 57–89.
Dretske, F. I. (1970). Epistemic operators. *Journal of Philosophy, 67*(24), 1007–1023.
Dretske, F. I. (1972). Contrastive statements. *Philosophical Review, 81*(4), 411–437.
Fabricius-Hansen, C. (2020). (Non)restrictive nominal modification. In D. Gutzmann (Ed.), *The Wiley Blackwell companion to semantics* (pp. 1–32). Wiley.
Gerken, M., & Beebe, J. R. (2016). Knowledge in and out of contrast. *Noûs, 50*(1), 133–164.
Hamblin, C. L. (1973). Questions in Montague English. *Foundations of Language, 10*(1), 41–53.
Hughes, M. (2013). Problems for contrastive closure: Resolved and regained. *Philosophical Studies, 163*(3), 577–590.
Kallestrup, J. (2009). Knowledge-WH and the problem of convergent knowledge. *Philosophy and Phenomenological Research, 78*(2), 468–476.

Karjalainen, A., & Morton, A. (2003). Contrastive knowledge. *Philosophical Explorations, 6*(2), 74–89.
Lewis, D. K. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic, 8*(1), 339–359.
Lewis, D. K. (1996). Elusive knowledge. *Australasian Journal of Philosophy, 74*(4), 549–567.
Montminy, M. (2008). Can contextualists maintain neutrality? *Philosophers' Imprint, 8*, 1–13.
Morton, A. (2012). Contrastive knowledge. In M. Blaauw (Ed.), *Contrastivism in philosophy* (pp. 101–115). Routledge.
Potts, C. (2005). *The logic of conventional implicatures*. Oxford University Press.
Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics, 5*, 1–69.
Rooth, M. (1985). *Association with Focus*. PhD thesis, Dept. of Linguistics, University of Massachusetts, Amherst.
Rooth, M. (1992). A theory of focus interpretation. *Natural Language Semantics, 1*(1), 75–116.
Schaffer, J. (2004). From contextualism to contrastivism. *Philosophical Studies, 119*(1/2), 73–103.
Schaffer, J. (2005). Contrastive knowledge. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 1, pp. 235–271). Oxford University Press.
Schaffer, J. (2007). Closure, contrast, and answer. *Philosophical Studies, 133*(2), 233–255.
Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research, 75*(2), 383–403.
Schaffer, J. (2008). The contrast-sensitivity of knowledge ascriptions. *Social Epistemology, 22*(3), 235–245.
Schaffer, J. (2009). Knowing the answer redux. *Philosophy and Phenomenological Research, 78*(2), 477–500.
Schaffer, J., & Knobe, J. (2012). Contrastive knowledge surveyed. *Noûs, 46*(4), 675–708.
Schlenker, P. (forthcoming). Supplements without Bidimensionalism. *Linguistic Inquiry*.
Simons, M., Beaver, D., Roberts, C., & Tonhauser, J. (2017). The best question: Explaining the projection behavior of factives. *Discourse Processes, 54*(3), 187–206.
Sinnott-Armstrong, W. (2004). Classy pyrrhonism. In W. Sinnott-Armstrong (Ed.), *Pyrrhonian skepticism* (pp. 188–207). Oxford University Press.
Sinnott-Armstrong, W. (2008). A contrastivist manifesto. *Social Epistemology, 22*(3), 257–270.
Steglich-Petersen, A. (2014). Knowing the answer to a loaded question. *Theoria, 80*(1), 97–125.