

Knowing whether A or B

Maria Aloni · Paul Égré · Tikitou de Jager

Received: 5 December 2008 / Accepted: 17 July 2009 / Published online: 4 September 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract The paper examines the logic and semantics of knowledge attributions of the form “s knows whether A or B”. We analyze these constructions in an epistemic logic with alternative questions, and propose an account of the context-sensitivity of the corresponding sentences and of their presuppositions.

Keywords Knowledge · Alternative questions · Contextualism · Awareness · Relevant alternatives · Attention · Presupposition · Accommodation

1 Introduction

The aim of this paper is to discuss the semantics of knowledge attributions of the form “s knows whether A or B”, where the complement “whether A or B” expresses an alternative question, as in “Bob knows whether Mary is French or Italian”. Our goal is to provide an account of the logic of these epistemic constructions, and of the context-sensitivity of the corresponding sentences.

It is standard in linguistic theory to distinguish the polar and alternative readings of disjunctive questions, e.g. (von Stechow 1991). Under the polar reading, a direct question of the form “Is Mary French or Italian?” calls for a yes or no answer. The

M. Aloni (✉) · T. de Jager (✉)
ILLC, Amsterdam, The Netherlands
e-mail: M.D.Aloni@uva.nl

T. de Jager
e-mail: S.T.deJager@uva.nl

P. Égré (✉)
Institut Jean-Nicod (ENS/EHESS/CNRS),
Département d’Etudes Cognitives de l’ENS, Paris, France
e-mail: paulegre@gmail.com

polar reading can be forced in English by asking “Is Mary either French or Italian?”. For the alternative reading, by contrast, the question cannot be answered by yes or no and has to be answered by a sentence like “Mary is French”, or “Mary is Italian”, namely by providing information about the truth and falsity of the respective disjuncts.

There is still some debate in the literature about the answerhood conditions of alternative questions, and in consequence, about the conditions under which a subject can be said to know whether A or B. In a recent paper, Schaffer (2007) has argued that in a context in which a subject *s* sees someone on TV, who is actually George Bush, but such that *s* is not able to discriminate between Bush and Will Ferrell (because Ferrell is a very good Bush impersonator), and yet is able to see that it is not Janet Jackson, (1-a) below should be judged false, but (1-b) should count as true:

- (1) a. *s* knows whether George Bush or Will Ferrell is on TV.
 b. *s* knows whether George Bush or Janet Jackson is on TV.

The intuitive reason for the truth of (1-b), according to Schaffer, is that the question “Is Bush or Janet Jackson on TV?” is easier for *s* to answer than the question “Is Bush or Will Ferrell on TV?”. In our view, however, although (1-a) should be incontrovertibly false in the scenario, ordinary intuitions about (1-b) are less stable. In our opinion, all that *s* really knows is that *Janet Jackson is not on TV*, which may not be sufficient to answer the question “Is Bush or Jackson on TV?”.

More formally, assuming the partition theory of questions of Groenendijk and Stokhof (1984), an answer of the form “Janet Jackson is not on TV” counts only as a *partial answer* to the question “Is Bush or Janet Jackson on TV?”. For *s* to know the complete answer he should know more, namely also that Bush is on TV. The partial answer “Janet Jackson is not on TV” would count as complete if one presupposed that exactly one of the two disjuncts had to be true. In principle, however, there is no more reason to think that “*s* knows whether Bush or Janet Jackson is on TV” is true than there is to think that “*s* knows whether Ferrell or Janet Jackson is on TV” is true in this scenario. In other words, *s*’s ignorance about who exactly *is* on TV seems to override *s*’s partial knowledge about who is *not* on TV.

Despite this, we agree with Schaffer that there is a sense in which, if *s* is allowed to *ignore* the possibility that Ferrell might be on TV, then *s* can be said to know whether Bush or Janet Jackson is. In a previous paper (Aloni and Égré 2008) we discussed several ways of making sense of Schaffer’s contextualist idea, and offered to clarify the meaning of “knowing whether” constructions. Two issues were left for further investigation. The first concerns the grounds on which an agent is allowed to ignore possibilities (from an internal perspective, and from the perspective of an external ascriber). In this paper we propose to relate the status of “knowing whether” sentences to *attention to possibilities*, to clarify in which cases an agent who is not attending to a possibility can still safely ignore that possibility, and be ascribed knowledge (Lewis 1996). A second, more linguistic issue that we only partially investigated in our previous work concerns the nature of the *presuppositions* that go with the use of alternative questions under the verb *know*.

We examine both these issues in turn in the paper. Before doing that, we lay out the main elements we need in order to analyze “knowing whether” constructions in epistemic logic.

2 Knowing whether

Logicians are familiar with the usual paraphrase in epistemic logic of expressions of the form “Bob knows whether (or not) Mary is French” in terms of “knowing that” and disjunction, namely as “Bob knows that Mary is French or Bob knows that Mary is not French”:

- (2) a. Bob knows whether (or not) Mary is French.
 b. $Ka \vee K\neg a$

When considering an expression of the form “Bob knows whether Mary is French or Italian”, however, the situation is more complicated than it may seem at first. The reason is that such a sentence is typically ambiguous. One way of paraphrasing the sentence, by analogy with the previous case, is as meaning: “Bob knows that Mary is French or Italian, or Bob knows that it is not the case that Mary is French or Italian”.

- (3) a. Bob knows whether (or not) Mary is French or Italian.
 b. $K(a \vee b) \vee K\neg(a \vee b)$

This paraphrase corresponds to what is known as the *polar reading* of embedded disjunctive questions. Bob knows whether Mary is French or Italian in that sense if Bob can answer by “Yes” or “No” to the direct question “Is Mary French or Italian?”. If pressed to answer what Mary’s nationality is, however, Bob may only be able to respond “French or Italian”, without knowing *which*.

A second, distinct reading for these disjunctive questions is what is known as the *alternative reading*. The expected answer for the direct disjunctive question in this case is not “Yes” or “No”, but will typically be “She is French” or “She is Italian”. In the embedded case, “whether” can no longer be replaced by “whether or not” to express alternative questions, as “or not” forces the polar reading (Larson 1985). Surprisingly, analyses diverge on what the logical form of “Bob knows whether Mary is French or Italian” ought to be in the alternative reading.

One obvious first candidate is the plain disjunction “Bob knows that Mary is French or Bob knows that Mary is Italian”:

- (4) a. Bob knows whether Mary is French or Italian.
 b. $Ka \vee Kb$

This analysis requires that Bob be able to respond to the question by giving at least one of the answers “French” or “Italian”. This account seems inadequate, however. For instance, suppose that Mary is neither French nor Italian, but German, and Bob knows it. Then what Bob knows can be represented by the formula $K(\neg a \wedge \neg b)$. In that case, (4-b) will be false. Intuitively, however, we would not want to say that (4-a) is false, given that he can give the correct and maximally informative answer “Neither”.

A dual problem arises if Mary happens to be both French and Italian. Let us assume that Bob only knows Mary to be French, but fails to realize she is also Italian. He would

answer “French” to “Is Mary French or Italian?”. In such a situation, (4-b) will indeed be true, but it is not obvious that (4-a) should count as true. Although Bob’s answer is not wrong in that case, it would still be *incomplete* with respect to the situation under discussion, and possibly incorrect in other cases. Indeed, imagine that Bob knows Mary is French but wrongly believes that she is not Italian: (4-b) will still count as true, but then we would definitely be reluctant to say that Bob knows whether Mary is French or Italian if his complete answer was “She is French and not Italian”.

For these two reasons, our favored analysis of alternative questions under *know* follows the lead of the partitionial analysis of questions proposed more generally by Groenendijk and Stokhof (1984). To know whether Mary is French or Italian, in their framework, means to be able to respond with one of the answers “French (and not Italian)”, “Italian (and not French)”, “Both” or “Neither”, whichever of these options is the true one. Thus, our desired paraphrase for embedded alternative questions is:

- (5) a. Bob knows whether Mary is French or Italian.
 b. $K(a \wedge \neg b) \vee K(\neg a \wedge b) \vee K(a \wedge b) \vee K(\neg a \wedge \neg b)$

One of the features of this analysis of embedded alternative questions is that it is truth-conditionally equivalent to $(Ka \vee K\neg a) \wedge (Kb \vee K\neg b)$ (assuming the standard Kripke semantics for the *K* operator, see Sect. 2). In other words, the analysis predicts that knowing whether *a* or *b* is truth-conditionally equivalent to knowing whether *a* and knowing whether *b*. Thus, Bob knows whether Mary is French or Italian if and only if Bob can correctly answer each of the polar questions “Is Mary French?” and “Is Mary Italian?”. As the reader can check, this analysis avoids both of the problems we raised for the analysis (4-b).

Before considering further examples, we should mention a third candidate analysis for “knowing whether *a* or *b*” under the alternative reading. A common interpretation of “Bob knows whether Mary is French or Italian” is “Bob knows which of the two nationalities ‘French’ and ‘Italian’ Mary has, *assuming she has only one of them.*” This can be expressed as:

- (6) a. Bob knows whether Mary is French or Italian.
 b. $K(a \wedge \neg b) \vee K(b \wedge \neg a)$

As the reader can check, this analysis can be seen as a strengthening of (4-b), in which the two alternatives *a* and *b* are taken to be mutually exclusive. It is also a strengthening of the partitionial analysis in (6-b). Unlike (4-b), but like (5-b), this analysis entails that if you know whether *a* or *b*, you know whether *a* and you know whether *b*. But again the objections we raised against (4-b) can be raised in principle against (6-b). If Bob knows that Mary is both Italian and French, then (6-b) is simply false, but logically speaking we would not want to say that Bob does not know whether Mary is Italian or French in that case.

The intuitive appeal behind the strengthened truth-conditions in (6-b) may be due to the fact that, in ordinary language, we take an alternative question to implicate or to presuppose that exactly one of the alternatives is true. Thus, someone who asks “Is Mary French or Italian?” generally assumes that Mary has at least one and at most one of the two nationalities. On our account, this fact relates to the use and pragmatics of alternative questions, and should not directly be part of their semantics, nor of the

truth conditions of “knowing whether” constructions. We say more about the status of these assumptions in Sects. 4 and 6 below.

For the sake of completeness, two further candidates might be considered for the analysis of “knowing whether a or b ”. These are intermediate between analysis (4-b) and analysis (6-b), and are therefore asymmetrical with regard to the space of possible answers:

- (7) a. $K(a \wedge \neg b) \vee K(\neg a \wedge b) \vee K(\neg a \wedge \neg b)$
 b. $K(a \wedge \neg b) \vee K(\neg a \wedge b) \vee K(a \wedge b)$

Analysis (7-a) corresponds to the idea that “knowing whether a or b ” implies to know that at most one of the alternatives is true, while (7-b) implies to know that at least one of the alternatives should be true. Some arguments have been made in the literature on questions that are compatible with an analysis like (7-a) in particular, based on the consideration of a possible asymmetry between the “Both” answer and the “Neither” answer to an alternative question (see Higginbotham 1991). However, our reasons to favor (5-b) over (4-b) or (6-b) simply carry over to those cases, because each of these paraphrases excludes either the “Both” answer or the “Neither” answer. As it turns out, this objection would not hold against a weaker variant of (7-a) such as $Ka \vee Kb \vee K(\neg a \wedge \neg b)$, nor against a weaker variant of (7-b) such as $Ka \vee Kb \vee K(a \wedge b)$. Both of those can be seen to weaken analysis (4-b). But such candidates would remain problematic, on the grounds of incomplete information we already explained: for they would predict that “Bob knows whether Mary is French or Italian” is true in a case where Mary is both French and Italian, and where Bob only knows that Mary is French, but wrongly believes she is not Italian.

Of the three main analyses (4-b), (6-b), and (5-b) and their variants that we have discussed so far, the partitional analysis (5-b) therefore is the most general. We may wonder if it can be even further generalized. A further weakening can indeed be extracted from Karttunen’s (1977) semantics for questions. On Karttunen’s approach, to know whether a or b means to know that a if a is true, and to know that b if b is true, namely:¹

- (8) a. Bob knows whether Mary is French or Italian.
 b. $(a \rightarrow Ka) \wedge (b \rightarrow Kb)$

It is easy to check that (5-b) entails (8-b), since (5-b) is equivalent to $(a \rightarrow Ka) \wedge (\neg a \rightarrow K\neg a) \wedge (b \rightarrow Kb) \wedge (\neg b \rightarrow K\neg b)$. The converse entailment does not hold, however. If Mary is neither French nor Italian, (8-b) trivially entails that Bob knows whether she is French or Italian. Furthermore, if Mary is French and not Italian, and Bob only knows that she is French (but would wrongly believe, for instance, that she is not Italian), then here too (8-b) predicts that Bob knows whether Mary is French or Italian. Both of these predictions can be seen to improve under additional assumptions, but as we shall see in greater detail in Sect. 6, Karttunen’s analysis only appears natural if specific presuppositions are satisfied.

¹ Actually, (8-b) is the so called simplified Karttunen analysis, see (Heim 1994). (8-b) follows from Karttunen’s analysis of alternative questions plus the clause: a knows Q iff a believes all the true answers to Q . Tucked away in footnote 11, p. 18, (Karttunen 1977) proposes a more complex analysis which avoids some of the problems of the simplified version.

For these various reasons, our working assumption in this paper is that the partitional analysis (5-b) captures exactly the truth-conditions of “knowing whether” constructions with alternative questions. If we agree to call (4-b) the *weak analysis* (WA) of embedded alternative questions, (5-b) the *partitional analysis* (PA), (6-b) the *strong analysis* (SA), and call (8-b) *Karttunen’s analysis* (KA), we can summarize the entailment relations between these various analyses as follows (assuming standard properties of the knowledge operator K):

$$\begin{array}{c} \text{(WA)} \\ \uparrow \\ \text{(SA)} \Rightarrow \text{(PA)} \Rightarrow \text{(KA)} \end{array}$$

Once again, all of these analyses treat the alternatives mentioned in the question symmetrically (unlike (7-a) and (7-b)), but (PA) is the only analysis on which a subject knows whether a or b exactly when the subject can give the maximally informative answer relative to both a and b . In the next section, we introduce a combined logic of knowledge and questions in which the partitional analysis of “knowing whether a or b ” constructions can be derived more formally. We return to a more detailed comparison of the partitional analysis (PA) with (WA), (SA) and (KA) in Sect. 6.

3 An epistemic logic with alternative questions

The logic we introduce in this section will serve as a baseline for the rest of the paper. It is a relatively simple system using standard techniques in formal semantics, although the particular application is we believe novel. We will see in later sections that it falls short of the phenomena we wish to capture, and we will propose a number of extensions and alterations that can be ‘layered onto’ it; it is in this sense that this simple system establishes a baseline for everything that follows.

3.1 Syntax

The logic is a version of propositional epistemic logic with questions. The language $\mathcal{L}(\mathcal{C})$ is built up from a countable set \mathcal{C} of propositional constants, with representative elements a and b , and a countable set \mathcal{V} of propositional variables, with representative elements p and q . Besides the epistemic operator “ K ”, the language is special only insofar as it involves a symbol “ $=$ ” of intensional identity between propositions, and a question-forming operator “ $?$ ” which can bind propositional variables.

Definition 1 (Syntax)

$$\phi ::= a \mid p \mid (\neg\phi) \mid (\phi \wedge \psi) \mid (\phi = \psi) \mid (?p.\phi) \mid (K\phi)$$

Parentheses are left off when they serve no disambiguating purpose. “ $\phi = \psi$ ” can be read as: “ ϕ and ψ express the same proposition”. An *approximate* gloss for “ $?p.\phi$ ”, where p can occur free in ϕ , is: “For which propositions p is ϕ true?”. “ $K\phi$ ” means

“The agents knows ϕ ” when ϕ is a propositional formula, and “The agent knows the answer to $?p.\psi$ ” when ϕ is a question of the form “ $?p.\psi$ ”.

Depending on the formula, however, “ $?p.\phi$ ” should be paraphrased by the intended question in natural language. For instance, if a is to mean “Mary speaks French”, and the operator K is referred to John, “ $?p.a$ ” means “does Mary speak French?/whether Mary speaks French”, and “ $K?p.a$ ” means “John knows whether Mary speaks French”. Based on this example, the reader may wonder why we need to write $?p.\phi$ instead of simply $?\phi$. Propositional quantification is introduced to handle more complex questions than yes/no questions, and in particular to give a uniform account of alternative and polar readings of disjunctive questions.

Before turning to the semantics, emphasis must be put on the fact that our suggested paraphrase for “ $?p.\phi$ ” is distinctively non truth-functional. The reason for that is that unlike other formulae of the logic, such formulae only get a truth-value relative to two indices. Formally, this reflects the consideration that on the present account a *wh*-clause, unlike a *that*-clause, denotes an *index-dependent* proposition (Groenendijk and Stokhof 1982). For instance, a *that*-clause like “that Mary speaks French” denotes a proposition that is the same relative to every world of evaluation. By contrast a *wh*-clause like “whether Mary speaks French” denotes a *variable* proposition, the proposition *that Mary speaks French* in worlds in which Mary speaks French, and the proposition *that Mary does not speak French* in worlds in which Mary does not speak French. This dependence of the propositional denotation of a question on the world of evaluation is what motivates the use of double indexing in what follows. Because of that, a question formula of the form “ $?p.\phi$ ” may be viewed equivalently as denoting a propositional concept (Stalnaker 1978), namely a function from worlds to propositions, which to each world w associates the set of worlds that agree exactly with w on the ways of making ϕ true.

3.2 Semantics

A model M for $\mathcal{L}(\mathcal{C})$ is a triple (W, R, V) , where W is a non-empty set of possible worlds, R is an equivalence relation on W , and $V : \mathcal{C} \rightarrow \wp(W)$ is a valuation function which associates a subset of W (a proposition) to each propositional constant a in the language. An assignment function g likewise associates a subset of W to each propositional variable p in \mathcal{V} ; $g[p/\alpha]$ denotes the assignment function exactly like g except at most on p where it assigns the proposition α .

The semantics uses double-indexing to capture question meanings (Lewis 1982; Groenendijk and Stokhof 1982). This is spelled out in terms of a satisfaction relation \models_g , holding between a model M and a pair of worlds w, v , on the one hand, and a formula ϕ , on the other.

Definition 2 (*Semantics*)

$$\begin{aligned}
 M, w, v \models_g a & \text{ iff } w \in V(a) \\
 M, w, v \models_g p & \text{ iff } w \in g(p) \\
 M, w, v \models_g \neg\phi & \text{ iff not } M, w, v \models_g \phi \\
 M, w, v \models_g \phi \wedge \psi & \text{ iff } M, w, v \models_g \phi \ \& \ M, w, v \models_g \psi
 \end{aligned}$$

$$\begin{aligned}
M, w, v \models_g \phi = \psi & \text{ iff } \forall w' : M, w', v \models_g \phi \text{ iff } M, w', v \models_g \psi \\
M, w, v \models_g ?p.\phi & \text{ iff } \forall \alpha \subseteq W : M, w, v \models_{g[p/\alpha]} \phi \text{ iff } M, v, v \models_{g[p/\alpha]} \phi \\
M, w, v \models_g K\phi & \text{ iff } \forall w' : wRw' \Rightarrow M, w', w \models_g \phi
\end{aligned}$$

Disjunction \vee , and material implication \rightarrow can be standardly defined in terms of \neg and \wedge . The double-indexing plays a role only in the clause for questions. If ϕ does not contain any questions, we can easily prove that for all $w, v \in W$: $M, w, v \models_g \phi$ iff $M, w, w \models_g \phi$. Truth and entailment are defined as follows:

Definition 3 (*Truth and entailment*)

- (i) $M, w, v \models \phi$ iff $\forall g : M, w, v \models_g \phi$;
- (ii) $\phi_1, \dots, \phi_n \models \psi$ iff $\forall M, w, v : M, w, v \models \phi_1, \dots, M, w, v \models \phi_n \Rightarrow M, w, v \models \psi$.

The clause for questions is obviously the least transparent aspect of our semantics and can be further illustrated. Intuitively, $?p.\phi$ holds at an ordered pair w, v if and only if w and v agree on which propositions answer the question, namely those p that make ϕ true. In the case of a question of the form $?p.a$, where a is a propositional constant representing a sentence such as “Mary speaks French”, quantification over propositions is vacuous, and therefore the ordered pair w, v satisfies the question if and only if w and v agree on the status of a , that is if they would give the same “Yes” or “No” answer to “Does Mary speak French?”. For more complex questions, the satisfaction of a question can be checked to define a reflexive, symmetric and transitive relation over the set of worlds of a model. As a result, two worlds satisfy a question if they belong to the same cell of the partition induced by the question (Groenendijk and Stokhof 1984).

Here are the representations we use for the two types of question we are investigating:

- (9) a. Polar disjunctive questions: $?p.(a \vee b)$
- b. Alternative questions: $?p.(p \wedge (p = a \vee p = b))$

The alternative reading may be paraphrased as: “Which of either proposition is true: a or b ?” (see Larson (1985), p. 225, and further historical elements by Jespersen there on the link between “which of either” and “whether”). This paraphrase may suggest that knowing whether a or b , on the alternative reading, will prove equivalent to $Ka \vee Kb$, but the logic derives the stronger, partitional meaning that we have argued for in the previous section. For convenience we define some shorthand notation.

Definition 4 (*Shorthand notation*)

$$\begin{aligned}
?(a \vee_A b) & := ?p.(p \wedge (p = a \vee p = b)) \\
?\phi & := ?p.\phi \text{ when } p \text{ does not occur free in } \phi
\end{aligned}$$

Using this simpler notation, we can observe the following equivalences (where $\phi \equiv \psi$ abbreviates $\phi \models \psi$ and $\psi \models \phi$):

- (10) a. $K?a \equiv Ka \vee K\neg a$
- b. $K?(a \vee b) \equiv K(a \vee b) \vee K\neg(a \vee b)$
- c. $K?(a \vee_{\wedge} b) \equiv K?a \wedge K?b$
- d. $K?a \wedge K?b \equiv K(a \wedge \neg b) \vee K(\neg a \wedge b) \vee K(a \wedge b) \vee K\neg(a \vee b)$

As far as the expressiveness of our logic is concerned, these equivalences may seem to obviate the need for a questioning operator at all, at least in the cases we are dealing with in this paper. Equivalently, they may be seen as reduction axioms showing how to rewrite formulae in order to eliminate question operators. However, the point of having a separate semantics for questions is that these equivalences are derived compositionally in our framework, rather than merely stipulated. This means that we can uniformly derive the meaning of embedded questions (such as “John knows *whether* Mary called”) from the meaning of the question (“did Mary call?”) and the meaning of the knowledge operator, without having to rely on case by case conventions. Most of the time, when epistemic logicians need to paraphrase the meaning of knowing-*wh* constructions, they simply agree to use some of the equivalences in (10) for their purposes. In this regard, the fact that we can derive these equivalences within one and the same object language may be viewed as formally establishing the correctness of these conventions.

3.3 Illustration

We close this section with the proof of one direction of (10-c), which shows how to derive these equivalences and allows us to illustrate how the semantics works. Suppose $M, w, v \not\models_g K?p.a \wedge K?p.b$. This means either (a) or (b) is the case:

- (11) (a) $M, w, v \not\models_g K?p.a$
- (b) $M, w, v \not\models_g K?p.b$

Possibility (a) means that $\exists w' : wRw'$ such that $M, w', w \not\models_g?p.a$. This means that:

$$(12) M, w', w \models_g a \not\leftrightarrow M, w, w \models_g a$$

which means:

$$(13) w' \in V(a) \not\leftrightarrow w \in V(a)$$

which is equivalent to:

$$(14) w' \in V(a) \ \& \ (V(a) = V(a) \text{ or } V(a) = V(b)) \not\leftrightarrow \\ w \in V(a) \ \& \ (V(a) = V(a) \text{ or } V(a) = V(b))$$

From (14) it follows that there is proposition $\alpha \subseteq W$ such that:

$$(15) w' \in \alpha \ \& \ (\alpha = V(a) \text{ or } \alpha = V(b)) \not\leftrightarrow \\ w \in \alpha \ \& \ (\alpha = V(a) \text{ or } \alpha = V(b))$$

Therefore, by the clauses for $p, a, =, \wedge$ and \vee we can conclude that there is an $\alpha \subseteq W$ such that:

$$(16) M, w', w \models_{g[p/\alpha]} (p \wedge (p = a \vee p = b)) \not\leftrightarrow \\ M, w, w \models_{g[p/\alpha]} (p \wedge (p = a \vee p = b))$$

which means, by the clause for $?$, that:

$$(17) M, w', w \not\models_g ?p.(p \wedge (p = a \vee p = b))$$

But, since wRw' , this means that:

$$(18) M, w, v \not\models_g K?p.(p \wedge (p = a \vee p = b))$$

Parallel reasoning for (b).

4 Schaffer's puzzle

In the previous sections we have examined different ways of analyzing epistemic constructions with “whether”-complements. We now turn to the discussion of Schaffer's puzzle and the problem of context-sensitivity of “knowing whether” sentences. Schaffer's scenario is one in which a subject, Bob, sees a particular person on TV, who happens to be George Bush. By assumption, Bob is not able to discriminate between George Bush and a clever impersonator like Will Ferrell, but can discriminate between George Bush and Janet Jackson. In particular, Bob can see that the person on TV is not Janet Jackson. Schaffer's puzzle concerns the semantic judgments we should issue about the pair of sentences (1-a) and (1-b), here repeated as (19-a) and (19-b):

(19) a. Bob knows whether George Bush or Will Ferrell is on TV.

b. Bob knows whether George Bush or Janet Jackson is on TV.

Schaffer's original intuition about this example is that while (19-a) should be false, (19-b) should be true. That (19-a) should be false is intuitive: since s cannot discriminate between Bush and Ferrell, s does not know which of Bush and Ferrell is on TV. Arguably, Schaffer's intuition for the truth of (19-b) can in principle be established similarly: since s can discriminate in principle between Bush and Janet Jackson, it should follow that s knows which of Bush and Janet Jackson is on TV.

In Aloni and Égré (2008), we argued that the situation is probably more complicated as regards (19-b). Something we noticed is that the sentence can equally be judged false or infelicitous by competent and rational speakers. Nevertheless, we admit that it can also be judged true. How can this be? Here we present the different options in turn.

4.1 The “false” judgment

First, the partition semantics for alternative questions that we assumed in the previous section makes the basic prediction that (19-b) should in fact be *false* in the scenario under discussion. Indeed, as explained earlier, $K?(b \vee_A j)$ is true if and only if $K?b \wedge K?j$ is also true. But if indeed Bob cannot discriminate Bush from Ferrell, then this should entail $\neg K?b$, namely that Bob does not know whether Bush is on TV or not, and therefore $\neg K?(b \vee_A j)$. Another way to put it is to say that all Bob really knows, by assumption, is that the person on TV is not Janet Jackson, namely $K\neg j$. But this is not enough for Bob to know positively *who* is on TV. For $K?(b \vee_A j)$ to be true in this case, the stronger proposition $K(b \wedge \neg j)$ must be true, namely Bob must also know *that Bush is on TV*. But this is not the case, given Bob's inability to ascertain that it is not Ferrell.

4.2 The “undefined” judgment

The judgment that (19-b) is false does not necessarily match what logically competent speakers of English report about the situation. In fact, one can observe a contrast between:

- (20) a. It is not true that Bob knows whether George Bush or Janet Jackson is on TV.
 b. Bob does not know whether George Bush or Janet Jackson is on TV.

Some of the subjects we asked (including ourselves) feel more readily inclined to judge (20-a) true than (20-b). In the case of (20-b), what we feel is that the sentence is just as inappropriate as its unnegated counterpart (19-b). How can this be? One intuitive reason we see for the infelicity judgment is the following: it is odd to utter (20-b) if Bob knows that Janet Jackson is not on TV. Intuitively, (20-b) suggests that Bob should be equally uncertain about Bush and about Janet Jackson. In Aloni and Égré (2008) we observed that this fact can be derived if one adds to the basic truth-conditions of “knowing whether” sentences a presupposition of symmetry with regard to the alternatives present in an alternative question. In the case of two alternatives, the principle can be expressed as follows:

- (21) a. Bob knows whether A or B *presupposes* that Bob knows whether A iff Bob knows whether B.
 b. $K?(a \vee_{\wedge} b)$ is true or false in w, v provided $(K?a \leftrightarrow K?b)$ is true in w, v ; it is undefined otherwise.

It is immediate that (19-b) will be undefined, since by assumption $K?j$ is true and $K?b$ false in the scenario under discussion. One important consequence of adding this presupposition to our logic concerns the negation of knowledge sentences with alternative questions. In the logic we defined in the previous section, only the following weak entailment holds:

$$(22) \neg K?(a \vee_{\wedge} b) \models \neg K?a \vee \neg K?b$$

That is, “Bob does not know whether Bush or Janet Jackson is on TV” implies “Bob does not know whether Bush is on TV or Bob does not know whether Janet Jackson is on TV”. Generally, however, from “Bob does not know whether Bush or Janet Jackson is on TV”, we seem to infer the stronger consequence “Bob does not know whether Bush is on TV *and* Bob does not know whether Janet Jackson is on TV”.

In Sect. 6 we propose an extension of our basic logic in which logical consequence (marked \vdash) is made sensitive to presuppositions such as (21). In this framework, every model in which $\neg K?(a \vee_{\wedge} b)$ is defined and true relative to the symmetry presupposition $(K?a \leftrightarrow K?b)$ is a model in which $\neg K?a \wedge \neg K?b$ is true:

$$(23) \neg (K?(a \vee_{\wedge} b)_{(K?a \leftrightarrow K?b)}) \vdash \neg K?a \wedge \neg K?b$$

The symmetry presupposition is a particular case of a more general principle that Chemla calls the principle of epistemic similarity for disjunctive sentences (Chemla 2008). We shall not try to motivate it further here, but discuss it in greater detail in Sect. 6. The important fact to bear in mind, from a logical point of view, is that this presupposition is probably the weakest we can add to the partitional semantics in order to support the intuition that (19-b) ought to be undefined in the context of Schaffer’s scenario, rather than true or false.

4.3 The “true” judgment

Besides judgments of undefinedness and falsity for (19-b), there is a point to the intuition that the sentence can be true, as Schaffer submits. We gave an intuitive reason above already: because Bob is able in principle to discriminate between Bush and Janet Jackson, his knowledge that it is not Jackson may be sufficient relative to those two alternatives to conclude that the person on TV is Bush. In Aloni and Égré (2008), we described this using a mechanism of *topical restriction*. Bob’s knowledge of which of the alternatives holds is evaluated in his epistemic state restricted to the propositions mentioned in the question. A static approximation of the context-sensitive truth-conditions stated in Aloni and Égré (2008) is the following:

$$(24) \quad M, w, v \models_g K?(a \vee_A b) \text{ iff } \forall w' : wRw' \text{ and } (M, w', v \models a \text{ or } \\ M, w', v \models b) \Rightarrow M, w', v \models_g?(a \vee_A b)$$

This semantics predicts that “Bob knows whether the person on TV is Bush or Jackson” is true then, even though Bob cannot discriminate in principle between Bush and Ferrell. And it preserves the fact that “Bob knows whether the person on TV is Bush or Ferrell” is false, since the restriction then is idle.

The restriction mechanism in (24) raises two issues. The first is linguistic and concerns the pragmatic or semantic character of this restriction. If such a restriction takes place, does it take place systematically? If it were so, one should be able to say in the same context:

- (25) a. Bob knows whether the person on TV is Bush or Jackson, although he does not know whether the person on TV is Bush or Ferrell.
 b. $K?(b \vee_A j) \wedge \neg K?(b \vee_A f)$

But this should imply that “Bob knows whether A or B” is synonymous with “Assuming A or B is true, Bob knows which of the two is true”. This is not obviously the case, however, and (25-a) (placing the two conjuncts in one and the same context) sounds like a near contradiction. This suggests that the mechanism in question is fundamentally pragmatic, and that more contextual dependence is at play (Aloni and Égré 2008).

The second issue is epistemological. Suppose Bob pays attention to all the alternatives present, namely Bush, Jackson and Ferrell. Then Bob’s uncertainty regarding the Bush-Ferrell pair might affect his confidence that the person on TV is Bush when asked “Is it Bush or Jackson?”. But are the judgments the same if Bob is not even aware of the existence of Ferrell, but only of that of Jackson and Bush? Then (25) may more easily be judged true. This suggests that more should also be said about the attention the agent pays to the alternatives.

In the following sections our aim will thus be to characterize more tightly the contextual parameters that influence our judgments of truth value for a sentence like (19-b). In Sect. 5, we deal first with the epistemological issue we just raised. In Sect. 6 we provide a detailed treatment of the notion of presupposition and discuss the status of the “undefined” judgment specifically. To a large extent the two sections explore different paths and can be read independently of each other. In particular, Sects. 5 and 6 may be seen to give different accounts of the “true” judgment for Schaffer’s examples.

In Sect. 5, the “true” judgment is derived in a bivalent system, and judgments of undefinedness are left as an extension; in Sect. 6, by contrast, judgments of undefinedness come first, and a mechanism of presupposition accommodation is used to derive the “true” judgment.

5 Attention and relevance

As is well known from the epistemological literature, one should distinguish precisely between subject’s factors and attributor’s factors in discussing the semantics of knowledge sentences (Stanley 2005). In this section we examine Schaffer’s scenario in the light of two such factors that we call *attention* (a subject’s factor) and *relevance* (an attributor’s factor). The surprising conclusion of this section will be that a significant part of Schaffer’s puzzle need not have anything to do with the particular properties of knowing-*whether* constructions after all. A model of attention and relevance, motivated primarily by the semantics of knowing-*that* sentences, is enough to give us both the “true” and “false” judgements we considered in the previous section. We set aside the undefinedness judgments until the end of this section; that is, we make no use of the symmetry principle (21) and keep the logic entirely bivalent for the purpose of teasing apart the “true” and “false” judgments first.

5.1 Contextualism and proper ignorance

Attention is a subject-based parameter, which concerns whether or not a subject is consciously attending to a particular epistemic possibility. For instance, Bob may fail to attend to the possibility that the person on TV is Ferrell, either because he does not think about Ferrell, or because he has never heard of Ferrell. Relevance on the other hand is a normative, ascriber-based parameter, which concerns whether it is permissible or not, from an external ascriber’s standpoint, to ignore a particular possibility.

As contextualists such as Dretske and Lewis have argued (Dretske 1970; Lewis 1996) in some cases some epistemic possibilities can be irrelevant and may be “properly ignored” (for instance very remote far-fetched sceptical possibilities). Cases in which a subject fails to attend to *relevant* possibilities undercut his knowledge, but conversely, when these possibilities may be properly ignored, the subject’s lack of attention to them will leave his knowledge unscathed. To account for the force of particular instances of the sceptic’s argument, if the subject does not in fact ignore a possibility (for instance when it is mentioned as part of a sceptical argument) he is thereby no longer properly ignoring it and it may likewise undercut his knowledge.

Based on this idea, the point of this section is to show that Schaffer’s intuition, according to which the sentences $\neg K?(b \vee_A f)$ and $K?(b \vee_A j)$ can be true together, should be examined in relation to these two parameters; in particular, relative to whether the possibility that Ferrell is on TV can be properly ignored or not, and relative to whether it is attended to by Bob or not.

According to the partitioned semantics of Sect. 2, for Schaffer’s intuition to hold we must have simultaneously $\neg K?f$, $K?b$ and $K?j$. The prediction we shall derive is

that this is only possible (under our running assumptions) when the Ferrell possibility is both irrelevant and not attended to. In all other cases, the relevance of f , or the fact that it is attended to by Bob, forces both $K?b$ and $K?f$ to be false together.

5.2 Modelling attention and relevance

In order to model attention and relevance, we need to enrich the models introduced so far. Our framework is based on the logic of awareness of Fagin and Halpern (1988), who use a syntactic notion of attention to sentences. It departs from their account by the introduction of two additional parameters, namely by distinguishing possibilities that are *relevant* on the one hand, and possibilities that are actually *entertained* on the other. Thus relevance in our approach is modelled by associating to each world a set of relevant alternatives, possibilities that cannot be properly ignored at the world. Attention, on the other hand, is defined by two components: by the possibilities the agent is actually entertaining, and also by the sentences the agent consciously considers. Attention to sentences is needed to avoid the classic problem of logical omniscience, or of epistemic closure (Fagin and Halpern 1988). For instance, Bob may attend to the ‘metaphysical truth’ that Bush is not Janet Jackson without attending to the fact that Bush is not Ferrell, even though both are true in exactly the same worlds (namely all of them). Attention to possibilities, on the other hand, serves to represent what the agent takes for granted or *assumes* to hold in a given situation (see Franke and de Jager, forthcoming). Attention to possibilities is needed besides attention to sentences to capture this idea, because no obvious mapping exists from unawareness of sentences to unawareness of possibilities. Basically, making knowledge relative to only attended sentences and relevant possibilities might in some cases be too strict on what the agent ignores, and in other cases be too liberal on what the agent knows.²

So let us define an *attention and relevance* model (for a single agent) as a structure $(W, R, S, E, \mathcal{A}, V)$, where (W, R, V) is as in Sect. 2. As usual, R represents the agent’s capacity to discriminate between possibilities, irrespective of whether he is attending to them or not. $S: W \rightarrow \wp(W)$ associates with each world w the set $S(w)$ of worlds the agent *should* entertain in w (relevant alternatives; possibilities the agent may not properly ignore); for each w we require $w \in S(w)$ (without this constraint contextualist knowledge is no longer factive). $E: W \rightarrow \wp(W)$ associates with w the set $E(w)$ of possibilities the agent actually *entertains*; finally \mathcal{A} is a function which to each world w associates a set of propositional constants $\mathcal{A}(w)$, which our agent consciously *attends to*. Let $\mathcal{L}(\mathcal{A}(w))$ be the set of sentences generated from $\mathcal{A}(w)$; we say that the agent attends to ϕ at w if $\phi \in \mathcal{L}(\mathcal{A}(w))$.

² For instance, the semantics we state below in Definition 5 would make unwanted predictions for the left column of the table presented in Subsect. 5.3 if besides the usual epistemic accessibility relation (R) we just retained the A and S parameters, or even simply the A parameter, as in more standard logics of awareness. Furthermore, from the example discussed there, it may appear that the set $E(w)$ of attended possibilities can be defined in terms of the set $\mathcal{A}(w)$ of attended sentences, by stipulating that if $f \notin \mathcal{A}(w)$ then f holds nowhere in $E(w)$. This would not be correct in general, however. Consider the propositional atom h , for “Bush still has his own hair.” If Bob fails to wonder whether h ($h \notin \mathcal{A}(w)$) he should not thereby assume that Bush is wearing a toupée.

Using these models we can define a new logic, with a new satisfaction relation \models_g^* . This relation is just like \models_g , except for the clause for knowledge:

Definition 5

$$M, w, v \models_g^* K\phi \text{ iff } \phi \in \mathcal{L}(\mathcal{A}(w)) \text{ and } \forall w' \in E(w) \cup S(w) : wRw' \Rightarrow M, w', w \models_g^* \phi$$

The definition says that an agent knows that ϕ if firstly he attends to ϕ (we describe *conscious*, rather than *potential* knowledge), and then if ϕ holds in every epistemic alternative that *is* entertained by the agent and in every alternative that *should* be entertained, by the normative criterion of relevance.³

Our original logic is in a sense a special case of this one. Suppose $M = (W, R, S, E, \mathcal{A}, V)$ is an attention and relevance model in which $S(w) = E(w) = W$ and \mathcal{A} associates with each world the *full* set of propositional constants \mathcal{C} . Then \models_g^* interpreted on M coincides with \models_g interpreted on (W, R, V) .

5.3 Predictions

We have to consider four scenarios, in each of which the person on TV is Bush. Each time, Bob cannot discriminate in principle between the actual world w_b and the world w_f in which Ferrell is on TV, but is able to discriminate w_j from the other two (the relation R is assumed reflexive, symmetric and transitive). Our scenarios vary along two axes. Along the attention axis, Bob may attend to f or not; if he does, he entertains w_f , otherwise he does not (this is a natural assumption in this scenario, which may be seen to follow from standard properties of awareness logics). Along the relevance axis, the possibility w_f may either be properly ignored or not.

To get the right intuitions about the relevance axis, imagine two different contexts in which Bob might believe he sees Bush on television. In the first, he is watching a live broadcast from the White House, with Bush addressing the American people. In this setting the idea that Bush is being impersonated by Will Ferrell seems a mere sceptical possibility, along the same lines as the possibility that Bob is hallucinating the broadcast. We say therefore that $S_1(w_b) = \{w_b\}$ (there is no normative requirement that Bob attends to any non-Bush possibilities).

In a different scenario, he sees footage of Bush making fun of his administration and his own policies. The fact that this behavior is completely outrageous for an American president but perfectly appropriate for a comic impersonator means that Bob would be naive, at best, to discount the possibility that he is seeing Ferrell. We represent this with $S_2(w_b) = \{w_b, w_f\}$ (he is normatively required to entertain the possibility w_f).

Cutting across these two scenarios are the two states of attention Bob may be in. We'll assume he attends to b and j , and entertains w_b and w_j , in both states. But in one

³ This semantics will over-generate for complex nested operator constructions such as $K?K?p$. A partial solution is to add the attention and relevant alternative sets as free parameters rather than world-based functions. In either case, however, a number of constraints on possible combinations of these parameters with R and with each other are required. For reasons of space we omit the details.

he fails to attend to f and ignores w_f : $\mathcal{A}_1(w_b) = \{b, j\}$ and $E_1(w_b) = \{w_b, w_j\}$. In the other he consciously considers the possibility that it is Ferrell: $\mathcal{A}_2(w_b) = \{b, j, f\}$ and $E_2(w_b) = \{w_b, w_j, w_f\}$. All scenarios are summarized in the following table, in which the world of evaluation is w_b :

	$S_1 = \{w_b\}$	$S_2 = \{w_b, w_f\}$
$E_1 = \{w_b, w_j\}$	$K?(b \vee_A j), K?b$	$\neg K?(b \vee_A j), \neg K?b$
$\mathcal{A}_1 = \{b, j\}$	$\neg K?(b \vee_A f), \neg K?f$	$\neg K?(b \vee_A f), \neg K?f$
$E_2 = \{w_b, w_j, w_f\}$	$\neg K?(b \vee_A j), \neg K?b$	$\neg K?(b \vee_A j), \neg K?b$
$\mathcal{A}_2 = \{b, j, f\}$	$\neg K?(b \vee_A f), \neg K?f$	$\neg K?(b \vee_A f), \neg K?f$

Of the four combinations we get, only one supports Schaffer’s judgement: S_1 with \mathcal{A}_1 and E_1 . In this combination Bob ignores w_f , and according to S_1 he is allowed to ignore it. This supports the judgement $K?b \wedge K?f$; since he doesn’t attend to f , on the other hand, $K?f$ is false.

In the other three combinations Bob fails to know whether Bush is on TV, but interestingly enough, for different reasons. Whenever he attends to f he is consciously *uncertain* whether Ferrell or Bush is on TV, however when \mathcal{A}_1 and E_1 combine with S_2 he ignores a possibility that he ought to attend to: the agent correctly believes that Bush is on TV, but this is not enough to ascribe knowledge.⁴

More generally, Bob may fail to know at w that some formula ϕ holds for three quite different reasons:

1. Because he is consciously uncertain about ϕ (the most standard reason to fail to know); this underpins $\neg Kb$ whenever $w_f \in E(w)$.
2. Because he fails to consider some non- ϕ possibility that he ought to consider (he may well believe he knows ϕ , but wrongly so). Our scenario combining S_2 with E_1 and \mathcal{A}_1 supports $\neg Kb$ for this reason.
3. Because he does not attend to ϕ (he cannot even wonder *whether* ϕ). This is the reason $K(b \rightarrow \neg f)$ does not hold for any scenario with \mathcal{A}_1 , in which he does not attend to f .

Worth emphasizing in our approach is the fact that the agent’s first-person uncertainty about a possibility will override the fact that the latter is not deemed relevant on normative grounds. For such cases, the agent’s conscious uncertainty will undercut her knowledge. This is as it should be in our opinion, for one may come to entertain possibilities that are not yet considered relevant by a third party or by the community, but that would turn out to be relevant at a later stage of the inquiry.

5.4 Summary

The strategy of explanation we followed in this section differs significantly from the one we sketched with (24) earlier. The topical restriction strategy outlined in (24)

⁴ Belief that ϕ in our framework may be defined by simply rephrasing Definition 5 without invoking S . Note that such a definition would preserve the entailment from knowledge to belief.

predicts that $K?(b \vee_A j) \wedge \neg K?(b \vee_A f)$ can be true without modification of the semantics for K . But we saw that it makes predictions that are intuitively inconsistent. In particular, $K?(b \vee_A j) \wedge \neg K?(b \vee_A f)$ should then be compatible with both $\neg K?b$ and $\neg K?f$, unless topical restriction should carry over to *knowing whether* sentences involving polar questions.⁵

In contrast, the account put forward in the present section rests on a more fine-grained semantics for knowledge, but makes no modification of the semantics of *knowing whether* clauses proper. The strength of the present account is that Schaffer's conjunction $K?(b \vee_A j) \wedge \neg K?(b \vee_A f)$ is predicted to be true only when f is a possibility that is both ignored and irrelevant, and otherwise false. Furthermore, it derives judgments of truth and falsity about Kb and Kf in a principled way. On this account, the truth of $K?(b \vee_A j) \wedge \neg K?(b \vee_A f)$ is always incompatible with that of $\neg K?b \wedge \neg K?f$.

This account has however still not explained the “undefined” judgements of Sect. 4. In the next section we give a presuppositional analysis that captures these cases. The semantics we develop there is compatible with the one we defined in this section, in the sense that it would allow us to predict judgments of undefinedness. For instance, in all four scenarios we distinguished in the present section, $K\neg j$ and therefore $K?j$ holds, and likewise $\neg K?f$ holds everywhere. In all conditions except for the S_1 - E_1 - \mathcal{A}_1 condition, on the other hand, $\neg K?b$ holds. If we adopted a presupposition of symmetry such as explained in (21) across the board, then $\neg K?(b \vee_A f)$ would remain true everywhere in the present setting, except in the first scenario, where it would come out as undefined rather than true (because $K?b$ and $\neg K?f$ would violate symmetry). Conversely, $K?(b \vee_A j)$ would remain true in the first scenario, where symmetry between $K?b$ and $K?j$ is respected, but would come out as undefined rather than false in the other three scenarios.

In the next section, we first describe the effect of adding presupposition to our basic logic rather than to the present and richer system. One motivation to do so is that if we set aside the effect of relevance and attention parameters, an alternative explanation of the “true” judgment can still be provided on the basis of a specific theory of presupposition accommodation.

6 Pragmatic inference and presuppositions

Presupposition is a pragmatic inference notion based on the idea that the utterance of a sentence may ‘presuppose’ (or take for granted) certain ‘background information’ (for a survey of approaches see Beaver 1997). A standard example is “The king of France is bald”, which presupposes that there exists a unique king of France. Unlike entailments, presuppositions are preserved under negation: “The king of France is *not* bald” has just the same presupposition. Indeed, a standard characterization of the notion (which we will implement formally below) is just that the presupposition of ϕ is

⁵ A dynamic mechanism of topical restriction is presented in Aloni and Égré (2008), making room for this possibility. We do not discuss it here, as our focus is on static aspects of the context-sensitivity of knowing whether sentences.

whatever can be inferred both from ϕ and from $\neg\phi$. Another standard notion we make use of is a truth-value gap representing *infelicity*: a sentence whose presuppositions are not met (such as “The king of France is not bald”) may strike hearers as neither true nor false but ‘odd’ or infelicitous.

In this section we implement the notion of a presupposition in our basic system (Sect. 6.1). The result will be a partial logic where presuppositions are represented by formulas like ϕ_ψ (read as ‘ ϕ with presupposition ψ ’) which introduce undefinedness whenever the presupposition ψ is not true. In Sect. 6.2 we discuss two presuppositions that have been proposed in the literature for alternative questions (the symmetry presupposition already introduced in Sect. 4 and the ‘exactly one’ presupposition), and we show how they can be used to generate both the judgements “undefined” (presupposition failure) and “true” (via presupposition *accommodation*) in Schaffer’s scenario described in Sect. 4. It turns out that the differences between the various analyses of alternative questions we gave in Sect. 2 collapse when we perform such a presuppositional analysis. It might seem, then, that any of these analyses would do as well as the partitional one we favour. In the remaining part of the section we will argue to the contrary. The inference patterns supporting these presuppositions have a pragmatic character—both can be *cancelled* under certain circumstances, which means that in cancellation cases we need the full partition analysis for questions to get the right truth conditions.

6.1 A partial logic for presupposition

In this section we present one way to implement presuppositions in the basic logic introduced in Sect. 3. First of all, we extend the language with a binary presupposition connective (ϕ_ψ) (Beaver 1997).

Definition 6 (Syntax)

$$\phi =: a \mid p \mid (\neg\phi) \mid (\phi \wedge \psi) \mid (\phi = \psi) \mid (?p.\phi) \mid (K\phi) \mid (\phi_\psi)$$

The formula (ϕ_ψ) should be read as ‘ ϕ with the presupposition ψ ’.⁶ So, for example, ‘John knows whether a or b ’ could translate in this logic as $(K?(a \vee_A b))_{(K?a \leftrightarrow K?b)}$ with explicit reference to the symmetry presupposition. In what follows we will leave out superfluous parenthesis and adopt the convention to write $K\phi_\pi$ for $(K\phi)_\pi$.

Models and assignments for such presuppositional language are defined as in Sect. 3.

In the definition of the partial semantics, $M, w, v \vdash_g \phi$ abbreviates ϕ is **true** in M, w, v with respect to g ; $M, w, v \dashv_g \phi$ abbreviates that ϕ is **false** in M, w, v with respect to g ; and $M, w, v \mid_g \phi$ stands for $M, w, v \vdash_g \phi$ or $M, w, v \dashv_g \phi$.

⁶ In principle, the extended language allows us to presuppose direct questions: $(\phi_{?p.\psi})$ is a well-formed formula. In such a static system, however, the presupposition of a question, does not make much sense. See Aloni et al. (2007) for a dynamic account of presupposed questions.

Definition 7 (*Semantics*)

$$M, w, v \vdash_g a \quad \text{iff } w \in V(a)$$

$$M, w, v \nvdash_g a \quad \text{iff } w \notin V(a)$$

$$M, w, v \vdash_g p \quad \text{iff } w \in g(p)$$

$$M, w, v \nvdash_g p \quad \text{iff } w \notin g(p)$$

$$M, w, v \vdash_g \neg\phi \quad \text{iff } M, w, v \nvdash_g \phi$$

$$M, w, v \nvdash_g \neg\phi \quad \text{iff } M, w, v \vdash_g \phi$$

$$M, w, v \vdash_g \phi \wedge \psi \quad \text{iff } M, w, v \vdash_g \phi \ \& \ M, w, v \vdash_g \psi$$

$$M, w, v \nvdash_g \phi \wedge \psi \quad \text{iff } M, w, v \nmid_g \phi \ \& \ M, w, v \nmid_g \psi \ \& \ \text{not } M, w, v \vdash_g \phi \wedge \psi$$

$$M, w, v \vdash_g \phi = \psi \quad \text{iff } M, w, v \nmid_g \phi \ \& \ M, w, v \nmid_g \psi \ \&$$

$$\forall w' \in W : M, w', v \vdash_g \phi \ \text{iff } M, w', v \vdash_g \psi$$

$$M, w, v \nvdash_g \phi = \psi \quad \text{iff } M, w, v \nmid_g \phi \ \& \ M, w, v \nmid_g \psi \ \& \ \text{not } M, w, v \vdash_g \phi = \psi$$

$$M, w, v \vdash_g ?p.\phi \quad \text{iff } \forall \alpha \subseteq W : M, w, v \nmid_{g[p/\alpha]} \phi \ \&$$

$$M, w, v \vdash_{g[p/\alpha]} \phi \ \text{iff } M, w, v \vdash_{g[p/\alpha]} \phi$$

$$M, w, v \nvdash_g ?p.\phi \quad \text{iff } \forall \alpha \subseteq W : M, w, v \nmid_{g[p/\alpha]} \phi \ \& \ \text{not } M, w, v \vdash_g ?p.\phi$$

$$M, w, v \vdash_g K\phi \quad \text{iff } M, w, v \nmid_g \phi \ \& \ \forall w' : wRw' \Rightarrow M, w', w \vdash_g \phi$$

$$M, w, v \nvdash_g K\phi \quad \text{iff } M, w, v \nmid_g \phi \ \& \ \text{not } M, w, v \vdash_g K\phi$$

$$M, w, v \vdash_g \phi\psi \quad \text{iff } M, w, v \vdash_g \phi \ \& \ M, w, v \vdash_g \psi$$

$$M, w, v \nvdash_g \phi\psi \quad \text{iff } M, w, v \nvdash_g \phi \ \& \ M, w, v \nvdash_g \psi$$

Truth and entailment are defined as follows:

Definition 8 (*Truth and entailment*)

(i) $M, w, v \vdash \phi$ iff $\forall g : M, w, v \vdash_g \phi$;

(ii) $\phi_1, \dots, \phi_n \vdash \psi$ iff $\forall M, w, v : M, w, v \vdash \phi_1, \dots, M, w, v \vdash \phi_n \Rightarrow M, w, v \vdash \psi$.

Given this notion of entailment as preservation of truth, a standard definition of presupposition is as follows:

Definition 9 (*Presupposition*) ϕ presupposes ψ iff $\phi \vdash \psi$ & $\neg\phi \vdash \psi$

(It is easy to see from the clauses for ϕ_ψ and for negation that ϕ_ψ always presupposes ψ in this sense, as we of course would like.)

For simplicity, following Karttunen's (now standard) terminology, all connectives and operators are treated here as *holes* (allowing presupposition projection):⁷

- (26) a. $\neg(\psi_\pi)$ presupposes π
 b. $\phi \wedge (\psi_\pi)$ presupposes π
 c. $\phi = (\psi_\pi)$ presupposes π
 d. $?p.(\psi_\pi)$ presupposes π
 e. $K(\psi_\pi)$ presupposes π

Conjunction and negation are interpreted here as in Kleene's weak system. Note that $K(\psi_\pi)$ does not presuppose $K\pi$, because $\neg K(\psi_\pi)$ does not entail $K\pi$. Thus, for example, "John knows whether the queen is bald" presupposes that there is a unique queen, but does not presuppose that John knows that there is a unique queen.

The formula ϕ_ψ introduces undefinedness whenever ψ is not true. This is meant to capture judgments of infelicity in cases of presupposition failure. In ordinary conversation, however, presupposition failure can be (and often is) repaired by *accommodation* (Lewis 1979). That is, a cooperative hearer may respond to presupposition failure by changing his beliefs in order to make the presupposition hold, rather than complaining about the problem.

One way to model global accommodation in this logic is by making the notion of entailment sensitive to pragmatics. This notion is sometimes called Strawson entailment (von Stechow 1999), and it is defined by restricting attention to models that satisfy the presuppositions of all the sentences involved.

Definition 10 (*Strawson entailment*) Let π be the conjunction of the presuppositions of ϕ_1, \dots, ϕ_n and ψ .

$$\phi_1, \dots, \phi_n \vdash_S \psi \text{ iff } \forall M, w, v : M, w, v \vdash \pi, M, w, v \vdash \phi_1, \dots, \\ M, w, v \vdash \phi_n \Rightarrow M, w, v \vdash \psi$$

Two presuppositions have been proposed in the literature for alternative questions: the *symmetry presupposition* and the *exactly one presupposition*. In the following section we show that assuming these two presuppositions within this logic and with this model of accommodation, we will be able to capture both the judgements "undefined" (presupposition failure) and "true" (via accommodation) of Schaffer's scenario presented in Sect. 4.

⁷ Note that conjunction is probably better analyzed as a *filter* (the presupposition may be modified if projected). We leave the discussion of alternative treatments for another occasion, however, since the projection of presuppositions is not our main focus in this paper.

6.2 Two presuppositions

6.2.1 Symmetry

Assuming S knows Napoleon was not born in 1869, it may appear altogether infelicitous to say:

(27) S does not know whether Napoleon was born in 1769 or in 1869.

More generally, “S does not know whether A or B” seems to imply “S does not know whether A and S does not know whether B”. However, in the situation under discussion, S knows that Napoleon was not born in 1869.

To account for this intuition, Aloni and Égré (2008) proposed the following presupposition:

(28) **Symmetry presupposition**

- a. S knows whether A_1 or ... or A_n presupposes
- b. S knows whether $A_1 \leftrightarrow \dots \leftrightarrow$ S knows whether A_n

As explained in Sect. 4, symmetry failure can also account for the “undefined” judgement in Schaffer’s scenario. Sentence (29-a) is naturally judged infelicitous because Bob has different epistemic attitudes towards the two alternatives mentioned (he knows *that*, so also whether, Jackson is not on TV, but he fails to know whether Bush is on TV).

- (29) a. Bob knows whether Bush or Jackson is on TV.
 b. $K?(b \vee_A j)_{K?b \leftrightarrow K?j}$

But when presuppositions fail, people can still accommodate, and this is precisely what the notion of Strawson entailment is meant to represent.⁸ It is easy to see that the following Strawson entailment holds:

$$(30) K \neg j \vdash_s K?(b \vee_A j)_{K?b \leftrightarrow K?j}$$

because entailment (31) straightforwardly holds:

$$(31) K \neg j, K?b \leftrightarrow K?j \vdash K?(b \vee_A j)_{K?b \leftrightarrow K?j}$$

It seems then that by adding a mechanism of presupposition accommodation we also have an explanation of the “true” judgement discussed in Sect. 4. Knowing that Jackson is not on TV pragmatically entails knowing whether Jackson or Bush is on TV.

There is something missing though in this explanation. From “Bob knows that Jackson is not on TV” and the alternative reading of “Bob knows whether Jackson or Bush is on TV”, the default may be to conclude that “Bob knows *that* Bush is on TV”.

⁸ Strawson entailment encodes a *static* notion of accommodation. It says what follows when people are willing to accommodate, but is silent on the question of when (in what contexts) they will be willing, or what kind of dynamic effects accommodation will have on their information. In fact, though, a person’s willingness to accommodate should be affected by the information they already hold, in ways that Strawson entailment cannot capture; in Schaffer’s scenario, someone who believes that $\neg K?b$ and $K?j$ should probably be unwilling to accommodate the symmetry presupposition $K?b \leftrightarrow K?j$. Representing this would require a properly dynamic account, with distinct information states for the interlocutors who do the presupposing and the accommodating. Assessing how problematic this is for our story will have to wait for such an account.

Our pragmatic analysis so far, instead, can only derive that “Bob knows *whether* Bush is on TV”. This fact suggests that there is another pragmatic effect of the use of alternative questions to which we turn in the following subsection.

6.2.2 *Exactly one*

Many authors have observed that the use of an alternative question suggests that one and at most one of the alternatives is true (see e.g. Cornulier (1982) for a discussion of this point). If Bob knows that Mary’s only nationality is German, then under the alternative reading it seems infelicitous to say either of:

- (32) a. Bob knows whether Mary is French or Italian.
 b. Bob doesn’t know whether Mary is French or Italian.

The same infelicity appears if Mary happens to be both French and Italian. These judgments can be explained in terms of a failure of an *exactly one* presupposition (e.g. Karttunen 1977):

- (33) **Exactly one presupposition**
 a. Bob knows whether A_1 , or . . . or A_n presupposes
 b. (i) *At least* one of the A_1, \dots, A_n is true
 (ii) *At most* one of the A_1, \dots, A_n is true

It is easy to see that if we assume both the symmetry and *exactly one* presuppositions, we have a full explanation of the “true” judgement in Schaffer’s scenario. If we are ready to accommodate the presupposition of “Bob knows whether Bush or Jackson is on TV”, then from “Bob knows that Jackson is not on TV” we can conclude that “Bob knows that Bush is on TV”. In what follows we write σ for the symmetry presupposition and π for the exactly one presupposition, that is, in $K?(a \vee_A b)_{(\sigma \wedge \pi)}$, σ is $K?a \leftrightarrow K?b$ and π is $(a \vee b) \wedge \neg(a \wedge b)$. The following pragmatic inferences hold:

- (34) a. $K\neg j \vdash_S K?(j \vee_A b)_{(\sigma \wedge \pi)}$ (by symmetry)
 b. $K\neg j, K?(j \vee_A b)_{(\sigma \wedge \pi)} \vdash_S Kb$ (by *at least* presupposition)

On the other hand, from “Bob knows that Jackson is not on TV” we cannot conclude “Bob knows whether Bush or Ferrell is on TV”; this explains Schaffer’s intuition of the contrast in difficulty between the two alternative questions:

- (35) a. $K\neg j \not\vdash_S K?(f \vee_A b)_{(\sigma \wedge \pi)}$
 b. $K\neg j, K?(f \vee_A b)_{(\sigma \wedge \pi)} \not\vdash_S Kb$

The following are three more interesting results we are now able to prove (the subscript S is present only when Strawson entailment is needed):

- (36) $K?(a \vee_A b)_{(\sigma \wedge \pi)} \equiv K(a \wedge \neg b) \vee K(b \wedge \neg a)$
 a. $K?(a \vee_A b)_{(\sigma \wedge \pi)} \vdash K(a \wedge \neg b) \vee K(b \wedge \neg a)$ (by *exactly one*)
 b. $K(a \wedge \neg b) \vee K(b \wedge \neg a) \vdash K?(a \vee_A b)_{(\sigma \wedge \pi)}$
- (37) $K?(a \vee_A b)_{(\sigma \wedge \pi)} \equiv_S Ka \vee Kb$
 a. $K?(a \vee_A b)_{(\sigma \wedge \pi)} \vdash Ka \vee Kb$ (by *at least*)
 b. $Ka \vee Kb \vdash_S K?(a \vee_A b)_{(\sigma \wedge \pi)}$ (by *at most* and symmetry)

- (38) $K?(a \vee_{\wedge} b)_{(\sigma \wedge \pi)} \equiv_S (a \rightarrow Ka) \wedge (b \rightarrow Kb)$
 a. $K?(a \vee_{\wedge} b)_{(\sigma \wedge \pi)} \vdash (a \rightarrow Ka) \wedge (b \rightarrow Kb)$
 b. $(a \rightarrow Ka) \wedge (b \rightarrow Kb) \vdash_S K?(a \vee_{\wedge} b)_{(\sigma \wedge \pi)}$ (by *at least* and symmetry)

If we assume the *exactly one* presupposition, then “knowing whether A or B” is equivalent to “knowing (A and not B) or knowing (B and not A)”. If we are ready to accommodate both the symmetry and *exactly one* presuppositions, then “knowing whether A or B” is equivalent to “knowing A or knowing B”. And, finally, if we accommodate the *at least* and symmetry presuppositions, then we have pragmatic equivalence between our partitional analysis and Karttunen’s analysis. Under these circumstances, then, the various alternative notions of knowing-*wh* discussed in Sect. 2 and repeated in (39) turn out to be equivalent:

- (39) a. Partition analysis (PA) : $K?a \wedge K?b$
 b. Strong analysis (SA) : $K(a \wedge \neg b) \vee K(b \wedge \neg a)$
 c. Weak analysis (WA) : $Ka \vee Kb$
 d. Karttunen’s analysis (KA) : $(a \rightarrow Ka) \wedge (b \rightarrow Kb)$

Should we then conclude that the semantic analyses of knowing whether A or B discussed there are correct after all? No. In a moment we will provide evidence that both the *symmetry* and the *exactly one* inferences are cancellable. Each of the analyses discussed above, apart from the partitional analysis, fails to account for cases of violations of whichever presuppositions are used to prove their pragmatic equivalence to the partition theory. Thus, KA fails to account for cases of violation of the *at least* or the symmetry presupposition; SA of the *at least* and *at most* presuppositions; and WA of all three presuppositions. We turn now to the cancellation cases.

6.2.3 Cancellations

Symmetry. Presuppositions, like entailments, are not cancellable in positive sentences; the following sentence is contradictory:

- (40) The king of France is bald, but there is no king of France.

Unlike entailed inferences, however, presuppositions can sometimes be cancelled under negation:⁹

- (41) a. It is not true that the king of France is bald, because there is no king of France.
 b. ? The king of France is not bald, because there is no king of France.

Our symmetry inference seems to follow the same pattern. Sentence (42) is contradictory under the alternative reading of the question, but (43-a) is true.

⁹ In our presupposition logic, sentences like (41-a) could be represented in terms of a denial operator # with the following interpretation (see again Beaver 1997):

$$M, w, v \vdash_g \# \phi \text{ iff not } M, w, v \neg_g \phi$$

$$M, w, v \neg_g \# \phi \text{ iff } M, w, v \vdash_g \phi.$$

- (42) Bob knows whether Mary is Italian or French. Bob knows whether she is Italian but he doesn't know whether she is French.
- (43) a. It is not true that Bob knows whether Mary is Italian or French, because Bob knows whether Mary is Italian but he doesn't know whether she is French.
 b. ? Bob doesn't know whether Mary is Italian or French because Bob knows whether Mary is Italian, but he doesn't know whether she is French.

Example (43-a) is a clear case of symmetry cancellation and could be truly used in a situation where Bob knows that Mary is Italian, but he wonders whether she is French as well. If we assume the symmetry presupposition, all theories in (39) make the same predictions about these cases. Examples (42) and (43-b) are predicted to be undefined (they violate symmetry), and (43-a), represented using the denial operator #, is predicted to be true (the sentence embedded by # is undefined). Our intuitions about these sentences, however, are more subtle than that. Compare (42) with (43-b). Example (42) is much worse than simply odd, it is contradictory; (43-b), instead, is a pragmatically marked way of saying something true. Let us check what would be the predictions of our semantic theories in the absence of the symmetry presupposition. Assume Mary is Italian and not French. Then the *exactly one* presupposition is satisfied in this case. SA and PA would give us the right predictions then: (42) would be false in this situation, and (43-b) would be true. WA, and Karttunen's analysis, instead, which rely on symmetry, would give us the wrong predictions: (42) would be true in this situation, and (43-b) would be false.

Both constructions are odd here (because they violate symmetry) but still we seem to have semantic intuitions about them. Our partition semantics and SA allow us to explain these intuitions; a weaker semantic analysis would fail in these cases.

At least one. Consider now the following sentence:

- (44) Bob knows whether Mary is Italian or French.

Suppose Mary is German. Then, according to KA, (44) is true irrespective of Bob's actual belief state.¹⁰ According to SA and WA, instead, (44) is false, again, irrespective to Bob's actual belief state. This cannot be correct. Suppose Bob believes that Mary is French, then (44) is a pragmatically marked way of saying something false, but if he knows that she is German, then (44) is a pragmatically marked way of saying something true. For these cases, only PA gives us the right predictions.

At most one. Consider now the following sentence, an overt cancellation of the *at most* inference:

- (45) Bob knows whether Mary is Italian, French or both.

Suppose Mary is both Italian and French. SA and WA, that rely on the *at most* presupposition, make the wrong predictions for this case. A natural extension of SA and WA to the case of three disjuncts is the following:

- a. SA: $K(a \wedge \neg b \wedge \neg c) \vee K(\neg a \wedge b \wedge \neg c) \vee K(\neg a \wedge \neg b \wedge c)$
 b. WA: $Ka \vee Kb \vee Kc$

¹⁰ Karttunen himself was aware of this shortcoming of the simplified version of his theory. See footnote 1.

But then, substituting $(a \wedge b)$ for c to mean “both”, on SA (45) is false in the situation described. WA is again too weak. Suppose Bob knows that Mary is Italian, but not that she is French. In this case, (45) is false, but WA predicts it to be true. KA and PA, instead, give the correct analysis here.

To conclude, although the sentences considered in this section may be pragmatically marked, we still seem to have semantic intuitions about them. PA is the only semantic analysis among those considered in this article which can account for all these intuitions.

6.2.4 A stronger presupposition

In the remainder of this section we would like to discuss an alternative analysis of the presupposition of knowing whether A or B. Instead of the combination of the symmetry and the *exactly one* presuppositions, we could have assumed just one stronger presupposition:

- (46) **Strong presupposition:**
- a. Knowing whether A or B *presupposes*
 - b. Knowing (A or B) and knowing not (A and B)

By accommodation of only (46) we would have accounted for Schaffer’s “true” judgment, just because the strong presupposition is stronger than the conjunction of the other two presuppositions:

$$(47) K(a \vee b) \wedge K\neg(a \wedge b) \vdash (a \vee b) \wedge \neg(a \wedge b) \wedge K?a \leftrightarrow K?b$$

But we have two arguments against assuming the strong presupposition. One is of an empirical nature, the other is more theoretical.

Empirical argument. Suppose you wonder whether Mary is French or Italian. You know she is one of the two, and you also know that Bob wrongly believes that she is either German or both Italian and French. Now consider the following sentence:

- (48) Bob doesn’t know whether Mary is French or Italian.

According to our theory your utterance would be felicitous here. Both the symmetry and the exactly one presuppositions are satisfied. According to the strong presupposition theory, instead, (48) is infelicitous in the described situation. Intuitions are subtle here. But we believe the sentence is felicitous. The strong presupposition theory is too strong.

Theoretical argument. We have independent evidence for our three pragmatic inferences. All three can be related to more general principles ruling the use of disjunction. The symmetry presupposition can be seen as a particular case of Chemla’s (2008) principle of epistemic similarity for disjunction, according to which a sentence $\phi(a \vee b)$ is felicitous provided the speaker believes $\phi(a)$ if and only if the speaker believes $\phi(b)$. The *at most* inference seems very similar to the exclusive readings of plain disjunction: these are generalized pragmatic effects, but still easy to cancel (cf. example (45)). Lastly, the *at least* inference relates to what Zimmermann (2000) refers to as exhaustive closure of possibilities triggered by *closed* disjunction uses. (Zimmermann, 2000, p. 267) distinguishes between closed and open disjunctions:

- Closed disjunction: A or B [low phrase-final tone] \mapsto the choice is between A or B only, therefore *at least* one of A or B must obtain.
- Open disjunction: A or B [high phrase-final tone] \mapsto the choice is between A or B or something else, no *at least* inference obtains.

Closed disjunctions end with a low phrase-final tone and claim to cover the space of all possibilities. *Open* disjunctions, instead, end on a high phrase-final tone and express the possibility of each disjunct without making any claim of completeness. Compare

- (49) a. Do you want COFFEE or TEA [low phrase-final tone] ?
 b. Do you want COFFEE or TEA [high phrase-final tone] ?

Only (49-a) triggers an *at least* inference, suggesting perhaps that such an inference should not be encoded in the semantics of alternative questions, but possibly in the interpretation of the final falling contour.

To conclude, we have independent evidence for our three pragmatic inferences. The strong presupposition instead is hard to relate to more general principles. Furthermore assuming the strong presupposition would not leave space to account for the possibility that the symmetry presupposition is of a different nature to the other two pragmatic inferences. As the last discussion seems to suggest, it is possible that only the first is a genuine case of presupposition; the other two inferences might be just cases of generalized conversational effects.

7 Conclusion

Let us summarize the main results of this paper. We compared different possible analyses of *knowing whether* sentences and showed how to integrate the partitionial analysis of alternative questions within epistemic logic. We examined different ways of making sense of the context-sensitivity of sentences of the form $K?(a \vee_A b)$. As argued in Sect. 4, a surprising fact about such sentences is that the same scenario can support the intuition that one and the same sentence can be true, false, or undefined in principle. However, we have seen that these judgments can be teased apart once we isolate the relevant parameters. By introducing a semantics of knowledge that distinguishes attention and relevance parameters, we saw that Schaffer's judgments can be derived without resorting to a special mechanism of accommodation of the alternatives raised by the question. Independently, we saw how judgments of undefinedness can be related in a precise way to the pragmatics of alternative questions. One thing we have not proposed here is a detailed comparison with the dynamic approach of the problem proposed in Aloni and Égré (2008), where the role of direct questions and the order in which they are posed was used to explain how the subject's attention can widen up. Likewise, we confined our logic to the case of a single agent, leaving open how to extend it to the multi-agent case. Both directions should be pursued in further work. Meanwhile, following Schaffer we hope to have established a fruitful connection between the broad issue of contextualism in epistemology on the one hand, and the articulation between attention and questions on the other.

Acknowledgements We thank E. Chemla, H. van Ditmarsch, F. Roelofsen, B. Spector, and P. Schlenker for critical comments and helpful suggestions, as well as two anonymous reviewers. We also thank W. van der Hoek for his editorial assistance. Part of the research reported here was financially supported by the Dutch Organization for Scientific Research (NWO).

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Aloni, M., Beaver, D., Clark, M., & van Rooij, R. (2007). The dynamics of topic and focus. In M. Aloni, A. Butler, & P. Dekker (Eds.), *Questions in dynamic semantics*, Vol. 17 of *CRiSPI* (pp. 123–145). Amsterdam: Elsevier.
- Aloni, M., & Égré, P. (2008). Alternative questions and knowledge attributions. *The Philosophical Quarterly* (forthcoming).
- Beaver, D. (1997). Presupposition. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language*. Amsterdam: North Holland.
- Chemla, E. (2008). Similarity: Towards a unified account of scalar implicatures, free choice permission and presupposition projection. *Semantics and Pragmatics* (to appear).
- Cornulier, B. de (1982). Sur le sens des questions totales et alternatives. *Langages*, 16(67), 55–109.
- Dretske, F. I. (1970). Epistemic operators. *The Journal of Philosophy*, 67(24), 1007–1023.
- Fagin, R., & Halpern, J. Y. (1988). Belief, awareness and limited reasoning. *Artificial Intelligence*, 34, 39–76.
- Groenendijk, J., & Stokhof, M. (1982). Semantic analysis of Wh-complements. *Linguistics and Philosophy*, 5(2), 175–233.
- Groenendijk, J., & Stokhof, M. (1984). Studies in the semantics of questions and the pragmatics of answers. Ph.D. thesis, Universiteit van Amsterdam.
- Heim, I. (1994). Interrogative semantics and Karttunen's semantics for *know*. In R. Buchalla & A. Mittwoch (Eds.), *Proceedings of IATL 1* (pp. 128–144). Jerusalem: Akademon.
- Higginbotham, J. (1991). Either/or. In T. Sherer (Ed.), *Proceedings of NELS 21* (pp. 143–155). Amherst, Ma: GLSA.
- Karttunen, L. (1977). Syntax and semantics of questions. *Linguistics and Philosophy*, 1, 3–44.
- Larson, R. K. (1985). On the syntax of disjunction scope. *Natural Language and Linguistic Theory*, 3, 217–264.
- Lewis, D. K. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8, 339–359.
- Lewis, D. K. (1982). 'Whether' report. In T. Pauli (Ed.), *Philosophical essays dedicated to Lennart Åqvist on his fiftieth birthday*. Filosofiska Studier.
- Lewis, D. K. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, LXXV(2), 383–403.
- Stalnaker, R. (1978). Assertion. In R. Stalnaker (Ed.), *Context and content* (pp. 78–95). Oxford: Oxford University Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- von Fintel, K. (1999). NPI-licensing, Strawson-entailment, and context-dependency. *Journal of Semantics*, 16(2), 97–148.
- von Stechow, A. (1991). Focusing and background operators. In W. Abraham (Ed.), *Discourse particles*, Vol. 12 of *Pragmatics and beyond* (pp. 37–84). Amsterdam: John Benjamins Publishing Company.
- Zimmermann, T. E. (2000). Free choice disjunction and epistemic possibility. *Natural Language Semantics*, 8, 225–290.