In: Graham Katz, Sabine Reinhard, and Philip Reuter, eds. (2002), *Sinn & Bedeutung VI, Proceedings of the Sixth Annual Meeting of the Gesellschaft für Semantik,* University of Osnabrück

ISSUES AND AWARENESS

Joris Hulstijn
Vrije Universiteit, Amsterdam
joris@cs.vu.nl

Abstract

In this paper we propose to use the structure provided by issues in a dialogue context to model the notion of *awareness*, which can then be used to account for (some aspects of) the *problem of deduction.* The resulting properties of awareness correspond to intuitions in the literature.

1 Introduction

In recent theories of questions and answers, a dialogue for cooperative information exchange can be specified as a process of *raising* and *resolving* issues (Groenendijk 1999). Just like a proposition models the content of an assertive utterance, a so called *issue* models the content of a question. And just like we accumulate factual information in dialogue, we keep track of the *current issues* or *questions under discussion* (Ginzburg 1995). The issues in a dialogue context provide structure: they indicate how detailed the distinction among possible worlds must be for the current purposes of the dialogue. Issues can be raised explicitly by asking a question or by the use of contrastive focus, or implicitly by introducing a new topic, which triggers issues about the identity of the topic, or by inference from the current task (Hulstijn 2000).

One of the drawbacks of a possible worlds semantics for knowledge or belief is the so called *problem of deduction* (Stalnaker 1984). In some versions, it is known as the problem of logical omniscience (Fagin et al. 1995, ch 9). Because the set of beliefs of an agent is closed under deduction, the agent will automatically believe all consequences of its beliefs. The problem is then how to explain that the conclusion of a deductive reasoning process can be informative. We can solve this dilemma if we drop the assumption that agents are perfect reasoners; they are considered to be *resource bounded*. Reasoning steps take time and effort. If these resources are insufficient, the agent will not have reached the conclusion by the time the resources run out. This approach is taken in recent theories of belief revision (Wasserman 1999). However, the approach requires an explicit mental computation model of agents.

A more general solution makes use of the concept of *awareness*. An agent is aware of some knowledge or belief, when its truth is actively considered. Consider the following example. You know that Clinton is taller than me. After all, you can see that I am of average height, and from television footage you know that Clinton is tall. So in none of the worlds you consider possible I will be taller than Clinton. So you know that Clinton is taller than me, but you only became aware of it, because the issue was brought up just now! Awareness can be formalised in a modal logic of belief, which separates explicit from implicit belief (Fagin and Halpern 1988). As suggested by Fagin and Halpern, awareness can be accounted for by different sources. They mention awareness by perception (Barwise and Perry 1983), versions of resource boundedness (Simon 1955; Wasserman 1999), distributed awareness among multiple agents with divergent areas of expertise (Fagin et al. 1995) or different contextually dependent domain models (Levesque 1984).

In this paper we propose to use the issues in a dialogue context to account for awareness. In this way, we can explain how raising an issue will make agents aware of implicit beliefs. Additionally, it may explain why asking a question can be relevant, even though the corresponding issue could already be inferred from the general task or dialogue topic. Moreover, it provides a semantics for attitudes like wonder or doubt, that have an issue as their complement. The paper is therefore structured as follows. After explaining the idea of raising and resolving issues in section 2, we define a logic for issues and awareness in section 3. The logical properties of the resulting notion of awareness are examined in section 4. The adequacy of these properties is discussed in section 5.

2 Raising and resolving issues

Before we can explain awareness by issues, we must first explain how issues are raised. We make the following assumptions. The dialogue context is modelled in a very abstract way: by the information contained in it. The context resembles a version of the common ground, established on the basis of public information about the domain, the task and about what was said. One can think of the context as the information state of an objective observer monitoring the dialogue. We abstract over differences between the information states of speaker and addressee. Crucially, the dialogue context consists of *data* structured by *issues*. The data corresponds to the facts known by the dialogue participants. Structure comes in many forms. Surface aspects of utterances like word order and intonation convey how new information relates to information already present in the dialogue context. Also the topics addressed in a dialogue and the apparent task of the speaker can trigger issues to be raised. The main idea of the model is that these different types of structure can be captured by a single mechanism, issues, because they share a distinction among possible alternative ways of resolving the issue.

First, an issue can be expressed by an explicit question. We have strong intuitions about what counts as an answer to a question. Each answer is an alternative way of resolving the issue. A question is associated with a question relation: an expression that specifies a 'gap' in the information of the asker. The answerer is invited to fill that gap. For example, the question relation of 'Whom does John love?' is given by the following lambda abstract: $\lambda x(person(x) \wedge love(john,x))$. In this case any description of a set of persons loved by John would count as an answer: 'nobody', 'Mary', 'Mary and Sue' or 'the girls with red hair'. The dialogue act of asking a question requires some sort of response. 'I don't know' or 'Why do you ask?' are not genuine answers, but are possible responses.

Second, the surface form and intonation of an utterance often indicate what a speaker takes to be the relevant issue. Rooth's (1992) alternatives semantics for focus is similar to a semantics of questions and answers. The element in focus is used to resolve an issue, which is presupposed by the word order or intonation contour. For example, if a child complains 'No, I don't want a *red* bike!' the apparent issue in that context is the colour of the bike. The child selects the answer 'not red' from other alternatives.

Third, dialogue topics provide a different kind of structure. The topic of a dialogue segment specifies the object, possibly abstract, that the segment 'is about'. A topic suggests a hierarchical topic space of related objects and attributes. The value of an attribute may not always be known. As new topics are introduced, unknown attribute values trigger new issues to be discussed. Consider the opening of a fairy tale: 'Once upon a time there was a little prince.' This makes us curious. Who is this prince? What is going to happen to him? Exactly which issues are raised depends on the genre of the text or the dialogue type.

Finally, for task oriented dialogues the dialogue structure depends partly on the task structure. The task domain defines which objects are available as potential topics and how they are related. The underlying activity and the goals of the dialogue participants determine what would count as a coherent contribution to the dialogue. Participants consider which plans or actions are available to accomplish the task. In order to choose the most appropriate action often more information is needed. Different actions have different preconditions. So a decision problem as to which action to take, triggers an issue about the truth of the preconditions. Often such issues will then be asked explicitly. For task oriented dialogues we can explain when asking a question is relevant (Hulstijn 2000).

Thus we can picture a dialogue for information exchange as a process of raising and resolving issues. Obviously, issues may also be 'lowered': as the task progresses or as the topic shifts, other issues may become relevant. Ginzburg (1995) and subsequent implementations of his theory in the Trindi project (Traum et al. 2000), therefore treat the 'questions under discussion' as a data-structure like a stack or a prioritised queue. Operations on this queue provide a (very crude) model of cognitive processing during dialogue. Typically, issues which are not re-activated will after some time be replaced by more salient issues. In this paper however we assume that issues are either raised or not. Once raised the only way to get rid of them, is to resolve them by factual information.

3 A Logic for issues and awareness

The logic combines a modal logic of explicit and implicit belief (Fagin and Halpern 1988) with an update semantics of questions and answers (Groenendijk 1999; Hulstijn 2000). The semantics makes a distinction between *propositions*, expressing the content of an assertive utterance or attitude, and so called *issues*, representing the content of an interrogative utterance or attitude. Just like propositions are modelled by a set of possible worlds, issues are modelled by either a partition of a set of possible worlds, or equivalently by an equivalence relation over a set of possible worlds. The intuition is that each 'block' in the partition corresponds to an alternative way of resolving the issue.

Because we are interested in the logical properties of awareness, we will restrict to a propositional language in this paper. The language is defined in three stages. Language $\mathcal{L}_0$ extends a propositional language with modal operators L for implicit belief, B for explicit belief, A for awareness and W for the attitude of wondering. Language $\mathcal{L}_1$ introduces an operator '?' for interrogatives. Note that W takes an issue as its complement. The operators L, B and A express assertive attitudes, which take a proposition. Sequences of assertive or interrogative utterances are combined by sequential conjunction ';'.

Definition 1 (Syntax $\mathcal{L}_0, \mathcal{L}_1, \mathcal{L}$)
Given a vocabulary $\mathcal{P} = \{p, q, ...\}$ define

$\mathcal{L}_0:$	$\varphi ::= \ p \mid \neg\varphi \mid (\varphi \vee \psi) \mid (\varphi \to \psi) \mid (\varphi \wedge \psi) \mid \mathsf{L}\varphi \mid \mathsf{A}\varphi \mid \mathsf{B}\varphi \mid \mathsf{W}\chi$	(proposition)
$\mathcal{L}_1:$	$\chi ::= \ ?\varphi$	(issue)
$\mathcal{L}:$	$\zeta ::= \ \varphi \mid \chi \mid (\zeta;\eta) \qquad (p \in \mathcal{P};\ \varphi, \psi \in \mathcal{L}_0;\ \chi \in \mathcal{L}_1;\ \zeta, \eta \in \mathcal{L})$	(sequence)

The semantics of $\mathcal{L}_0$ is based on a Kripke model $M = \langle W, \mathcal{R}, V \rangle$ where W is a set of possible worlds and V a valuation function. Since differences between worlds must be expressible in the language, we identify each world w with an assignment $V(w)$ of truth values to atomic expressions from $\mathcal{P}$. The semantics of $\wedge$ and $\neg$ are standard. The semantics of $\vee$ and $\to$ can be defined in terms of these: $(\varphi \vee \psi) \equiv \neg(\neg\varphi \wedge \neg\psi)$ and $(\varphi \to \psi) \equiv \neg(\varphi \wedge \neg\psi)$. We assume a transitive, serial and Euclidean epistemic accessibility function $\mathcal{R}$ to model implicit belief of some agent. This gives a logic of KD45 for operator L. The semantics of B, A and W will be dealt in the subsequent sections[1].

Definition 2 (Satisfaction ($\models$))
Given a model M as above, define for each $\varphi, \psi \in \mathcal{L}_0$

$M, w \models p$	iff	$V(w)(p) = 1$	$(p \in \mathcal{P})$
$M, w \models \neg\varphi$	iff	$M, w \not\models \varphi$	
$M, w \models (\varphi \wedge \psi)$	iff	$M, w \models \varphi$ and $M, w \models \psi$	
$M, w \models \mathsf{L}\varphi$	iff	$M, v \models \varphi$, for all $v \ \langle w, v \rangle \in \mathcal{R}$	

If needed, the language can be extended to a predicate logical language with substitution. The vocabulary $\mathcal{P}$ contains sentences $R(c_1, ..., c_n)$, where R is a predicate symbol and c_i a constant. Now $?love(john, mary)$ expresses the issue whether John loves Mary; $\mathsf{B}taller(clinton, author)$ expresses the explicit belief that Clinton is taller than me; $\mathsf{W}?taller(clinton, author)$ that the agent is wondering whether Clinton is taller than me, and $\mathsf{A}taller(clinton, official)$ that the agent is aware that Clinton is taller than some official. Add variables and clauses $\exists x\varphi$ and $\forall x\varphi$ to $\mathcal{L}_0$. Formulas of the form $?\vec{x}\varphi$ are added to $\mathcal{L}_1$ to express wh-issues, where $\vec{x}$ is short for a sequence of zero or more distinct variables. For example $?xy(love(x,y))$ expresses the issue 'who loves whom'. Definition 2 is extended with the clause: $M, w \models \exists x(\varphi)$ iff $M, w \models \varphi\{a/x\}$, substituting x for some constant a. As usual we define $\forall x(\varphi) \equiv \neg\exists x(\neg\varphi)$. Think of a world w combined with a substitution θ as a so called 'epistemic possibility' w_θ, which plays the same role as a possible world in the framework.

3.1 Information states

In general, an update semantics makes use of *information states* to model the information of an agent. As the dialogue progresses, new information is added. Information states $\sigma = \langle S, I \rangle$ monitor two aspects of information: data and structure. The *data set S* contains the worlds considered possible by the agent given its current beliefs. The *issue structure* is represented by an equivalence relation I. Two worlds are I-related whenever they are indistinguishable with respect to the facts that are currently 'at issue'. In other words, they agree on all their answers to the 'questions under discussion'. Initially, all worlds are indistinguishable. As more issues are raised, the sets of indistinguishable worlds or 'blocks' in the partition become more constrained.

Formally, an *update system* is a tuple $\langle \Sigma, \mathbf{0}, \mathbf{1}, \sqsubseteq, [.] \rangle$. Here Σ is a set of information states. The initial information state $\mathbf{0}$ expresses total ignorance. It consists of the whole logical space W as data set and the universal relation over W as issue structure. The inconsistent or absurd information state $\mathbf{1}$ is modelled by an empty data set: no worlds are compatible with inconsistent information. It's issue structure does not matter: each equivalence relation over W is allowed.

Definition 3 (Information states ($\Sigma, \mathbf{0}, \mathbf{1}$))
Given a set of possible worlds W, define

Σ	$=$	$\{\langle S, I \rangle \mid S \subseteq W;\ I \subseteq W \times W$ refl, sym, trans$\}$	(information states)
$\mathbf{0}$	$=$	$\langle W, W \times W \rangle$	(initial state)
$\mathbf{1}$	$=$	$\langle \varnothing, R \rangle$, for any $R \subseteq W \times W$ refl, sym, trans	(absurd state)

[1] If needed, we can model the attitudes of separate agents by distinguishing for example $\mathsf{L}_a, \mathsf{A}_a, \mathsf{L}_b$ etc, for agents a, b etc. These are then related to specific accessibility relations $\mathcal{R}_a, \mathcal{R}_b$ etc. in the model.

The relation $\sqsubseteq$ represents an information order on states: $\sigma \sqsubseteq \tau$ whenever τ contains at least as much information as σ. Because more information means less possibilities or less indistinguishable pairs, the information order is modelled by the subset relation on both S and I.

Definition 4 (Information order ($\sqsubseteq$))
For $\langle S_1, I_1\rangle, \langle S_2, I_2\rangle \in \Sigma$ define $\langle S_1, I_1\rangle \sqsubseteq \langle S_2, I_2\rangle$ iff $S_1 \supseteq S_2$ and $I_1 \supseteq I_2$.

The update function [.] maps information states and formulas onto new information states. We write $\sigma[\zeta]$ for the result of updating σ with a formula ζ. An update with a proposition φ restricts S to those worlds that are compatible with φ. An update with an issue $?\varphi$ on the other hand eliminates all pairs of worlds that do not agree on their answer to the question 'whether φ'. Because of the use of 'iff', relation I nicely remains an equivalence relation.

Definition 5 (Update function ([.]))
For each $\sigma = \langle S, I\rangle \in \Sigma, \psi \in \mathcal{L}_0, ?\varphi \in L_1$ and $\zeta, \eta \in \mathcal{L}$ define

$$\begin{array}{lcl}
\langle S, I\rangle[\psi] & = & \langle \{v \in S \mid M, v \models \psi\}, I\rangle \\
\langle S, I\rangle[?\varphi] & = & \langle S, \{\langle v, w\rangle \in I \mid M, v \models \varphi \text{ iff } M, w \models \varphi\}\rangle \\
\langle S, I\rangle[\zeta;\eta] & = & \langle S, I\rangle[\zeta][\eta]
\end{array}$$

Note that propositions do not affect the issue structure I, and that issues do not affect the information set S. In other versions of the logic (Groenendijk 1999; Hulstijn 2000) both aspects are modelled by a single equivalence relation I, which remains restricted to worlds in S. However, in this paper we want to separate an agent's private beliefs S from public issues I. An issue may remain under discussion, even though a particular agent has already resolved it. It is useful to have a notation for the set of worlds that satisfies a proposition. Given a model M, define $[\![\psi]\!]_M = \{v \mid M, v \models \psi\}$. For issues we define $[\![?\varphi]\!]_M = \{\langle v, w\rangle \mid M, v \models \varphi \text{ iff } M, w \models \varphi\}$, the *indistinguishability relation* induced by $?\varphi$. Reference to a model M will often be dropped.

In a predicate logic version, the clause for issues would be: $\langle S, I\rangle[?\vec{x}\varphi] = \{S, \{\langle v, w\rangle \in I \mid M, v \models \varphi\{\vec{a}/\vec{x}\} \text{ iff } M, w \models \varphi\{\vec{a}/\vec{x}\}$, for all $\vec{a}$ such that $M, u \models \varphi\{\vec{a}/\vec{x}\}$ with $u \in S$. It eliminates all pairs of worlds which do not agree on the issue, using substitutions which are considered possible with respect to S. We can also define the indistinguishability relation over epistemic possibilities w_θ. In that case we get blocks corresponding to different values of variables.

As is clear from definitions 4 and 5 each update results in an increase of information. Revisions or contractions have not been analysed in this model. Typically, it holds that $\mathbf{0} \sqsubseteq \sigma \sqsubseteq \sigma[\eta] \sqsubseteq \mathbf{1}$ for arbitrary $\sigma \in \Sigma$ and $\eta \in \mathcal{L}$. This can be shown by induction. For propositions $S \cap [\![\psi]\!] \subseteq S$ and for issues $I \cap [\![?\varphi]\!] \subseteq I$. Sequential conjunction is then dealt with by the induction hypothesis. It indicates that this update semantics for issues does not differ much from the logic of questions that can be found in Groenendijk and Stokhof (1996).

3.2 Support and entailment

For each update system there exists a *support relation* '$\Vdash$' stating which formulas are 'supported' or 'accepted' by an information state. It is defined by means of a fixed point: a formula is supported by an information state, whenever an update with it will not increase the information. Note that this holds for interrogatives and assertions alike.

Definition 6 (Support ($\Vdash$))
Information state σ supports utterance $\zeta \in \mathcal{L}$, $\sigma \Vdash \zeta$, whenever $\sigma[\zeta] \sqsubseteq \sigma$.

Because $\sqsubseteq$ is anti-symmetric, the 'increase of information' result discussed above makes the definition stronger. In fact $\sigma \Vdash \zeta$ iff $\sigma[\zeta] = \sigma$.

We define a non-standard notion of entailment based on support. A sequence of utterances conforms to a valid argument, when, after having applied the premises in the right order to an arbitrary information state, the conclusion is supported.

Definition 7 (Entailment ($\models$))
$\zeta_1;...;\zeta_n \models \eta$ iff $\sigma[\zeta_1]...[\zeta_n] \Vdash \eta$, for any $\sigma \in \Sigma$.

Because the notion of support works both for interrogatives and assertions, the same holds for entailment. Issues can appear both as premises or as conclusion. This gives four possible combinations. First, a proposition entailing another proposition is the usual case. In this version of the logic, entailment behaves as in proposition logic. So $\varphi \models \psi$ when $[\![\psi]\!]_M \subseteq [\![\varphi]\!]_M$. Second, an issue entails another issue, $?\varphi \models ?\psi$, when each solution to $?\varphi$ counts as a solution to $?\psi$. By definition 4, 5 this means that $[\![?\psi]\!]_M \subseteq [\![?\varphi]\!]_M$. In other words, the partition induced by $?\varphi$ is the same or more fine-grained than that induced by $?\psi$. Example (1) illustrates such a logical dependency between issues. The question whether to bring an umbrella or not, depends on what the weather forecast is like: $?x(\mathit{forecast}(x)) \models ?\mathit{bring}(\mathit{umbrella})$. Therefore B's question is considered a relevant response. Although it does not resolve A's issue, it does contribute to it.

(1) A: Shall I bring an umbrella?
B: What was the forecast like?

Third, in the logic an issue cannot entail a proposition. In definition 5 the effect of interrogatives is confined to the issue structure and propositions only affect the data. In natural language however, interrogative utterances may add information by means of presupposition accommodation. All definite expressions mentioned in the question, as well as presuppositions induced by factual verbs or other lexical connotations restrict the data set S. Presupposed material, indicated by the ∂ symbol, can be dealt with by a special kind of 'cautious' update (Beaver 1996). It only takes place in case the resulting information state remains consistent: $\sigma[\partial\psi] = \sigma[\psi]$ if $\sigma[\psi] \neq \mathbf{1}$, undefined otherwise. Moreover, in a predicate logic version, presuppositions may bind variables, adding their possible values to the information state.

(2) Is the king of France bald? $\partial(\ \exists! x.\mathit{king\ of\ france}(x)\ ); ?\mathit{bald}(x)$

Fourth, a proposition cannot entail an issue. However, assertive utterances do trigger issues to be raised. We discussed focus-ground constructions, the introduction of a topic or a task, which may introduce or presuppose issues. Such expressions must be translated to explicit issues in the logical language. Thus, both propositions and issues may be presupposed (Groenendijk 1999). The child's utterance discussed in the introduction can for example be translated as follows.

(3) I don't want a *red* bike! $\partial(\ \exists x\ \mathit{bike}(x) \wedge \mathit{want}(x); ?y\ \mathit{colour}(x,y)\ ); \neg(y = \mathit{red})$.

3.3 Resolution

We model a dialogue as a process of raising and resolving issues. Raising issues is modelled by interrogative updates. But what about resolution? Earlier, we used the metaphor of an issue as a set of alternatives. Now roughly an issue is said to be resolved, when only one alternative remains open. An issue is partially resolved, when at least one alternative is removed, as a result of considering the information. To express this we introduce some notation. An equivalence relation R over some set A generates equivalence classes $[a]_R$ for $a \in A$, which form a partition called the quotient set $A/R = \{[a]_R \mid a \in A\}$.

An information state σ is said to resolve an issue $?\varphi$, written $\sigma \Vdash ?\varphi$, when its data set entails one of the equivalence classes induced by the issue. Resolution can also be taken as the basis for a kind of entailment. A sequence of propositions $\psi_1;\ldots;\psi_n$ is said to resolve $?\varphi$, written as $\psi_1;\ldots;\psi_n \models\!\!\equiv ?\varphi$ when an arbitrary σ updated with $\psi_1;\ldots;\psi_n$ resolves $?\varphi$ in the earlier sense. In practice we will use the fact that $\psi \models\!\!\equiv ?\varphi$ iff $[\![\psi]\!]_M \subseteq W/[\![?\varphi]\!]_M$. Note that for all tautologies and inconsistencies ψ we have $\models\!\!\equiv ?\psi$, because trivial issues, which separate possible from logically impossible worlds, are resolved by any answer.

Definition 8 (Resolution ($\Vdash$))
Information state $\sigma = \langle S, I\rangle$ resolves $?\varphi$, $\sigma \Vdash ?\varphi$, iff $S \subseteq V$ for some $V \in W/[\![?\varphi]\!]$.

Definition 9 (Resolution ($\models\!\!\equiv$))
$\psi_1;\ldots;\psi_n \models\!\!\equiv ?\varphi$ iff $\sigma[\psi_1]\ldots[\psi_n] \Vdash ?\varphi$, for arbitrary $\sigma \in \Sigma$.

An issue is said to be partially resolved in some information state, when at least one of the blocks in the partition is eliminated by the data set. This notion can be used to define a technical account of relevance, as in definition 11 below. A proposition is called relevant in context, when it partially resolves some non-trivial issue in the context (Hulstijn 2000). As we discussed above, an issue may be called relevant, when there is a logical dependency relation that relates it to some existing issue in the context, which will indirectly help reduce the partition. If the weather forecast is already known in example (1), A's question is answered. If it is not known, B's question gives a hint where to look for the answer. Note that such a response only makes sense if the participants are not yet aware of this dependency.

Definition 10 (Partial Resolution)
Information state $\sigma = \langle S, I\rangle$ partially resolves $?\varphi$ when for some $U \in W/[\![?\varphi]\!]$, there is no $V \in S/[\![?\varphi]\!]$ such that $V \subseteq U$.

Definition 11 (Relevance)
Proposition ψ is relevant to issue $?\varphi$, when $\sigma[\psi]$ partially resolves $?\varphi$, for arbitrary $\sigma \in \Sigma$.
Issue $?\psi$ is relevant to issue $?\varphi$, when for all ψ' such that $\psi' \models\!\!\equiv ?\psi$, $\sigma[\psi']$ partially resolves $?\varphi$, for $\sigma \in \Sigma$.

According to the above notion, any inconsistent proposition is relevant: it reduces the partition. There is nothing to stop the speaker from being over-informative. Groenendijk (1999) defines a constraint that does just that: licensing. An utterance is licensed when it does not provide more information than is required for current purposes. This corresponds to the second clause of Grice's maxim of quantity (Grice 1975, p 45). How much information is required, is specified by contextual issues. This means that updates should proceed 'block wise'. If some worlds are eliminated, all worlds that are indistinguishable with respect to the current issues, should also be eliminated.

Definition 12 (License)
Proposition ψ is licensed with respect to issue $?\varphi$, when for all $\sigma = \langle S, I\rangle \in \Sigma$ such that $\sigma[\psi] = \langle S', I'\rangle$ it holds that if $v \in S$ then for all w such that $\langle v, w\rangle \in [\![?\varphi]\!]$, $v \in S'$ iff $w \in S'$.

Now you may ask if issues can or must be licensed too? We might say that an issue is licensed in case it is relevant to some general issue in the context, and its typical answers are licensed. However, that would mean that given the initial dialogue topic or issue, no other issues may be raised. This is only true of a very restricted set of dialogue types, such as interrogations or examinations.

3.4 Implicit belief

To continue with our account of awareness we must link implicit belief and information states. The link is based on the following observation. Because of the KD45 properties of $\mathcal{R}$, the set of worlds that are accessible from some world w by $\mathcal{R}$ forms a set $S(w)$ (Fagin et al. 1995, p 62). This set models the implicit beliefs of our observer agent in world w, which can act as the data set of an information state. Similarly, the issues of our agent in w are modelled by an indistinguishability relation $I(w)$. From now on a model M is defined as $M = \langle W, \sigma, V \rangle$, where for each w we have $\sigma(w) = \langle S(w), I(w) \rangle$, such that $S(w) = \{v \mid \langle w, v \rangle \in \mathcal{R}\}$ and $I(w)$ is an equivalence relation over W.

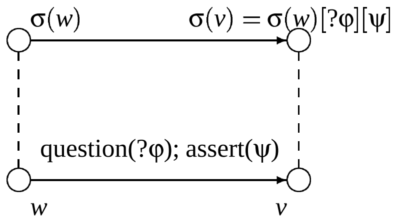

Figure 1: Utterances and information state updates

As the dialogue progresses the world changes. This is depicted in figure 1. After an sequence of events, for example a question $?\varphi$ followed by an assertion ψ, the information state of an agent in world w must be updated accordingly. So $\sigma(w)$ becomes $\sigma(v) = \sigma(w)[?\varphi][\psi]$. We can now show that the notion of implicit belief and of support in update semantics coincide.

Proposition 1 (Implicit belief and support)
For each $\varphi \in \mathcal{L}_0$ the following holds: $M, w \models \mathsf{L}\varphi$ iff $\sigma(w) \Vdash \varphi$.
Proof: Suppose $M, w \models \mathsf{L}\varphi$, so $M, v \models \varphi$ for all v such that $\langle w, v \rangle \in \mathcal{R}$ (def 2), which are in fact all $v \in S(w)$. Therefore $\{v \in S(w) \mid M, v \models \varphi\} = S(w)$ itself, and because φ only affects the data (def 5) $\langle S(w), I(w) \rangle[\varphi] = \langle S(w), I(w) \rangle = \sigma(w)$ itself, so $\sigma(w) \Vdash \varphi$ (def 6).

We assume that models M guarantee correct updates. That means that agents are assumed to be cooperative in the sense that they take over all issues raised by other dialogue participants, and trusting in the sense that they update with all propositions asserted or implied by assertions of the other participants. Obviously, to model dialogue proper we would need to distinguish between speakers, addressees and overhearers, and make sure addressees only update if the information is compatible with what they knew or otherwise revise, and if so indicate this by some positive acknowledgement. Speakers or overhearers only update their version of the common ground if the utterance has been *grounded* by means of some acknowledgement from the addressee (Clark and Schaefer 1989). We realise that in general, correctness is too strong an assumption to make. Typically, dialogue participants display strategic behaviour. They may have private reasons for only pretending to accept what was said, without actually updating their information states. However, we believe that the correctness assumption does make sense for grounded utterances observed by an objective observer, who models what participants ought to do in a cooperative information exchange. For other participating agents or for other types of dialogue like debates or negotiations, these assumptions must be replaced by weaker update procedures. McBurney and Parsons (to appear) use such update rules to model the commitments expressed by utterances of participants in a debate.

3.5 Explicit belief

Following Fagin and Halpern (1988) we characterise the notion of explicit belief as implicit belief plus awareness.

Definition 13 (Explicit Belief (B))
For each $w \in W, \varphi \in \mathcal{L}_0$ define $M, w \models \mathsf{B}\varphi$ iff $M, w \models \mathsf{L}\varphi$ and $M, w \models \mathsf{A}\varphi$

This leaves us to characterise awareness. To study the properties of awareness, Fagin and Halpern introduce an *awareness set* $\mathcal{A}(w)$, which contains the formulas that an agent is actively aware of in w. The A operator can then be defined as follows: $M, w \models \mathsf{A}\varphi$ iff $\varphi \in \mathcal{A}(w)$. It is quite crucial to their approach that the awareness set is characterised syntactically.

Rather than mixing syntactic and semantic approaches, Konolige (1986) suggests a purely deductive approach to modelling belief sets. This can be used to study the explicit beliefs of an agent that is resource bounded. The agent may not able to calculate the deductive closure of its beliefs before its resources run out. Such a procedural notion is syntactic by nature. For example, Wasserman (1999) defines a syntactic notion of *relatedness* between formulas. Groups of largely interrelated formulas typically form clusters, or topics. Such clusters often do remain closed under logical consequence. On the other hand, an agent may never realise the logical consequences of combining unrelated parts of its tacit knowledge.

Other notions of awareness can be studied too. The perceptual situations of Barwise and Perry (1983) can be understood as partial possible worlds; the agent is not aware of all the distinctions made in a world. This partiality is motivated by the limited nature of perception. Below we will compare these situations with a suggestion of Levesque (1984), who restricts worlds to a limited vocabulary $\mathcal{P}$, making them partial. Another interesting form of awareness is based on the distribution of knowledge and expertise among agents in a network. Although the agents together might have some information, they are not aware of this until they have communicated the information. See for example Fagin et al. (1995).

In the terminology of Konolige (1986), one could say that the deductive approaches to belief use awareness as a *derivator*. By contrast, we apply awareness as a *filter*; it separates explicit from implicit belief. The structure of issues raised in a dialogue is used as a source of awareness. An agent is said to be aware of proposition φ in case it is considering an issue $?\varphi$[2]. We could define an awareness set in the style of Fagin and Halpern: $\mathcal{A}(w) = \{\varphi \mid \sigma(w) \Vdash ?\varphi\}$. However, since issues are a semantic notion, there is no need for a syntactic awareness set in our approach. Therefore we define the semantics of A directly in terms of issues.

Definition 14 (Awareness (A))
For each $w \in W, \varphi \in \mathcal{L}_0$ define $M, w \models \mathsf{A}\varphi$ iff $\sigma(w) \Vdash ?\varphi$.

The motivation for this approach comes from several sources. Consider again the Clinton example. By asking ‘Is Clinton taller than me?’ I force you to consider the alternative answers (yes, he is taller; no, he is not taller; you are of equal height) or responses (I don’t know). This mechanism shows the feasibility of what might be called a ‘Socratic dialogue’ in which a subject learns or discovers what she knew all along, guided by the questions of a teacher. So at least part of the function of a question is to make agents aware of certain facts. A similar conclusion was made on the basis of example (1), where B’s question indicates awareness of the dependency between the weather forecast and whether to take an umbrella or not.

[2] When an agent is considering $?\varphi$ it will be aware of both the positive and the negative answers, φ and $\neg\varphi$. See proposition 2.

In all the examples discussed so far there is a separation between *background* and *foreground*. Raising an issue brings certain tacit distinctions to the foreground. It is an interesting empirical question how this mechanism may be used to account for linguistic data about the focus-ground distinction. The similarity between the semantics of questions and the alternatives semantics of focus (Rooth 1992) indicates that this is a fruitful way to go. In example (3) the issue of the colour of the bike provides a background against which the scope of the negation can be decided. A different way to apply the foreground-background distinction, is by relating recently raised issues to the topic of the dialogue (Van Kuppevelt 1995). Roughly, the topic can be seen as the object that an utterance, or larger dialogue segment, is *about*. The topic is usually contrasted with the comment: that part of the utterance or segment that raises new issues, or provides new or contrasting information about the topic. We could define the topic therefore as the (set of) objects that re-occur in each of the alternative solutions to an issue. So the topic is what remains invariant over the possible answers to a question. As such, it is similar to the presupposition of a question.

3.6 Wondering and knowledge-wh

As an additional benefit, our approach provides a semantics for attitudes that take issues as their complement, such as considering, wondering or doubting. As a sample we define the attitude of wondering, expressed by W. We say that Sherlock is wondering who did it, when he is actively considering that issue.

Definition 15 (Wondering (W))
For each model M, world w and issue $?\varphi \in \mathcal{L}_1$ $\;M,w \models \mathsf{W}?\varphi$ iff $\sigma(w) \Vdash ?\varphi$

There is a direct relation between awareness and wondering: $M,w \models \mathsf{W}?\varphi$ iff $M,w \models \mathsf{A}\varphi$. Also other attitudes can be accounted for in this way. Attitudes like 'knowing who, what or whether' are called *resolutive* because their complement needs to be resolved (Ginzburg 1996). We can only say that Sherlock *thinks* he knows who did it, when he has information that resolves the issue: $M,w \models \mathsf{B}?\varphi$ iff $\sigma(w) \mathrel{|\!\!\Vdash} ?\varphi$. To make it knowledge, this information must be correct with respect to the current world of evaluation. If we define explicit knowledge-that by $M,w \models \mathsf{K}\varphi$ iff $M,w \models \mathsf{B}\varphi$ and $M,w \models \varphi$, then explicit knowledge-wh is defined by $M,w \models \mathsf{K}?\varphi$ iff $M,w \models \mathsf{K}\psi$ and $\psi \,|\!\!\equiv ?\varphi$. Because in general both $\varphi \,|\!\!\equiv ?\varphi$ and $\neg\varphi \,|\!\!\equiv ?\varphi$ we can derive that knowledge-that implies knowledge-wh, but not the opposite.

4 Properties

A number of logical properties of awareness can be derived from the nature of issues. The properties discussed in this section are taken from the literature on awareness (Fagin and Halpern 1988; Huang and Kwast 1991). For some types of awareness, they are judged to be desirable.

4.1 Closure under negation

The first property states that awareness must be closed under negation. If $\mathsf{A}\varphi$ then also $\mathsf{A}\neg\varphi$. This property does indeed hold for our issue related notion of awareness. If I am aware of one alternative solution to an issue, should be aware of the other alternatives too.

Proposition 2 (Closure under Negation)
Awareness is closed under negation. If $M,w \models \mathsf{A}\varphi$ then $M,w \models \mathsf{A}\neg\varphi$.
Proof. Suppose $M,w \models \mathsf{A}\varphi$. By definition $\sigma(w) = \langle S(w), I(w)\rangle \Vdash ?\varphi$. Because of definition 5, 6, $I(w) \subseteq \{\langle u,v\rangle | M,u \models \varphi \text{ iff } M,v \models \varphi\}$. Because of 'iff' $I(w) \subseteq \{\langle u,v\rangle | M,u \not\models \varphi \text{ iff } M,v \not\models \varphi\}$ and because of '$\neg$' (def 2) $I(w) \subseteq \{\langle u,v\rangle | M,u \models \neg\varphi \text{ iff } M,v \models \neg\varphi\}$. So we have $\sigma(w) \Vdash ?\neg\varphi$ again, and $M,w \models \mathsf{A}\neg\varphi$ (def 14). □

The proof is not based on characteristics of σ, so we can show that in general $?\varphi \models ?\neg\varphi$ and vice versa. Please note that in natural language a question and its negative counterpart are usually not equivalent. Consider 'Are you married?' and 'Are you unmarried?'. Such formulations have different answers [3].

4.2 Closure under conjunction

Another property for awareness notions is closure under conjunction. If you are aware of two issues separately, you must be aware of their combination.

To shorten the proofs we introduce the following notation: φ_w means $M,w \models \varphi$ and $\overline{\varphi_w}$ means $M,w \not\models \varphi$. Worlds w and v are said to *agree* on φ, written $w \sim_\varphi v$, when $M,w \models \varphi$ iff $M,v \models \varphi$.

Proposition 3 (Closure under Conjunction)
Awareness is closed under conjunction. If $M,w \models \mathsf{A}\varphi$ and $M,w \models \mathsf{A}\psi$ then $M,w \models \mathsf{A}(\varphi\wedge\psi)$.
Proof. Suppose that $M,w \models \mathsf{A}\varphi$ and $M,w \models \mathsf{A}\psi$, so $\sigma(w) \Vdash ?\varphi$ and $\sigma(w) \Vdash ?\psi$ (def 14). This means $u \sim_\varphi v$ and $u \sim_\psi v$ for all $\langle u,v\rangle \in I(w)$ as above (def 5,6). Therefore (φ_u,φ_v) or $(\overline{\varphi_u},\overline{\varphi_v})$ and also (ψ_u,ψ_v) or $(\overline{\psi_u},\overline{\psi_v})$ (def 'iff'). By propositional 'meta' logic, this gives four possible combinations: $(\varphi_u,\varphi_v),(\psi_u,\psi_v)$ or $(\varphi_u,\varphi_v),(\overline{\psi_u},\overline{\psi_v})$ or $(\overline{\varphi_u},\overline{\varphi_v}),(\psi_u,\psi_v)$ or $(\overline{\varphi_u},\overline{\varphi_v}),(\overline{\psi_u},\overline{\psi_v})$. By changing the order we get: $(\varphi_u,\psi_u),(\varphi_v,\psi_v)$ or $(\varphi_u,\overline{\psi_u}),(\varphi_v,\overline{\psi_v})$ or $\underline{(\overline{\varphi_u},\psi_u),(\overline{\varphi_v},\psi_v)}$ or $(\overline{\varphi_u},\overline{\psi_u}),(\overline{\varphi_v},\overline{\psi_v})$. This is in fact the truth condition of $(\varphi\wedge\psi)_u,(\varphi\wedge\psi)_v$ or $\overline{(\varphi\wedge\psi)_u},\overline{(\varphi\wedge\psi)_v}$. So $u \sim_{(\varphi\wedge\psi)} v$ for $\langle u,v\rangle \in I(w)$, so $\sigma(w) \Vdash ?(\varphi\wedge\psi)$ (def 5,6) and $M,w \models \mathsf{A}(\varphi\wedge\psi)$ (def 14). □

A more intuitive explanation is that the equivalence classes resulting from the $?\varphi$ issue and the $?\psi$ issue are combined using set intersection (def 5), producing four possible equivalence classes. This is illustrated in figure 2. The class of (φ,ψ)-worlds corresponds to the positive answer to the question $?(\varphi\wedge\psi)$, while the $(\varphi,\overline{\psi}),(\overline{\varphi},\psi)$ and $(\overline{\varphi},\overline{\psi})$-worlds form the negative answer. In other words, the partition induced by $?\varphi$ and $?\psi$ separately, subsumes the partition of $?(\varphi\wedge\psi)$. Because this argument is independent of σ, we can strengthen it to the general case: $?\varphi;?\psi \models ?(\varphi\wedge\psi)$. From the picture it becomes clear that the opposite is not true. So $?(\varphi\wedge\psi) \not\models ?\varphi$. Therefore our notion of awareness is not decomposable under conjunction.

Proposition 4 (Not Decomposable under Conjunction)
Awareness is *not* decomposable under conjunction. If $M,w \models \mathsf{A}(\varphi\wedge\psi)$ not always $M,w \models \mathsf{A}\varphi$.
Proof. Construct a counter example on the basis of figure 2. Suppose that $M,w \models \mathsf{A}(\varphi\wedge\psi)$, so $u \sim_{\varphi\wedge\psi} v$ for all pairs $\langle u,v\rangle \in I(w)$. Take one such pair, u',v', such that $M,u' \models (\varphi\wedge\neg\psi)$ and $M,v' \models (\neg\varphi\wedge\psi)$. Although $u' \sim_{\varphi\wedge\psi} v'$ (both false), clearly not $u' \sim_\varphi v'$. □

[3] This is reason for Groenendijk and Stokhof (1996) to reject an extensional version of the logic for questions. The logic does not meet their adequacy criteria.

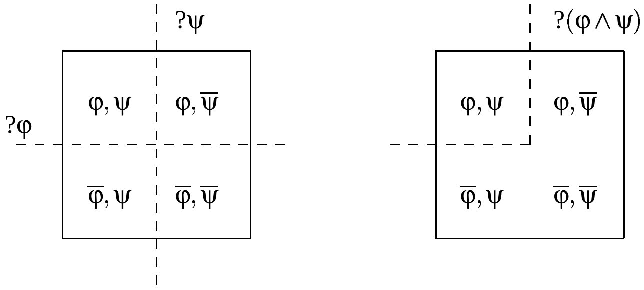

Figure 2: Issue structures of $?\varphi$ and $?\psi$ combined, and of $?(\varphi \wedge \psi)$.

4.3 Closure under sub-formulas

Decomposition under conjunction is a special case of closure under sub-formulas. Under some interpretations awareness must be closed under sub-formulas. In particular, if you take a deductive approach in which a formula is seen as a representation, and you want to deal with representations in a compositional way, being aware of a formula involves being aware of its sub-parts. To re-install closure under sub-formulas, we would have to add the following sub-formula constraint on models M: if $M, v \models ?\varphi$ then also $M, v \models ?\psi$ for ψ a sub-formula of φ.

In case we strengthen this constraint to a bidirectional one, i.e. $M, v \models ?\varphi$ iff $M, v \models ?\psi$ for all sub-formulas ψ, awareness of a formula φ can be reduced to being aware of the proposition letters that occur in φ (Fagin and Halpern 1988, p. 341). This corresponds to the approach of Levesque (1984), where in each world w only a particular vocabulary $\mathcal{P}(w)$ of proposition letters may be used. An agent may implicitly believe φ, but not explicitly, because some of the terms in the formula φ are unknown. Effectively, this reduces possible worlds to situations; the total valuation functions over $\mathcal{P}$ have become partial valuations in which $V(w)(p)$ is undefined when $p \notin \mathcal{P}(w)$.

Proposition 4 shows that without the constraint, our awareness is not closed under sub-formulas. This may seem counterintuitive, but can be explained by the semantic nature of our approach. Although a conjunction may be logically satisfiable, it does not have to be present in the representation. Agents may not be aware that a concept is composed of more basic parts. For example, the unemployment rate of a country is defined as the number of unemployed, divided by the total number of people able to work. This number is rising when either the number of unemployed is increasing, or the number of people able to work is decreasing, or both. For participants who are not aware of this definition B's response question may come as a surprise

(4) A: Is the unemployment rate rising?
B: I don't know. Has the number of people able to work, changed?

4.4 Closure under equivalence

Because we have chosen a semantic approach, awareness must remain closed under logical equivalence. If the semantics does not distinguish between issues $?\varphi$ and $?\psi$ and the agent is considering $?\varphi$, the agent must be considering $?\psi$ too, regardless of the way they are formulated. This gives rise to some of the usual problems with substitution in intensional contexts. For example, under the usual interpretation the following question would never be relevant, because the answer, 'none', is always supported.

(5) Which bachelors are unmarried?

However, the tautology can only be enforced by a meaning postulate like $\forall x(bachelor(x) \leftrightarrow unmarried(x) \wedge male(x))$. There is nothing in the logic itself to stop an agent from thinking that bachelors are people who are always drunk, for instance. In other words, the agent may not be aware of the meaning of a bachelor. All agents must of course be aware of equivalences that have to do with the logical connectives.

Logical equivalence means entailment in both directions, so $\eta \equiv \zeta$ iff $\eta \models \zeta$ and $\zeta \models \eta$. By definition 5,6 and 7, for propositions it turns out to be: $\varphi \equiv \psi$ whenever $M, w \models \varphi$ iff $M, w \models \psi$, for all $w \in W$.

Proposition 5 (Closure under Logical Equivalence)
Awareness is closed under logical equivalence. If $M, w \models \mathsf{A}\varphi$ and $\varphi \equiv \psi$ then also $M, w \models \mathsf{A}\psi$.
Proof. Suppose that $M, w \models \mathsf{A}\varphi$ and $\varphi \equiv \psi$. By definition $\sigma(w) \Vdash ?\varphi$, so by definition 5 it holds that $M, u \models \varphi$ iff $M, v \models \varphi$ for all $\langle u, v\rangle \in I(w)$. Because $\varphi \equiv \psi$ we get $M, w \models \psi$ iff $M, v \models \psi$ for all $\langle u, v\rangle \in I(w)$ too (transitivity of 'iff'), so $\sigma(w) \Vdash ?\psi$. $\square$

Similarly, an agent must be aware of all logically valid formulas. If $\models \varphi$ then also $\models \mathsf{A}\varphi$. For issues this means that asking for a tautology or contradiction does not add any structure. A tautological issue sets apart the possible worlds from the 'logically impossible worlds'. Since these are not included in W in the first place, all worlds are indistinguishable with respect to tautologies or contradictions.

Proposition 6 (Awareness of Logical Validities)
Awareness includes all logical validities. If $\models \varphi$ then also $\models \mathsf{A}\varphi$.
Proof. Suppose $\models \varphi$, which means $M, w \models \varphi$ for all M and $w \in W$. Take an arbitrary $\sigma = \langle S, I\rangle$. $\langle S, I\rangle[?\varphi] = \langle S, \{\langle w, v\rangle \in I \mid M, w \models \varphi \text{ iff } M, v \models \varphi\}\rangle$ equals σ itself. So $\sigma \Vdash ?\varphi$ and $\models \mathsf{A}\varphi$ $\square$

As a consequence of this property of awareness, an agent must explicitly believe all valid formulas. This is a particular aspect of the problem of deduction, called the *problem of logical omniscience*. It really becomes a problem if valid formulas are taken to include all mathematical truths. For if all mathematical truths are believed explicitly, how could you ever make mathematical discoveries? We believe that this must be solved in a different way. Not all mathematical truths are logical truths. The meaning of mathematical connectives must be dealt with explicitly by meaning postulates in the object language, similar to the bachelor example above. In that case the agent need not be aware of such meaning postulates. A mathematical discovery then means that some implicit belief has become an explicit belief.

4.5 Closure under deduction

The motivation of the concept of awareness is to solve at least some aspects of the problem of deduction: if explicit beliefs were closed under deduction, how could asking about logical consequences be relevant? Our answer is that explicit beliefs are not closed under deduction, because the agent may not be aware of the antecedent, or of the dependency between the antecedent and consequent. For this reason, we do not want awareness to be closed under deduction either. There are various ways to interpret deduction. We may mean entailment (as in definition 7), material implication, or valid implication. These interpretations are not necessarily equivalent when a non-standard semantics is used. In our logic however, the counter example of proposition 4 can be used to disprove each of these interpretations. This should come as no surprise, because decomposition under conjunction is a special case of deduction.

Proposition 7 (Non Closure under Entailment; Material Implication; Valid Implication)
Awareness is *not* closed under (1) entailment, (2) material implication or (3) valid implication. If $M,w \models \mathsf{A}\varphi$ and (1) $\varphi \models \psi$, (2) $M,w \models \varphi \to \psi$, or (3) $\models \varphi \to \psi$, not always $M,w \models \mathsf{A}\psi$.
Proof. Suppose $M,w \models \mathsf{A}(p \wedge q)$. Under both the classic and the dynamic notions of entailment, clearly (1) $(p \wedge q) \models p$, (2) $M,w \models (p \wedge q) \to p$ and (3) $\models (p \wedge q) \to p$. We use the counterexample of proposition 4. Take a pair of worlds u,v such that $M,u \models (p \wedge \neg q)$ and $M,v \models (\neg p \wedge q)$. Although $u \sim_{p \wedge q} v$ (both false) and $\langle u,v \rangle \in I(w)$, clearly not $u \sim_p v$. So $M,w \not\models \mathsf{A}p$. □

Based on the properties of awareness, we can deduce similar closure properties for explicit belief. Explicit belief is closed under negation, closed under equivalence but not closed under deduction, because awareness works like a filter. Some of the implicit beliefs of an agent are filtered out because the agent may not be aware of them until they are raised again in the form of issues.

Until now we have treated A as a kind of modal operator. An operator is modal in case it can be axiomatised using at least the propositional tautologies and the distribution axiom, which states that the modal operator distributes over the connective '$\to$'. With respect to propositional tautologies, we proved in proposition 6 that these are indeed covered. So whenever $\models \varphi$ also $\models \mathsf{A}\varphi$. However, the distribution axiom is not covered, unless we make additional assumptions about indistinguishability relations.

Proposition 8 (No Distribution)
Awareness does *not* distribute. If $M,w \models \mathsf{A}\varphi$ and $M,w \models \mathsf{A}(\varphi \to \psi)$, not always $M,w \models \mathsf{A}\psi$.
Proof. Suppose $M,w \models \mathsf{A}\varphi$ and $M,w \models \mathsf{A}(\varphi \to \psi)$. Now take $\langle u,v \rangle \in I(w)$ such that $M,u \models \neg\varphi \wedge \psi$ and $M,v \models \neg\varphi \wedge \neg\psi$. Clearly $u \sim_\varphi v$ and $u \sim_{\varphi \to \psi} v$, but not $u \sim_\psi v$. So $M,w \not\models \mathsf{A}\psi$. □

The counterexample can be explained using figure 2. The issue $?\varphi$ makes a horizontal divide; the issue $?(\varphi \to \psi)$ separates out the top right corner. However, a combination of these divisions does not produce a $?\psi$ divide. Only after a positive answer to the question $?\varphi$, restricting the data set to φ-worlds, we would get the closure property. It is easy to show that awareness does in fact distribute over dependency relations between questions: if $\mathsf{A}\varphi$ and $?\varphi \models ?\psi$ also $\mathsf{A}\psi$.

5 Discussion

We have shown that the issues raised in dialogue give rise to a particular notion of awareness, which solves some aspects of the problem of deduction. The previous sections were rather technical; we still need to discuss the adequacy of these notions.

5.1 Relevance of issues

We illustrated the notion of a logical dependency between issues by example (1). The idea was that B's question was a relevant response, because apparently it made A aware of the dependency. The argument rests on a general assumption that an utterance is relevant, if it directly or indirectly helps to partially resolve an existing issue in the context. Clearly, if awareness were closed under deduction, all relevant issues would immediately follow from such a general issue, and this explanation would become impossible. Here is another example.

(6) A: Who do you think will win?
B: Mm. Who is running?

A: (looking in the programme) Agnes Digital, Momentum, Native Driver, ...
B: Momentum has been doing well recently.

There is a dependency between B's question 'Who is running' and A's original question 'Who will win?'. This is not really a logical dependency; rather it describes a practical reasoning procedure. Finding out who is running is a prerequisite for predicting who will win, but it is not enough. Similarly, B's last remark is relevant, because of the presupposed principle that horses doing well in the past are likely to win again. An answer to the (implicit) question who did well recently, 'Momentum', can therefore count as a partial answer to the original question. This is an example of the second clause of definition 11. Although awareness is not closed under deductions, it seems to be closed under dependency relations between issues, provided the agent is or is made aware of them. There is some empirical evidence to support this idea. In a corpus study of a radio show about financial advice, Walker (1996) found many utterances that would be uninformative or redundant given a traditional account of the dialogue context. Often such utterances expressed a practical reasoning principle, for example about tax laws, or simply repeated the information on which to apply the principle, like the age group of the listener. Apparently, such 'redundant' utterances were used to reinforce participants' awareness of the information or reasoning principles. This usage also corresponds with accounts of argumentation and debate, in which moves and counter moves have the purpose of making underlying implicit reasoning principles explicit (McBurney and Parsons to appear).

5.2 Closure under linguistic substructures

Although we showed that awareness is not closed under sub-formulas, we can consider a weaker version. Earlier we discussed the so called question relation which represents a question at a level much closer to the linguistic surface structure. A DRS or SDRS would be another example. For such linguistic representations it makes sense to postulate that awareness must remain closed under subsumed representation structures. If language processing is compositional, the speaker and the hearer must be aware of the substructures of a representation in order to interpret or generate an utterance correctly. Closure under linguistic substructures would make sense for issues too. We make the following observation: expressions which are *pragmatically incomplete* trigger issues as to their completion. This is related to the Gricean maxims. An expression is pragmatically incomplete, when it is apparent from its form that it can only provide a partial solution to the contextual issue that would make the expression relevant. The non-resolved remainder of the issue is thereby brought to the foreground. The observation is illustrated by the following examples.

(7) Once upon a time there was a prince. Who is he? Is he blond?

(8) John will arrive on Saturday or on Sunday. When exactly?

(9) If Mary comes, John will not come to the party. Will she?

As we explained in the introduction, an indefinite introducing a topic, as in (7), typically triggers issues about its properties. A disjunction like in (8) is incomplete as a contribution to the underling contextual issue. Its form triggers the issue which of the disjuncts is actually true. By contrast, if the speaker would choose a non-disjunctive formulation with roughly the same meaning, like 'John will arrive in the weekend', no such issue is triggered[4]. If a conditional is uttered as a relevant contribution to the dialogue, it must be a partial answer to some question or issue. This can only be the case when its consequent is already 'at issue'. So example (9)

[4]We are not completely convinced by this data. Comments welcome.

is typically used in the context of a question like ‘Will John come to the party?’. Now the conditional can be used to deduce an answer, but only in case the antecedent is true. Thus, conditionals trigger the issue whether the antecedent is true.

To prove a closure property under substructures, we would have to add constraints on the behaviour of linguistic surface-level representations. Since we have chosen a semantic approach, in this paper we we lack the formal tools to enforce such a constraint.

6 Conclusion

In this paper we have attempted to solve some of the puzzles related to the problem of deduction. If the beliefs of an agent are closed under deduction, how could a conclusion be relevant? We hope to have shown that the answer lies in a notion of awareness, which separates implicit from explicit belief. Although agents may implicitly believe many facts, such as the fact that Chirac is taller than me, they will only realise this when the issue comes up. Therefore, we suggested to model awareness by a semantic notion of indistinguishability, which corresponds to the issues being raised in a dialogue context. We have shown that the resulting notion of awareness is closed under negation, is closed under conjunction, and closed under logical equivalence, but not decomposable under conjunction and therefore not closed under deduction. Although an agent is aware of all logical validities, our awareness notion can not be seen as a modal notion, because it does not distribute. As an additional benefit, we managed to extend the technical notion of relevance to issues, and have provided an intuitive characterisation of mental attributes with wh-complements, such as wonder or doubt.

Much more research is needed into the empirical relation between linguistic data about focus-ground, and ways of modelling the foreground-background distinction by issues in a formal semantics. In a way, this could provide a micro-level theory of relevance or salience at the utterance level. Further research is also needed into the empirical relationship between dialogue topics and issues. There are a lot of unexplored regularities between newly introduced topics and the relevance of subsequent utterances, that can be modelled using issues centred around the topic. This would then provide a macro level of relevance or salience, at the dialogue segment level. Finally, the relationship between issues and pragmatic principles needs to be investigated further. The discussion about closure under pragmatically incomplete contributions, about relevance and about dependency relations, shows we need to know more about the relation between uttering a dialogue act and the effects of its semantic content on the dialogue context.

Acknowledgements Many thanks to Zhisheng Huang for comments and helpful suggestions.

References

Barwise, Jon and John Perry. 1983. Situations and Attitudes. MIT Press.

Beaver, David. 1996. Presupposition. In Handbook of Logic and Language, ed. by Johan van Benthem and Alice Ter Meulen. Elsevier, Amsterdam.

Clark, Herbert H. and E.F. Schaefer. 1989. Contributing to discourse. Cognitive Science, 13:259–294.

Fagin, Ronald and Joseph Y. Halpern. 1988. Belief, awareness and limited reasoning. Artificial Intelligence, 34:39–76.

Fagin, Ronald, Joseph. Y. Halpern, Yoram Moses, and Moshe Vardi. 1995. Reasoning about Knowledge. MIT Press, Cambridge, Mass.

Ginzburg, Jonathan. 1995. Resolving questions, I. Linguistics and Philosophy, 18:459–527.

Ginzburg, Jonathan. 1996. Interrogatives: Questions, facts and dialogue. In The Handbook of Contemporary Semantic Theory, ed. by Shalom Lappin. Basil Blackwell, Oxford, ch 15.

Grice, H. P. 1975. Logic and conversation. In Syntax and Semantics 3, ed. by Peter Cole and Jerry L. Morgan. Academic Press, New York, pages 41–58.

Groenendijk, Jeroen. 1999. The logic of interrogation: classical version. In Proceedings of the Ninth Conference on Semantics and Linguistics Theory (SALT-9), Santa Cruz, ed. by T. Matthews and D.L. Strolovitch. CLC Publications. ILLC research report PP-1999-19.

Groenendijk, Jeroen and Martin Stokhof. 1996. Questions. In Handbook of Logic and Language, ed. by Johan Van Benthem and Alice Ter Meulen. North-Holland, Elsevier.

Huang, Zhisheng and Karen Kwast. 1991. Awareness, negation and logical omniscience. In Logic in AI, Proceedings of European Workshop on Logics in Artificial Intelligence (JELIA'90), ed. by J. van Eijck, LNCS 478. Springer-Verlag, pages 282–300.

Hulstijn, Joris. 2000. Dialogue Models for Inquiry and Transaction. Ph.D. thesis, University of Twente, Enschede.

Jäger, Gerhard. 1996. Only updates - on the semantics of the focus particle *only*. In Proceedings of the 10th Amsterdam Colloquium, ed. by P. Dekker and M. Stokhof, pages 387–406. ILLC, University of Amsterdam.

Konolige, Kurt. 1986. What awareness isn't: a sentential view of implicit and explicit belief. In Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowledge (TARK), ed. by J. Y. Halpern, pages 241–250.

Levesque, Hector L. 1984. A logic of implicit and explicit belief. In Proceedings of AAAI-84, Austin, Texas, pages 198–202.

McBurney, Peter and Simon Parsons. to appear. Games that agents play: A formal framework for dialogues between autonomous agents. Journal of Logic, Language and Information. Special Issue on Logic and Games.

Rooth, Mats. 1992. A theory of focus interpretation. Natural Language Semantics, 1:75–116.

Simon, Herbert A. 1955. A behavioral model of rational choice. Qualterly Journal of Economics, 69:99–118.

Stalnaker, Robert C. 1984. Inquiry. MIT Press.

Traum, David, Peter Bohlin, Johan Bos, Stina Ericsson, Staffan Larsson, Ian Lewin, Colin Matheson, and David Milward. 2000. Dialogue dynamics and levels of interaction. Technical report, TRINDI, LE4-8314, Dept. of Linguistics, Göteborg Universitet, Sweden.

Van Kuppevelt, Jan. 1995. Discourse structure, topicality and questioning. Journal of Linguistics, 31:109–149.

Walker, Marilyn A. 1996. Inferring acceptance and rejection in dialog by default rules of inference. Language and Speech, 39(2-3):265–304.

Wasserman, Renata. 1999. Resource bounded belief revision. Erkenntnis, 50(2-3):429–446.