

A modal interpretation of the logic of interrogation

RANI NELKEN (nelken@eecs.harvard.edu) and CHUNG-CHIEH SHAN
(ccshan@post.harvard.edu)

*Division of Engineering and Applied Sciences, Harvard University, 33 Oxford Street, Cambridge,
MA 02138*

Abstract. We propose a novel interpretation of natural-language questions using a modal predicate logic of knowledge. Our approach brings standard model-theoretic and proof-theoretic techniques from modal logic to bear on questions. Using the former, we show that our interpretation preserves Groenendijk and Stokhof’s answerhood relation, yet allows an extensional interpretation. Using the latter, we get a sound and complete proof procedure for the logic for free. Our approach is more expressive; for example, it easily treats complex questions with operators that scope over questions. We suggest a semantic criterion that restricts what natural-language questions can express. We integrate and generalize much previous work on the semantics of questions, including Beck and Sharvit’s families of subquestions, non-exhaustive questions, and multi-party conversations.

Keywords: Natural language semantics, questions, quantification, modal logic

1. Introduction

Several different approaches exist in the linguistic and logical literature for modeling natural-language questions. In linguistics, it has been popular to follow Hamblin (1973) and Karttunen (1977) (hereafter HK) in taking a question to denote its set of partial answers or partial true answers. For instance, for the *wh*-question *Who’s quitting?*, this set would contain answers such as *Alice is quitting* and *Alice and Bob are quitting*. Groenendijk and Stokhof (1984, 1997; hereafter GS) propose a more parsimonious approach, in which the answers in the set are required to be complete and mutually exclusive—in other words, a *partition* of possible worlds in the space of epistemic possibilities. For the same *wh*-question, these answers would be *Nobody is quitting*, *Just Alice is quitting*, *Only Alice and Bob are quitting*, and so on. Such classical approaches are firmly intensional, which causes complications when they try to handle more complex questions. By contrast, Nelken and Francez (2000, 2002; hereafter NF) propose an extensional interpretation: The meaning of the same question is *r* (“resolved”) if it is known for every person in the domain whether he or she is quitting. Otherwise, it is *ur* (“unresolved”).

Here, we propose a new interpretation of questions in modal predicate logic, presented in Section 2. The idea of interpreting questions using modal logic goes back to Hintikka (1976) and Åqvist (1965), who interpret a question as a request for knowledge: “bring it about that I know whether . . .”. Such a request is composed of an imperative part and an epistemic part. Focusing on the latter, we interpret a question as the *knowledge condition* required to answer it completely. We reduce the epistemic part of the meaning of both yes-no questions and *wh*-questions to



© 2005 Kluwer Academic Publishers. Printed in the Netherlands.

statements of the form “it is known whether ...” or “it is in the common ground that ...”. For instance, a yes-no question such as *Is Alice quitting?* means “it is known that Alice is quitting or it is known that Alice is not quitting”. It may seem that such approaches cannot deal with embedded questions, but we address this problem in Section 3.

The bulk of this paper bridges previous semantics of questions and combines their advantages. As GS emphasize, crucial to any semantics of questions are two entailment relations involving questions: *answerhood*, the relation between a question and its complete answers, and *question entailment*. Section 4 and Appendix A show that our modal interpretation exactly captures these entailments. Thus, Monz’s (2003) suggestion for basing practical question answering on logical inference relations can be straightforwardly implemented using existing inference procedures for modal logic. Like NF’s theory, our approach enjoys an extensional semantics, detailed in Section 5 and Appendix B.

Section 6 explains how our logic internalizes questions and is more expressive. In particular, it straightforwardly lets various operators scope over questions, which is notoriously difficult for previous theories. However, the logic is in a sense too expressive, allowing question meanings that natural-language questions cannot express. Section 7 proposes a simple semantic criterion to rule out such spurious questions: we hypothesize that a question must *license* an answer in GS’s sense.

Following GS, we start with a *strongly-exhaustive* interpretation of questions. In Section 8 we use knowledge conditions to also encode *weakly-* and *non-exhaustive* questions, as HK’s sets of answers can, and as Beck and Rullmann (1999) argue we must. This encoding lets us generalize families of *subquestions* (Beck and Sharvit 2002; Sharvit and Beck 2001) to weakly- and non-exhaustive questions, in Section 9. Finally, Section 10 extends Groenendijk’s *game of interrogation* (1999) to more than two players.

2. From knowing to asking

The basic ingredients of our proposal are found in the reader’s favorite first-order modal logic of knowledge. Given a necessity operator \Box , which can be read as “it is known that” or “it is in the common ground that”, assertions are formulae of the form $\Box\phi$. For example, for it to be asserted that Alice is quitting is for it to be in the common ground that Alice is quitting: $\Box Qa$.

We impose only minimal constraints on the logic, summarized in Table I. First, we require a normal modality to be able to reason with the logic. To use the logic as an epistemic one, we require that knowledge of a proposition implies its truth. To simplify reasoning, we further assume that the domain remains the same as in the real world, even when contemplating epistemic alternatives. For concreteness, we

Table I. Constraints we require of our logic. See footnote 1 for the axioms named in this table.

Constraint	Syntactically	Semantically
Normal modality	Necessitation rule and K axiom	Possible worlds and accessibility
Knowledge must be true	T axiom	Accessibility is reflexive
Barcan both ways	Barcan formula and its converse	Constant domain

assume that the underlying logic is S5, which is characterized by further validating that accessibility is transitive (the 4 axiom) and symmetric (the B axiom).¹

For any formula ϕ , we write $?\phi$ as shorthand for $\Box\phi \vee \Box\neg\phi$. Formulae of this form encode yes-no questions. For example, to the question *Is Alice quitting?* we assign the semantics $?Qa$, or $\Box Qa \vee \Box\neg Qa$. The intuition behind this assignment is that to know whether Alice is quitting is to either know that she is quitting or know that she is not. Thus, we directly encode the meaning of the question as its knowledge condition—what it takes to know a complete answer to the question. The intensional semantics of such formulae is the standard Kripke semantics. The meaning of $\Box Qa \vee \Box\neg Qa$ is that all the possible worlds seen from the current one agree on Qa . In other words, Qa is either uniformly true in all these worlds, or uniformly false in all of them. Similarly, we encode *wh*-questions as formulae of the form $?\vec{x}.\phi$, which is shorthand for $\forall\vec{x}.\phi$, where \vec{x} is zero or more variables.² For example, we take the meaning of *Who is quitting?* to be $\forall x.\Box Qx \vee \Box\neg Qx$. The intuition here is that to know who is quitting is to know for each person whether he or she is quitting. The intensional semantics is that all the worlds seen from the current world agree on the extension of Q —the set of people who quit must be the same in all the worlds. This approach is *strongly exhaustive* as in GS’s work; we refine this assumption in Section 8.

3. Question denotations

The main linguistic objection to reducing questions to knowledge conditions is that it seems to take knowledge as an integral part of the question meaning. This would seem to preclude treating embedding verbs such as *wonder* and *depend on*. More-

¹ For completeness, here is the list of axioms:

Necessitation If ϕ is provable with no assumption, then so is $\Box\phi$

K $\Box(\phi \rightarrow \psi) \rightarrow \Box\phi \rightarrow \Box\psi$

The T axiom $\Box\phi \rightarrow \phi$

Barcan both ways $(\Box\forall x.\phi) \leftrightarrow (\forall x.\Box\phi)$

The 4 axiom $\Box\phi \rightarrow \Box\Box\phi$

The B axiom $\phi \rightarrow \Box\neg\Box\neg\phi$

² The case of zero variables reduces to $?\phi$.

over, it does not seem to distinguish between the speech acts of asking a question and asserting knowledge of the question's answer.

To address this objection, it is important to note that the modality \Box in a formula such as $\Box Qa \vee \Box \neg Qa$ refers to the knowledge state of no particular agent or group. Neither does \Box quantify universally over all possible worlds in intensional logic. Rather, \Box is just an abstract modality. By slight abuse of notation, it is perhaps more accurate to say that the matrix question *Is Alice quitting?* and the embedded question *whether Alice is quitting* both denote the abstraction

$$(1) \quad \lambda \Box. \Box Qa \vee \Box \neg Qa,$$

in which \Box is bound by a lambda operator. We shall shortly explain the added $\lambda \Box$ model-theoretically. We posit that questions enter semantic composition as such a function. To finalize a sentence meaning, this function must be applied to some epistemic modality, in other words, to some knowledge state. Performing a matrix question applies the abstraction to the implicit conversational common ground \Box_{CG} . Likewise, a question-embedding verb such as *know* or *wonder* applies the same abstraction to other knowledge states. Thus, if to wonder is to want to know, then *wonder* denotes

$$(2) \quad \lambda q. \lambda x. 'x \text{ wants that } q(\Box_x)',$$

where \Box_x is the knowledge state of x in an alternative world, and $q(\Box_x)$ is the proposition that x knows (a complete answer to) q .

What does it mean for a function to take a modality as argument? A modality is specified by its accessibility relation (of type $\langle s, \langle s, t \rangle \rangle$), so a question meaning could simply take an accessibility relation as argument. Questions would then be of type $\langle \langle s, \langle s, t \rangle \rangle, \langle s, t \rangle \rangle$: functions from accessibility relations to propositions.

But our constraints on the accessibility relation can simplify this type. Because accessibility is transitive and symmetric in S5, and the truth value of a formula only depends on worlds at least indirectly accessible, we can simplify the argument to questions from being an accessibility relation to being merely a set of accessible worlds. Questions are then of type $\langle \langle s, t \rangle, \langle s, t \rangle \rangle$, so that $q(W)(w)$ holds just in case the set of worlds W agree on a complete answer to q at the world w . In other words, $q(W)(w)$ is true if knowing that the actual world is in W entails knowing a complete answer to q at w . Because accessibility is reflexive in S5, $q(W)(w)$ need only be defined when W contains w .

4. Entailment relations

GS describe two entailment relations involving questions that any semantics of questions should provide.

Table II. Entailment examples

Entailment	In words
$Qa \wedge \neg Qb \models ?x. Qx$	<i>Alice is quitting and Bob isn't quitting</i> answers <i>Who is quitting?</i> (This example assumes that the domain consists solely of <i>a</i> and <i>b</i> .)
$\exists x. Qx \not\models ?x. Qx$	<i>Someone is quitting</i> does not answer <i>Who is quitting?</i>
$\models ?(R \vee \neg R)$	<i>Is it either raining or not raining?</i> is trivially answered
$?x. Qx \models ?x. \neg Qx$	<i>Who is quitting?</i> entails <i>Who is not quitting?</i>
$?x. Qx \wedge Mx \not\models ?x. \neg Qx$	<i>Who is quitting and moving away?</i> does not entail <i>Who is not quitting?</i>

1. *Answerhood* is a relation between an indicative sentence and a question. An indicative is said to entail a question if the indicative completely answers the question.
2. *Interrogative entailment* is a relation between two questions. One question is said to entail another if knowing a complete answer to the first question entails knowing a complete answer to the second question.

Our theory reduces both these relations, as well as ordinary indicative entailment, to modal consequence, defined standardly as follows.

DEFINITION 1 (Modal consequence). $\psi \models^m \phi$ iff for any Kripke structure M , with a reflexive accessibility relation, if $\models_M^m \psi$, then $\models_M^m \phi$.

That is, if ψ holds in every world, then ϕ holds in every world.

Table II gives some easily verifiable examples to illustrate that our approach matches GS's empirical adequacy. In fact, we formally prove in Appendix A that our answerhood relation exactly preserves the exhaustive answerhood relation of GS. The result is quite general, and requires only that frames be reflexive. This is a welcome result, since it shows that we can achieve GS's explanatory power using a much simpler semantics (in fact, an extensional one, as described in Section 5). The intuition behind the proof is as follows. On the GS approach, a question partitions the logical space of possible worlds into equivalence classes, each corresponding to a possible answer to the question. On our approach, the (knowledge condition of the) question holds in each possible world w iff all the worlds accessible from w (that is, all the worlds compatible with the agent's knowledge) agree on the answer. In other words, to assign a question the value "true" is to entertain only epistemic possibilities that agree on the answer to the question. Thus, for a question to be resolved is for our accessibility relation to respect GS's partition boundaries, yielding answerhood preservation.

By assigning truth values (of modal formulae) to questions, we gain the simplicity of NF's analysis: a uniform consequence relation that subsumes indicative entailment, answerhood, and interrogative entailment. We also gain expressive

power with which to account for more complex questions, as explained in Section 6. Another advantage of our modal perspective is that it bridges the *what* of question answering, which linguistic semantics is concerned with, to the *how* of question answering, which practical systems must embody (Monz 2003): established computational techniques of modal logic can now be used to reason about questions. Previous approaches to implement GS’s logic (Bos and Gabsdil 2000) restricted their attention to a weakened variant of GS’s entailment relations. By making the relation with modal logic, we can apply inference procedures for modal logic to answerhood and interrogative entailment, extending ten Cate and Shan’s (2002) question-answering algorithm for GS’s partition semantics to any question meanings encoded as knowledge conditions. The procedure is sound and complete—given a question, it generates *only* and *all* answers. However, it does not always terminate—in some cases, there are an infinite number of possible answers, and no answer is most informative. In fact, we can directly use Cerrito and Cialdea Mayer’s (2001) proof procedure for first-order S4 logic, and its theorem prover implementation (Thion et al. 2002).

5. Extensional semantics

Both HK and GS take a question to mean its answers. Because answers are propositions, such an approach inherently attributes the added complexity of intensionality to interrogative sentences, but not to indicative sentences. While there may be good reasons to adopt such an asymmetry, it is less clear what level of intensionality is actually necessary to interpret questions. To examine this issue, we propose to try to construct an extensional semantics of questions, and see how far one can take it. Even if intensionality turns out to be ultimately necessary, we will better understand what phenomena in the interpretation of questions really overstep the extensional boundary and drive the need for intensionality.

NF propose an explicitly extensional interpretation of questions that assigns to each question one of 5 rather than 2 truth values, organized in an algebraic structure called a bilattice. Unfortunately, this interpretation is not quite empirically adequate. For instance, it does not capture the third example in Table II—repeated below—without further assumptions (Nelken and Francez 2002).

$$(3) \models ?(R \vee \neg R).$$

By contrast, our modal approach offers a new extensional interpretation of questions that preserves GS’s entailment relations, including (3) and the rest of Table II.

Recall that propositional modal logic has the *finite-model property*: if a collection of formulae has a model, then it has a model with a finite number of possible worlds. This useful property does not hold in general for first-order modal (predicate) logic, but we show in Appendix B that it does hold for implications among first-order modal formulae *that encode questions*. In fact, two worlds in the model

are enough: any entailment between questions is satisfiable iff it is satisfiable with two worlds, and falsifiable iff it is falsifiable with two worlds.

This result provides a novel extensional interpretation that exactly preserves the entailment relations of GS. We can easily simulate a two-world structure by a non-modal first-order model that assigns to each atomic formula one of 4 truth values: $\{FF, TF, FT, TT\}$. The truth value of a more complex formula is computed by applying the regular truth tables of the logical operators pointwise; for example, the disjunction of TF and FT is TT . The $?$ operator checks whether the two worlds agree: it maps TT and FF to TT ; and TF and FT to FF .

To illustrate this simulation, let us consider (3) once again. Since the logical operators operate pointwise, $R \vee \neg R$ has the value TT regardless of the value of R . Thus $?(R \vee \neg R)$ is TT , so the entailment holds in this 4-value logic, as desired.

This novel extensional interpretation nevertheless has an intensional flavor, in that the 4 truth values simply encode 2 possible worlds. This flavor raises the more philosophical issue of what constitutes an extensional semantics. On one hand, if we forget about possible worlds and just use a non-classical handful of truth values, is our semantics then extensional? On the other hand, if we encode more possible worlds, say 20 instead of 2, yielding an astronomical number of truth values, is the interpretation then intensional? Is intensionality a binary distinction, or a matter of degree?

6. Internalized questions

Because we let questions denote their knowledge conditions, which are just modal formulae, they can further combine with other formulae. In particular, our $?$ operator is *internalized*: as in NF's but not GS's analysis, it can apply anywhere in a question meaning, not just at the top level or under some top-level universal quantifiers. We can thus handle the following constructions using standard logical connectives.

Conjunction *Do you have a license and who (else) has one?* ($?L \wedge ?x.Lx$).

Disjunction *What's your social security number? Or what's your mother's maiden name?* ($?x.Sx \vee ?x.Mx$). *Is it getting warmer? Or is it just me?* Whereas conjunction is straightforward for GS, they get disjunction only by resorting to higher-order type-shifting from partitions to sets of sets of partitions, yielding a highly complex object.

There is some debate in the literature over whether disjoined questions are available in natural language. Whereas GS and NF accept them, Szabolcsi (1997) claims that they are unavailable. We offer a refined prediction in Section 7.

Conditional *If it's raining, who has an umbrella?* ($R \rightarrow ?x.Ux$). These questions are unavailable for GS.

As the formula indicates, the conditional questions we consider here—where an implication leads from an indicative to a question—are those that can be completely answered by falsifying the antecedent in the common ground. For example, the question above can be completely answered by announcing *it's not raining*. On one hand, when the antecedent is hard to falsify (for example, when it is next week's weather that is under discussion), a conditional question such as $R \rightarrow ?x.Ux$ is hard to distinguish from a question whose body is a conditional, such as $?x.R \rightarrow Ux$ (*Who, if it rains next week, will have an umbrella?*). On the other hand, when the antecedent is manifestly true (*If I'm sick, why did you take me on a 5-mile hike?*), the knowledge condition is equivalent to the unconditional question (*Why did you take me on a 5-mile hike?*), though obviously there is at least a pragmatic difference between the two questions.

Universal quantification *Who recommends each candidate?* ($\forall y.Cy \rightarrow ?x.Rxy$). In previous theories, because questions are not propositions, they are difficult to quantify over. For instance, Karttunen uses double negation to get these readings. Since conditional questions are unavailable for GS, so is universal quantification over conditional questions.

Also, by universally quantifying over conditional questions, we can express the contrast between *de dicto* and *de re* readings of *which*-questions (Groenendijk and Stokhof 1984), as in *Alice just discovered which spies are quitting*. The *de dicto* reading here ($\forall x.Sx \rightarrow ?Qx$) neither entails nor is entailed by the *de re* reading ($\forall x.?(Sx \wedge Qx)$).

Whereas HK and GS struggle to get these representations if they get them at all, our logic allows them straightforwardly. However, nothing in the expressivity of our logic bars other combinations that, while logically possible, do not correspond to natural-language questions. For instance, while a question-denoting formula can be negated to express that its knowledge condition does not hold, natural language does not recognize such a construction ($\neg ?R$: **Not is it raining?*). We turn to these cases in the next section.

7. Licensing answers

Groenendijk (1999) uses interrogative entailment to define a semantic notion of *licensing*.

DEFINITION 2 (Licensing). *A question Q licenses an answer A just in case $Q \models ?A$.*

For Q to license A does not require that A completely answer Q . That is, $Q \models ?A$ may hold even when $\Box A \models Q$ does not. Rather, licensing means that if Q is fully resolved then the question “whether A ” is also resolved.

To illustrate this difference, consider a *maybe* answer to a yes/no question such as *Is it raining?*³ Clearly, $\Diamond R \not\models ?R$, since the existence of a reachable world in which R holds does not entail that all reachable worlds have this property. However, yes/no questions do license *maybe* answers, as it is easy to verify that on transitive frames, $?R \models ?\Diamond R$.

As explained in Section 6, our logical language of questions is more expressive than GS’s. Using GS’s notion of licensing, we can circumscribe this expressive power to predict what question meanings are available in natural language. We propose the following criterion:

CLAIM 1. *A question is acceptable in natural language only if it licenses a non-trivial answer.*

More formally, if Q is a question, then we require that $Q \models ?A$ for some (non-modal) proposition A that is neither tautologically true nor tautologically false. We must exclude such trivial answers because every question entails $?T$ and $?F$. The following cases illustrate this claim.

Yes-no questions The question $?A$ is acceptable for any non-trivial indicative A , since we always have $?A \models ?A$. In other words, the question *whether A* always licenses the answer *that A*.

Simple *wh*-questions The question $?x.Px$ is acceptable for any non-trivial property or relation P , since we always have $?x.Px \models ?P\vec{a}$, where \vec{a} is a sequence of individuals.

Conjunction Let Q be a conjunction of questions, say $Q = Q_1 \wedge Q_2$, such that at least one of the conjuncts is acceptable, say Q_1 . Then there exists some answer A such that $Q_1 \models ?A$. It is easy to verify that $Q_1 \wedge Q_2 \models ?A$, so Q is also acceptable.

Disjunction As mentioned in Section 6, there is some controversy as to whether disjoined questions are acceptable. Our licensing criterion makes a novel prediction: disjoined questions is usually unacceptable but occasionally acceptable.

For example, to know a complete answer to *Did you just come out of a swimming pool, or is it raining?* ($?S \vee ?R$) is to either know a complete answer to *Did you just come out of a swimming pool?* ($?S$) or know a complete answer to *Is it raining?* ($?R$). If it is in the common ground that exactly one of S and R is true—for example, if the listener is clearly drenched, the

³ We thank an anonymous referee for this example.

only conceivable ways for the listener to be drenched are S and R , and they are mutually exclusive—then the disjunction forms an acceptable question, because it licenses such non-trivial answers as S and R . Formally, we have that $S \leftrightarrow \neg R$, $?S \vee ?R \models ?A$, for $A = S$ and $A = R$.

Similarly for *Would you like coffee or would you like tea?*, which we can formalize as $?C \vee ?T$. At least on the surface, this sentence seems to be a (perfectly acceptable) disjunction of two questions. Szabolcsi claims that, in an apparently disjoined question, the second question cancels the first, but that claim cannot be maintained here, since one can still answer the first question. Note the strong contextual presupposition that exactly one of the two choices is true ($C \leftrightarrow \neg T$). Where does this presupposition come from? It is easy to verify that, with the assumption $C \leftrightarrow \neg T$, the question does license the non-trivial answers C and T .

Knowing a complete answer to one of two questions usually does not entail knowing a complete answer to any particular yes-no question, so questions usually cannot be disjoined. For example, even though *Who is coming to the party?* and *Is it raining?* are both acceptable questions, we predict correctly that their disjunction is not.

Negation Not knowing a complete answer to one question never entails knowing a complete answer to another question. Hence we predict, correctly, that questions can never be negated.

Quantification The question *Who recommends each candidate?* can be formalized as $\forall y. Cy \rightarrow ?x. Rxy$. For this question meaning to license a non-trivial answer, there must be at least one individual, say Alice (a), whose candidacy is in the common ground. Assuming that Ca , the question then licenses the answer Raa . Conversely, if nobody's candidacy is in the common ground, then our proposal correctly rules out the question.

In this way, our licensing criterion allows a universally quantified question when at least one individual is commonly known to be in the domain of quantification, but other quantifiers such as *most* and *no* cannot take scope over a question. For example, the question *Who recommends most candidates?* cannot be answered by specifying the recommenders of most candidates in a so-called pair-list reading.

Presuppositions Presuppositions project out of questions. For example, the question *Did you stop smoking?* presupposes that you smoked, just as its indicative counterparts *You stopped smoking* and *You did not stop smoking* do. Similarly, *Which candidates stopped smoking?* presupposes that at least some of the candidates smoked. Our account expects these question presuppositions, because the licensing criterion places the question in the antecedent part of an entailment, where presuppositions project out. For instance, we must know for

each candidate whether they stopped smoking, and hence that they smoked in the past.

In proposing this constraint on natural-language questions, we diverge from previous methodology. Previous approaches start with a less expressive base, then add mechanisms to handle more complex constructions available in natural language. For other, unacceptable operations, these theories can just remain silent, and the issue why such operations are barred never arises. While formally adequate, such an approach cannot justify or even contemplate why certain boolean operators are acceptable with questions whereas others are not. By contrast, we do not delegate the range of available constructions to the syntax of the logic. Instead, we consider the full range of operations, then explain why certain combinations are impossible by proposing a semantic criterion that has to do not with the logical apparatus, but with the empirical properties of natural language.

8. Exhaustive questions versus complete answers

An assertion ϕ is a complete answer to a question ψ just in case ϕ entails ψ . For example, asserting that it is raining and nobody is quitting ($R \wedge \forall x. \neg Qx$) completely answers the question who is quitting ($?x. Qx$).

Completeness relates answers to questions. *Exhaustivity* is a separate notion that applies to *wh*-questions only.

1. The encoding of *wh*-questions in Section 2—as formulae of the form $\forall \vec{x}. ?\phi$ —is *strongly exhaustive* in that it universally quantifies over \vec{x} . To know who is quitting in this sense is to know, for each person x , either that x is quitting or that x is not.
2. By contrast, to know who is quitting in the *weakly exhaustive* sense is to know, for each person x who is quitting, that x is quitting: $\forall x. Qx \rightarrow \Box Qx$, or more generally $\forall \vec{x}. \phi \rightarrow \Box \phi$.
3. Finally, a question such as *Where can I get gas?*, in an appropriately desperate situation, is *non-exhaustive* and would be represented by existential quantification: $\exists x. \Box Gx$. Any assertion of a gas station's location qualifies as a complete answer.

Any strongly-exhaustive question $\forall \vec{x}. ?\phi$ can be recast as the weakly-exhaustive question

$$(4) \quad \forall \vec{x}. \forall y. (y = \phi) \rightarrow \Box (y = \phi)$$

(in which the variable y ranges over truth values), because both questions are equivalent to

$$(5) \quad \forall \vec{x}. (\phi \rightarrow \Box \phi) \wedge (\neg \phi \rightarrow \Box \neg \phi).$$

We can characterize the difference between strongly-, weakly-, and non-exhaustive questions semantically, as follows. All questions on our account are *monotonic*: it never hurts to know more (that is, to entertain fewer epistemic possibilities).

CLAIM 2 (Monotonic). *If q is a question, then $q(W_1)(w)$ implies $q(W_2)(w)$ whenever $W_1 \supseteq W_2$.*

Exhaustive questions are furthermore *additive*: a disjunction of complete answers is still a complete answer.

CLAIM 3 (Additive). *If q is an exhaustive question, then $q(W_1)(w) \wedge q(W_2)(w)$ implies $q(W_1 \cup W_2)(w)$ (similarly for infinite unions).*

Non-exhaustive questions are not additive: *Where can I get gas?* is completely answered by both *in Central Square* and *in Inman Square*, but not by *in either Central Square or Inman Square; I'm not sure which*. Strongly-exhaustive questions are not only additive but also constant across worlds. That is, $q(W)(w_1) = q(W)(w_2)$, if q is strongly exhaustive, for $w_1, w_2 \in W$.

Let q be an additive question denotation. Following Heim's terminology (1994), we define the *answer1* of q at a world w , written $A_w^1(q)$, to be the union of all world-sets W satisfying $q(W)(w)$. The additivity of q ensures that the answer1 is itself an answer, in that $q(A_w^1(q))(w)$ is true. Because q is monotonic, $q(W)(w)$ is true just in case $W \subseteq A_w^1(q)$. Thus q is determined by $A_w^1(q)$ at each world w , and we are tempted to further simplify the semantic type of exhaustive questions from $\langle\langle s, t \rangle, \langle s, t \rangle\rangle$ to $\langle s, \langle s, t \rangle \rangle$: given a world w , return the answer1 $A_w^1(q)$ at w . However, this simplification only works for exhaustive questions. Hence, Beck and Rullmann (1999; Section 7.1) argue that a question must sometimes denote a set of answer propositions, rather than always denoting its answer1 proposition. By contrast, we let strongly-, weakly-, and non-exhaustive questions uniformly denote the same type $\langle\langle s, t \rangle, \langle s, t \rangle\rangle$.

Note that weakly- and even non-exhaustive questions are allowed by the licensing criterion of Section 7. For instance, though $\exists x. \Box Gx$ does not entail $?Ga$ for any domain element a , it still entails the existential question *Can I get gas (somewhere)?*— $? \exists x. Gx$.

Van Rooy (2003) views the difference between exhaustive and non-exhaustive questions not as ambiguity but rather as underspecification.⁴ He assigns to questions a uniform meaning regardless of whether they are exhaustive or not, namely a family of (potentially overlapping) sets of possible worlds. These sets depend on the optimal answers to the asker's *decision problem* that motivates asking the question in the first place. This approach has the advantage of unifying questions of different exhaustivity levels. However, it does so at the price of introducing additional notions having to do with the speaker's hidden mental state (such as the underlying decision problem and the optimal utility of answers). Our knowledge

⁴ We thank an anonymous reviewer for directing our attention to this point.

conditions unify different exhaustivity levels in a simpler way. In fact, knowledge conditions can be seen as encoding the utility of an answer as a binary value: either useful (that is, entailing the knowledge condition) or not.

It is technically possible to reformulate van Roooy's approach in our terms by dividing a subset of the domain (consisting of those elements that are *relevant* according to van Roooy) into a family \mathcal{S} of possibly overlapping subsets. Intuitively, these are the subsets of equal utility for the asker. The meaning of a question then remains universal quantification, but with an added existential quantification over the subsets in \mathcal{S} . For instance, *Where can I get gas?* becomes $\exists D \in \mathcal{S}. \forall x \in D. \Box Gx$. The exhaustive case is when \mathcal{S} consists of a single set: the domain. The non-exhaustive case is when \mathcal{S} is a full partition of the domain into singletons. Presumably, \mathcal{S} can be determined by the same methods as van Roooy determines his division of the set of possible worlds.

9. Plurality of questions

As explained above, we interpret a *wh*-question by universally quantifying over individual questions. This strategy is tightly related to Beck and Sharvit's work (2002; Sharvit and Beck 2001) on *families of subquestions* for explaining *quantificational variability* effects. On Beck and Sharvit's analysis, sentences such as

- (6) Alice mostly knows who is quitting.
- (7) With few exceptions, Alice knows who is quitting.

quantify over a contextually salient family of subquestions of the question *who is quitting*. A family of subquestions is simply a set of questions, satisfying some conditions detailed below. For example, the family \mathcal{S} of subquestions might be

- (8) { is Alice quitting, is Bob quitting, is Carol quitting, ... },

and the sentences (6) and (7) mean

- (9) For most questions s in \mathcal{S} , Alice knows s .
- (10) For all but few questions s in \mathcal{S} , Alice knows s .

Beck and Sharvit argue that these sentences quantify over not persons or propositions but questions. Leaving these arguments aside, we show here how our modal perspective expresses their proposal, so as to generalize it to weakly- and non-exhaustive questions as introduced in Section 8.

It is tempting to view the *wh*-question formula $?x.\phi$ as explicitly encoding the set of yes-no questions $\{?\phi[\vec{d}/\vec{x}] \mid \vec{d} \in D^n\}$, where D is the domain and \vec{x} consists of n variables. Indeed, this set is a family of subquestions, used in the simplest cases of quantificational variability, such as those above. But as Beck and Sharvit show,

this set sometimes differs from the contextually salient family of subquestions. To allow for these cases, they propose that any set S of questions can be used as the family of subquestions of a question q , subject to the following three criteria.

1. Answering every question in S would answer q as well.
2. Every question s in S is a *subquestion* of q . Formally, if s and q both partition the set of possible worlds into classes, then some class of s is disjoint from some class of q . In other words, there are two possible worlds w_s and w_q such that the partition in s containing w_s is disjoint from the partition in q containing w_q .⁵ Thus s is a subquestion of q iff q is a subquestion of s .
3. No proper subset of S satisfies both criteria above.

The second criterion above mentions partitions: Beck and Sharvit only consider strongly exhaustive questions, which they follow GS in treating as partitions of the set of possible worlds. It turns out that we can apply the family-of-subquestions approach not only to strongly-exhaustive questions but also to weakly- and sometimes non-exhaustive questions. However, since partitions cannot model weakly- and non-exhaustive questions, we must reformulate the second criterion using only knowledge conditions. Instead of referring to the partitions containing w_s and w_q , we can look at an answer to s in w_s and an answer to q in w_q . These answers must be mutually exclusive. That is, for S to be a family of subquestions of q requires the following three criteria.

1. The conjunction of all questions in S entails q .
2. Every question s in S is a subquestion of q , in the sense that there exist two possible worlds w_s and w_q , such that any two propositions W_s and W_q are mutually exclusive whenever they answer s and q at w_s and w_q , respectively. In other words, for s and q to be subquestions of each other is for there to be worlds w_s and w_q , such that any two world-sets W_s and W_q are disjoint whenever $s(W_s)(w_s)$ and $q(W_q)(w_q)$ are both true.⁶
3. No proper subset of S satisfies these two criteria.

This reformulation has the virtue of being applicable even when the partition semantics of questions is not.

Despite this reformulation, non-exhaustive questions cannot ordinarily be interpreted as a family of subquestions, as reflected by their distinct (existential) representation. For example, if Alice knows the exact location of even a single gas station, then Alice knows exactly where to get gas. If Alice runs out of gas while

⁵ If, to the contrary, every partition of s intersects every partition of q , then s and q are unrelated questions.

⁶ Recall that, if q is a question, W is a set of worlds, and w is a world, then $q(W)(w)$ holds iff knowing that the actual world is in W entails knowing the complete answer to q at w .

driving, and she asks Bob whether he knows where to get gas, it would be odd for Bob to answer *With few exceptions, I do*, even if Bob knows the exact locations of all but a few gas stations in the area. This oddity is because no discrete family of subquestions is contextually salient. Bob could say *I mostly do*, not to mean that he knows the locations of most gas stations nearby, but to mean that he has somewhat vague recollections of how to navigate to a certain gas station nearby. The relevant family of subquestions here is how to navigate at each step.

10. Multi-party conversations

Because the partition theory of questions embeds into our semantics, we can recast Groenendijk's game of interrogation (1999) in our terms. In Groenendijk's original game, an *interrogator* and a *witness* take turns asking questions and asserting answers.

1. The interrogator must not ask superfluous questions. For example, having asked *Who is quitting?*, the interrogator must not ask *Who is not quitting?*, because of the fourth example in Table II above.
2. The witness must only assert licensed answers. For example, having been asked *Who is quitting and moving away?*, the witness must not assert *Alice is not quitting*. (The witness must also avoid assertions that are redundant, meaning entailed by the common ground, or incredible, meaning inconsistent with the common ground.)

From the modal perspective, we can extend this game from one interrogator and one witness to multiple, overlapping groups of participants.

Groenendijk keeps track of the knowledge and issues in the common ground using a *context* C , which is a partial equivalence relation over possible worlds, or equivalently, a partition of a subset of the set of possible worlds. A world appears in the context just in case it is not yet ruled out by the knowledge state in the common ground; two worlds are related by the context just in case it is not yet under discussion which one is real. In Groenendijk's game, the only knowledge state and knowledge conditions relevant to felicity in the conversation is the common ground between the interrogator and the witness. We can generalize this to knowledge state and knowledge conditions among overlapping groups by keeping track of one context per group. That is, for every group of participants G , we keep track of a context C_G for that group, still a partial equivalence relation over worlds. Entailment among contexts respects containment among groups; that is, if $G' \subseteq G$, then $C_{G'} \subseteq C_G$ throughout our generalized game.

Groenendijk's interrogator updates the context by removing equivalences; his witness updates the contexts by removing worlds. For a question q (denoting a partition of worlds) to be non-superfluous is for the context C to not entail q . For an

assertion ϕ (denoting a set of worlds) to be licensed is for the context C to entail $?\phi$, where $?\phi$ is the binary partition of worlds formed by ϕ and its complement $\neg\phi$. In our generalization, the group G of participants in the room can change from move to move. A question still removes equivalences, and an answer still removes worlds, but from *every* context $C_{G'}$ for subgroups G' of G . (It is easy to check that this update procedure preserves the invariant that entailment among contexts respects containment among groups.) Each question q must not be redundant: C_G better not entail q . (Because of the invariant, then, it never hurts to ask a question in a bigger group. For example, the same question may be asked again after a participant enters, but not leaves, the room.) Each answer ϕ must be licensed: C_G better entail $?\phi$. (Because of the invariant, it never hurts to assert an answer in a smaller group—except the assertion may then be redundant or incredible.)

The generalization just described of Groenendijk’s game may seem straightforward at first glance, and to the extent that it is straightforward, the link that this paper makes between question meanings and knowledge conditions is successful. However, we must point out two caveats that call for future work.

First and foremost, neither Groenendijk’s nor our game deals with questions that are not strongly exhaustive. The problem lies with the licensing condition for answers: according to Groenendijk, the witness can assert an answer ϕ just in case the context entails $?\phi$. But if the question ψ under discussion is weakly- or non-exhaustive, then even a complete answer ϕ to ψ (that is, so that $\Box\phi \models \psi$) is unlicensed (that is, $\psi \not\models ?\phi$)! This wrongly prevents any answer to a question such as *Where can I get gas?*. Hence this relevance criterion must be revised for non-strongly-exhaustive questions. Invoking the notion of subquestions, perhaps one could err on the side of permissiveness and allow any answer ϕ to be asserted as long as $\Box\phi$ completely answers any subquestion in any family of subquestions of ψ .

Second, as we consider more complex games of interrogation, it becomes less clear that our formal criteria for redundancy and relevance really correspond to intuition. In what sense do participants in these games *know* what issues are under discussion, and work towards the *goal* of resolving these issues? Not in any sense so far formally related to logics for epistemic actions. Grounding games of interrogation such as ours, in logics for epistemic actions such as Baltag et al.’s (1999), would help model subtleties such as the following. Suppose that Alice asks Bob and Carol whether it is raining, then Bob leaves the room. Can Carol still tell Alice that it is raining? This assertion does not further the goal of common knowledge among Alice, Bob, and Carol, yet our generalized game above allows it.

Acknowledgements

We thank Nissim Francez, Stuart Shieber, anonymous reviewers, and participants of the SALT 14 conference for their comments on earlier versions of this paper.

The authors are supported by United States National Science Foundation Grants IIS-0329089 and BCS-0236592.

References

- Åqvist, Lennart. 1965. *A new approach to the logical theory of interrogatives*. Uppsala: Almqvist & Wiksell.
- Baltag, Alexandru, Lawrence S. Moss, and Slawomir Solecki. 1999. The logic of public announcements, common knowledge and private suspicions. Tech. Rep. 534, Computer Science Department, Indiana University.
- Beck, Sigrid, and Hotze Rullmann. 1999. A flexible approach to exhaustivity in questions. *Natural Language Semantics* 7(3):249–298.
- Beck, Sigrid, and Yael Sharvit. 2002. Pluralities of questions. *Journal of Semantics* 19(2):105–157.
- Bos, Johan, and Malte Gabsdil. 2000. First-order inference and the semantics of questions and answers. In *Göteborg 2000: 4th workshop on the semantics and pragmatics of dialogue, June 15–17*, ed. M. Poesio and D. Traum, 43–50. Göteborg, Sweden: Gothenburg Papers in Computational Linguistics 00-5.
- ten Cate, Balder, and Chung-chieh Shan. 2002. Question answering: From partitions to Prolog. In Egly and Fermüller (2002), 251–265.
- Cerrito, Serenella, and Marta Cialdea Mayer. 2001. Free-variable tableaux for constant-domain quantified modal logics with rigid and non-rigid designation. In *Proceedings of IJCAR 2001: 1st international joint conference on automated reasoning*, ed. Rajeev Goré, Alexander Leitsch, and Tobias Nipkow, 137–151. Lecture Notes in Computer Science 2083, Berlin: Springer-Verlag.
- Egly, Uwe, and Christian G. Fermüller, eds. 2002. *Proceedings of TABLEUX 2002: Automated reasoning with analytic tableaux and related methods*. Lecture Notes in Computer Science 2381, Berlin: Springer-Verlag.
- Groenendijk, Jeroen. 1999. The logic of interrogation: Classical version. In *Proceedings from Semantics and Linguistic Theory IX*, ed. Tanya Matthews and Devon Strolovitch, 109–126. Ithaca: Cornell University Press.
- Groenendijk, Jeroen, and Martin Stokhof. 1984. Studies on the semantics of questions and the pragmatics of answers. Ph.D. thesis, Universiteit van Amsterdam.
- . 1997. Questions. In *Handbook of logic and language*, ed. Johan van Benthem and Alice ter Meulen, 1055–1124. Amsterdam: Elsevier Science.
- Hamblin, Charles Leonard. 1973. Questions in Montague English. *Foundations of Language* 10: 41–53.
- Heim, Irene. 1994. Interrogative semantics and Karttunen’s semantics for *know*. In *Proceedings of the Israeli association for theoretical linguistics*.
- Hintikka, Jaakko. 1976. The semantics of questions and the questions of semantics. *Acta Philosophica Fennica* 28(4).
- Karttunen, Lauri. 1977. Syntax and semantics of questions. *Linguistics and Philosophy* 1(1):3–44.
- Monz, Christof. 2003. From document retrieval to question answering. Ph.D. thesis, Universiteit van Amsterdam.
- Nelken, Rani, and Nissim Francez. 2000. A calculus of interrogatives based on their algebraic semantics. In *Proceedings of TWLT16/AMILP2000: Algebraic methods in language processing*, ed. Dirk Heylen, Anton Nijholt, and Giuseppe Scollo, 143–160.
- . 2002. Bilattices and the semantics of natural language questions. *Linguistics and Philosophy* 25(1):37–64.

- van Rooy, Robert. 2003. Questioning to resolve decision problems. *Linguistics and Philosophy* 26(6): 727–763.
- Sharvit, Yael, and Sigrid Beck. 2001. Subquestions and quantificational variability effects. In *Proceedings of the 20th West Coast Conference on Formal Linguistics*, ed. Karine Megerdooimian and Leora A. Bar-el, 510–523. Somerville, MA: Cascadilla.
- Szabolcsi, Anna. 1997. Quantifiers in pair-list readings. In *Ways of scope taking*, ed. Anna Szabolcsi, chap. 9, 311–347. Dordrecht: Kluwer.
- Thion, Virginie, Serenella Cerrito, and Marta Cialdea Mayer. 2002. A general theorem prover for quantified modal logics. In Egly and Fermüller (2002), 266–280.

Appendix

A. Answerhood preservation

In this appendix, we provide the proof that answerhood is preserved for reflexive frames. Recall that according to GS, an indicative ψ , in a structure M denotes a set of possible worlds, denoted by $\llbracket \psi \rrbracket_M$. An interrogative, ϕ , denotes a partition relative to a structure M . In each possible world it denotes a class of that partition, denoted $\llbracket \phi \rrbracket_{M,w}$. GS define ϕ to be an answer to ψ if the meaning of ψ is wholly included in some class of the partition.

DEFINITION 3 (GS exhaustive answerhood). *For indicative ψ and interrogative ϕ , $\psi \models^e \phi$ if and only if for every structure M , $\llbracket \psi \rrbracket_M \subseteq \llbracket \phi \rrbracket_{M,w}$ for some w .*

Using this definition, and Definition 1 we can now state the answerhood preservation theorem.

THEOREM 1. *For indicative ψ and interrogative ϕ , $\psi \models^e \phi$ iff $\psi \models^m \phi$*

To prove the theorem, we relate the structures used for the two kinds of semantics. We view modal-based structures as “sub-structures” of partition-based structures. For a Kripke structure M and formula ϕ , we would like to associate a subset $\mathcal{W}' \subseteq \mathcal{W}$ with ϕ . If ϕ is indicative, we take some set of worlds in which it is true. If ϕ is interrogative, we take a subset of one of the equivalence classes in the partition it induces.

DEFINITION 4 (ϕ -set). *Let M be a Kripke structure. Let ϕ be a formula. A non-empty set of possible worlds $\mathcal{W}' \subseteq \mathcal{W}$ is called a ϕ -set iff:*

- If ϕ is indicative then $\mathcal{W}' \subseteq \llbracket \phi \rrbracket_M$
- If ϕ is interrogative then $\mathcal{W}' \subseteq \llbracket \phi \rrbracket_{M,w}$ for some $w \in \mathcal{W}$

Now, given a formula ϕ and Kripke structure M , take a ϕ -set, \mathcal{W}' , throw on any reflexive accessibility relation, and view the result as a new Kripke structure M' .

DEFINITION 5 (Induced Kripke structure). *Let $M = \langle \mathcal{W}, R, D, a, V \rangle$ be a Kripke structure. Let $\mathcal{W}' \subseteq \mathcal{W}$, and let R' be a reflexive accessibility relation on \mathcal{W}' . The Kripke structure induced by M , \mathcal{W}' and R' is $M' = \langle \mathcal{W}', R', D, a, V|_{\mathcal{W}'} \rangle$*

Interestingly enough, M' satisfies ϕ 's translation:

LEMMA 1. *Let $M = \langle \mathcal{W}, R, D, a, V \rangle$ be a Kripke structure. Let ϕ be a formula. Let \mathcal{W}' be a ϕ -set, and choose $R' \subseteq \mathcal{W}' \times \mathcal{W}'$ to be a reflexive relation on \mathcal{W}' . If M' is the Kripke structure induced by M , \mathcal{W}' and R' , then $\models_{M'}^m \phi$*

Proof.

- If ϕ is indicative, then ϕ is true in all the worlds of $\llbracket \phi \rrbracket_M$. Hence, M' only contains worlds in which ϕ is true.
- If $\phi = ?\vec{x}.\eta$ is interrogative, then by our construction, \mathcal{W}' is a subset of one of the classes in the partition induced by ϕ . Hence, the extension of η is the same in all possible worlds $w \in \mathcal{W}'$. Thus for any tuple of domain elements \vec{d} over D of the same arity as \vec{x} , η is assigned the same truth value in all of the worlds in \mathcal{W}' , under the assignment of \vec{x} to \vec{d} . Consequently, by definition, $\models_{M'}^m \forall \vec{x}(\Box \eta \vee \Box \neg \eta)$.

For the other direction, let M be a Kripke structure with reflexive accessibility s.t. $\models_M^m \phi$. Choose any possible world $w \in \mathcal{W}$, and define $Rw =_{def} \{w' | wRw'\}$ as the set of possible worlds accessible from w . Rw is a ϕ -set.

LEMMA 2. *If $\models_M^m \phi$, then Rw is a ϕ -set for any w .*

Proof. First, note that for any w , Rw is non-empty since R is reflexive.

- If ϕ is indicative, then all the worlds in \mathcal{W} satisfy ϕ .
- If ϕ is interrogative, let $\phi = ?\vec{x}.\eta = (? \vec{x}.\eta) = \forall \vec{x} . (\Box \eta \vee \Box \neg \eta)$. Assume $\models_M^m \phi$. Let \vec{d} be a tuple of domain elements of the same arity as \vec{x} . Under the assignment $a[\vec{d}/\vec{x}]$, $\Box \eta \vee \Box \neg \eta$ is satisfied by M and w . Therefore one of the two disjuncts is satisfied by M and w . Hence for all $w' \in Rw$, $v_{w'}(\eta) = v_w(\eta)$, under this assignment. Thus the extension of η is the same over all Rw . Clearly, $Rw \subseteq \llbracket \phi \rrbracket_{M,w}$. Hence, Rw is a ϕ -set.

Based on these two lemmata, we can now prove the theorem:

Proof.

- Assume $\psi \models^e \phi$. Let $M = \langle \mathcal{W}, R, D, a, V \rangle$ be a Kripke structure with reflexive accessibility such that $\models_M^m \psi^\sharp$. We must show that $\models_M^m \phi$. By Lemma 2, Rw is a ψ -set for any $w \in \mathcal{W}$. In other words, $Rw \subseteq \llbracket \psi \rrbracket_M$ for any $w \in \mathcal{W}$. Thus, $R \subseteq \llbracket \psi \rrbracket_M \times \llbracket \psi \rrbracket_M$. By definition of exhaustive answerhood, $\llbracket \psi \rrbracket_M \subseteq \llbracket \phi \rrbracket_{M,w'}$ for some $w' \in \mathcal{W}$. Hence, $R \subseteq \llbracket \phi \rrbracket_{M,w'} \times \llbracket \phi \rrbracket_{M,w'}$ for some $w' \in \mathcal{W}$. By Lemma 1, $\models_M^m \phi$.

- For the other direction, assume $\psi \models^m \phi$. Let M be a Kripke structure.⁷ If $\llbracket \psi \rrbracket_M = \emptyset$, then we are done, since exhaustive answerhood holds trivially. Otherwise, choose $\mathcal{W}' = \llbracket \psi \rrbracket_M$, and choose $R' = \mathcal{W}' \times \mathcal{W}'$. Let M' be the Kripke structure induced by M , \mathcal{W}' and R' . By Lemma 1, $\models_{M'}^m \psi^\sharp$. Since $\psi \models^m \phi$, we have $\models_{M'}^m \phi^\sharp$. By Lemma 2, $R'w$ is a ϕ -set for any $w \in \mathcal{W}'$. In other words, $R'w \subseteq \llbracket \phi \rrbracket_{M,w}$. Now choose some $w \in \mathcal{W}'$. Since $R'w = \mathcal{W}' = \llbracket \psi \rrbracket_M$, we have $\llbracket \psi \rrbracket_M \subseteq \llbracket \phi \rrbracket_{M,w}$

B. Finite models in first-order modal logic

In this appendix we consider how many possible worlds are needed to satisfy or falsify a satisfiable or falsifiable formula ϕ in first-order modal (predicate) logic. In particular, we bound the number of worlds needed for question entailments.

DEFINITION 6 (Modal footprint). *The two columns of dual, mutually recursive equations below define the positive modal footprint $|\phi|_+$ and negative modal footprint $|\phi|_-$ for certain modal formulae ϕ .*

$$\begin{array}{ll}
|\phi|_+ = 0 & \text{if } \phi \text{ is atomic} & |\phi|_- = 0 & \text{if } \phi \text{ is atomic} \\
|\neg\phi|_+ = |\phi|_- & & |\neg\phi|_- = |\phi|_+ & \\
|\phi \wedge \psi|_+ = |\phi|_+ + |\psi|_+ & & |\phi \vee \psi|_- = |\phi|_- + |\psi|_- & \\
|\phi \vee \psi|_+ = \max(|\phi|_+, |\psi|_+) & & |\phi \wedge \psi|_- = \max(|\phi|_-, |\psi|_-) & \\
|\diamond\phi|_+ = 1 + |\phi|_+ & & |\square\phi|_- = 1 + |\phi|_- & \\
|\square\phi|_+ = 0 & \text{if } |\phi|_+ = 0 & |\diamond\phi|_- = 0 & \text{if } |\phi|_- = 0 \\
|\exists x.\phi|_+ = |\phi|_+ & & |\forall x.\phi|_- = |\phi|_- & \\
|\forall x.\phi|_+ = 0 & \text{if } |\phi|_+ = 0 & |\exists x.\phi|_- = 0 & \text{if } |\phi|_- = 0
\end{array}$$

Clearly, $|\phi|_+ = |\neg\phi|_-$ and $|\neg\phi|_+ = |\phi|_-$ for all ϕ .

THEOREM 2. *Suppose that M is a Kripke structure, and w is a world in M .*

- *If $M, w \models \phi$ for some formula ϕ for which $|\phi|_+$ is defined, then there exists a set of at most $|\phi|_+$ worlds in M , such that removing every world in M except w and these worlds gives a reduced Kripke structure M' such that $M', w \models \phi$.*
- *If $M, w \not\models \phi$ for some formula ϕ for which $|\phi|_-$ is defined, then there exists a set of at most $|\phi|_-$ worlds in M , such that removing every world in M except w and these worlds gives a reduced Kripke structure M' such that $M', w \not\models \phi$.*

⁷ No assumptions are made on the accessibility relation of M .

Proof. By mutual structural induction on ϕ .

For our purposes, it is important that removing worlds from a Kripke structure preserves the reflexivity, transitivity, and symmetry of its accessibility relation.

Suppose now that $\phi = \phi_1 \rightarrow \phi_2$, where ϕ_1 and ϕ_2 are both conjunctions of formulae of the form $?\vec{x}. \psi$. Then $|\phi|_+$ and $|\phi|_-$ are both at most 2. Therefore, given a Kripke structure M satisfying the constraints in Table I, and a world w in M where ϕ is satisfied (or falsified), we can use the theorem above to reduce M to w and at most two other worlds while still satisfying (respectively, falsifying) ϕ .

In the context of Section 5, this result tells us that GS's question semantics has an extensional interpretation using 3 possible worlds, that is, 2^3 truth values. In fact, because every atomic formula in ϕ is buried under a modal \Box , we can identify the current world w with one of the two other worlds and still satisfy (respectively, falsify) ϕ . In other words, GS's question semantics has an extensional interpretation using 2 possible worlds, that is, 2^2 truth values as described in Section 5.