

A DYNAMIC ANALYSIS OF EPISTEMIC POSSIBILITY

Chris Brumwell

University of Amsterdam, ILLC

`chrisbrumwell@gmail.com`

10/31/08

Contents

1	Introduction	3
2	A Non-Truth Conditional Theory of Epistemic Modality	5
2.1	Core Data	5
2.2	Three Theories	6
2.2.1	Contextualism	7
2.2.2	Relativism	9
2.2.3	Formal Articulations of Contextualism and Relativism . .	10
2.2.4	The Dynamic Theory	12
2.2.5	Formal Articulation of the Dynamic Theory	16
2.3	Complex Data	18
2.3.1	The Objective Pull	18
2.3.2	Warranted Assertability	22
2.3.3	Time Lag Problems	26
2.4	Conclusion	29
3	Salient Possibilities and Non-Informative Content	30
3.1	What is a Salient Possibility?	31
3.2	Epistemic Modals as Answers to Questions	34
3.3	Formal Theory of Salient Possibilities	39
3.3.1	Language and Syntax	39
3.3.2	Inquisitive Semantics for Information States	40
3.3.3	Salient Possibilities	45
3.3.4	Semantics for Contexts	48
3.4	Conclusion	52
4	Cognitive Aspects of Non-Informative Content	54
4.1	Why Belief States Contain Salient Possibilities	54
4.2	Yalcin’s Theory	58
4.3	Formal Theory of Attitude Reports	60
4.3.1	Belief States as Contexts	60
4.3.2	Semantics for Attitude Reports	61
5	Conclusion and Further Issues	64
5.1	General Account of Indirectly Relevant Answers	64
5.2	Epistemic Necessity	65
5.3	Conclusion	66

1 Introduction

Suppose you're on a peaceful raft floating down a quiet river, and all of a sudden your friend tells you that there might be crocodiles in the water. Neither of you knows for sure what's swimming around down there, and you glance over the side of your boat into the water. You don't see anything, and later find out that there have never been crocodiles in the river. Had your friend told you something true? After all, neither of you knew at the time that there weren't any crocodiles, so in a sense she was right to say it. Certainly, at least, you were told something non-trivial. But were you given any information? Did your friend's statement change your beliefs? You definitely thought twice about going in the water after what she said. If nothing else, she spoiled your peaceful boat ride.

This thesis will address these and other questions about epistemic modality, focusing mostly on epistemic possibility. Modality is the feature of linguistic meaning that expresses necessity or possibility, and it has many different incarnations in natural language. Possibility can be expressed relative to certain bodies of rules, sets of beliefs, certain states of affairs, and much else. Epistemic modality expresses necessity or possibility relative to some body of knowledge or evidence. The assertion that 'there might be crocodiles in the water' in the paragraph above expresses epistemic modality. In English, the modal expressions 'might,' 'may,' 'must,' 'perhaps,' 'can,' 'could,' 'it is possible that,' and their close relatives often express epistemic modality. We refer to statements expressing epistemic modality as 'epistemic modals'.

Below, we will argue for a non-truth conditional analysis of epistemic modality, and will further argue that the content of epistemic modals has a non-informative dimension. Both of our analyses will be presented in a dynamic semantics framework. While this thesis is primarily an analysis of epistemic modality, it can also be seen as a case study in the usefulness of the dynamic perspective on natural language. Hence, it is both an argument for a particular view of epistemic modality and an argument for a dynamic theory of meaning.

The thesis is divided into three main parts. First, we argue for a non-truth conditional analysis of epistemic modality. In this section, we also consider two popular truth conditional theories of epistemic modality. Our non-truth conditional theory is developed from a dynamic theory of meaning and communication, both of which are sketched in this section.

In the second part, we argue that the content of epistemic modals has a non-informative dimension. Specifically, we hold that this non-informative content of epistemic modals consists in their ability to draw attention or awareness towards certain possibilities. We support this claim with intuitive considerations regarding communication and conversation as well as data regarding universally quantified epistemic modal formulas and epistemic modals used to answer questions. Our considerations regarding communication and conversation develop the informal idea of a 'salient possibility', a notion we later formalize and put it to

crucial use in our formal theory of questions and answers, which accounts for both informative and non-informative answers to questions. We make some brief remarks concerning what the notion of a non-informative answer entails for the semantics of questions.

Third, we argue that there are cognitive aspects of the non-informative content of epistemic modals. Rather than simply being a tool for the management and organization of a discourse, we will argue that the non-informative content of epistemic modals plays a significant role in belief and thought. Here, we will focus mostly on data involving epistemic modals in attitude reports, though we also briefly discuss epistemic modals used to give advice and warnings.

Below, we assume minimal exposure to the central concepts of dynamic semantics. However, we assume some knowledge of possible worlds semantics and pragmatics and a modest understanding of the semantics and pragmatics of questions.

2 A Non-Truth Conditional Theory of Epistemic Modality

In this section, we present core pieces of linguistic data concerning epistemic modality and three theories that can account for this data: the contextualist theory, the relativist theory, and the dynamic theory. We will then examine more complex data, and evaluate both the relevance of this data and the various levels of success each theory has in accounting for it. We will conclude in favor of the dynamic theory.

2.1 Core Data

A fundamental fact about epistemic possibility is that individuals in different contexts can disagree about whether or not a certain statement ‘possibly ϕ ’ should be accepted without the individuals in either party forming false beliefs. In other words, there can be disagreement about what is epistemically possible without disagreement about the way the world is. The following example from DeRose (1991) nicely highlights this fact:

Cancer Test Case 1A

Background information for DeRose’s ‘Cancer Test Cases’: John has some symptoms indicative of cancer, so his doctors decide to run a diagnostic test on him. If the tests are negative, then John does not have cancer, but if they are positive then John could have cancer but more tests must be administered to find out.

Bill: I heard that John has cancer, is it true?

Jane (John’s wife): It’s possible that John has cancer. He has some of the symptoms. But it’s by no means certain that he’s got it. They’ve run a test on him which may rule cancer out, but they won’t tell us the results of the test until tomorrow.

John’s doctor (the same day): It’s impossible that John has cancer, so we should start planning tests for other diseases.

In Cancer Test Case 1A, Jane and John’s doctor disagree about whether or not John might have cancer. Let’s suppose that John’s doctor is right, and that John doesn’t have cancer. Crucially, then, Jane and John’s doctor don’t have any conflicting beliefs (at least not concerning John’s cancer): John’s doctor believes that John doesn’t have cancer, and although Jane lacks this belief she doesn’t believe that John has cancer. In order for Jane to accept John’s doctors beliefs she doesn’t need to give up any of her beliefs about the world. All she needs to give up is her lack of knowledge.

Thus, we have the following principle about epistemic possibility:

INCONSEQUENTIAL DISAGREEMENT: There can be disagreement concerning statements of the form ‘it is possible that ϕ ’ without there being disagreement about the state of the world.

The intuitive explanation of this principle is quite simple. There is disagreement concerning what is epistemically possible because judgments about what is epistemically possible are made relative to different bodies of knowledge. John’s doctor has access to a stronger body of knowledge than Jane, so understandably his appraisal of whether or not it is possible for John to have cancer differs from hers.

Inconsequential Disagreement is complemented by another equally fundamental fact about epistemic possibility and disagreement: speakers can correct and deny one another’s assertions of epistemic possibility. For example, the following dialogue is perfectly felicitous:

Boston

A: Joe might be in Boston.

B: No, I just saw Joe in Berkeley. He can’t be in Boston.

A: Oh ok, I guess I was wrong.

This dialogue suggests the following principle:

CORRECTABILITY: Speakers uttering statements of the form ‘it is possible that ϕ ’ can be corrected.

It’s important to note that not all kinds of assertions can be corrected by other speakers. For example, parallel facts do not seem to hold with respect to all assertions expressing personal taste.

A: Apples are tasty.

#B: No, I’ve never eaten a tasty apple.

#A: Oh ok, I guess I was wrong.

Speakers cannot easily be wrong about their own personal tastes, however as the data shows us above they can easily be wrong about what’s epistemically possible.

2.2 Three Theories

There is a bit of tension between the two facts we have seen so far. The first suggests that speakers can disagree on epistemic possibility statements but, in a sense, both be ‘right’. After all, the disagreement witnessed in Cancer Test 1A didn’t result in any significant disagreement between Jane and John’s doctor. However, when A and B disagree in the dialogue above, B is able to felicitously

correct A, suggesting that there is some sort of significant disagreement between them. The theories we will examine in this section all provide some way of accounting for both CORRECTABILITY and INCONSEQUENTIAL DISAGREEMENT, and hence resolve this apparent tension.

2.2.1 Contextualism

Contextualist theories hold that epistemic modals are indexical, i.e. that their meaning is determined by their context of utterance. More specifically, contextualists hold that the meaning of an epistemic modal uttered in a context C is determined by some body of knowledge that is salient in C. Just as the meaning of ‘Chris lives nearby’ depends on the location of the context in which it’s uttered, for contextualists the meaning of ‘Chris might have a dog’ depends on certain epistemic features of the context in which it’s uttered.

Contextualist theories also hold that statements expressing epistemic possibility have truth conditions. This feature of their theory is rarely motivated. Contextualist theorists, rather, claim that they find epistemic possibility claims intuitively true or false.

The essential features of contextualism are stated in the following:

Core of Contextualism.

- (i) A statement expressing epistemic modality of the form ‘it is possibly the case that ϕ ’ uttered in context C is true if and only if ϕ is consistent with some contextually salient body of knowledge K.
- (ii) A statement expressing epistemic modality of the form ‘it must be the case that ϕ ’ uttered in context C is true if and only if ϕ is entailed by some contextually salient body of knowledge K.
- (iii) The body of knowledge K is determined by the knowledge of some contextually salient community of people.

Clauses (i) and (ii) state the indexicality of epistemic modals and specify their truth conditions, and (iii) gives a (fairly weak) constraint on which bodies of knowledge can count as salient in a context.

The philosophical and linguistic literature on epistemic modality is full of contextualist theories. Perhaps the earliest contextualist theory is given in Moore (1962), which contains the following:

People in *philosophy* say: The propositions that I’m not sitting down now, that I’m not male, that I’m dead, that I died before the murder of Julius Caesar, that I shall die before 12 tonight are ‘logically possible’. But it’s not English to say, with this meaning: It’s possible that I’m not sitting down now etc. this only means ‘It is not certain that I am’ or ‘I don’t know that I am’.

Strictly speaking this isn't a contextualist theory because Moore does not appeal to the knowledge of a community of people, but only of the speaker. However, in other respects the theory is contextualist. Since Moore, the contextualist theory has undergone a great deal of refinement and re-articulation. Some especially noteworthy articles in the development of the contextualist theory are Hacking (1967), Teller (1972), and DeRose (1991), and Gillies and von Fintel have defended a theory very much like contextualism (which we discuss below) in several recent articles. These later theories mainly disagree about how a body of knowledge K is determined in a context of utterance. Particular contextualist theories that put constraints on K lead to a more empirically robust and explicit theory than 'Core of Contextualism' above; 'Core of Contextualism' is really the common commitment of all contextualist theories, and not a particular contextualist theory. However, in this paper we will argue against *any* contextualist theory, i.e. any theory committed to 'Core of Contextualism', and hence it is easier to think of contextualism about epistemic modality as being expressed by the general theory 'Core of Contextualism' as opposed to a more specific articulation that builds on this common commitment.

If we take propositions to be sets of possible worlds (see Stalnaker (1984)), contextualist theories are formally a special case of Angelika Kratzer's (Kratzer (1977)) classic theory of modals as quantifiers over possible worlds. In Kratzer's analysis of modality, a statement 'possibly ϕ ' evaluated relative to conversational background CB is true if and only if ϕ is true in some world $w \in CB$, and dual truth conditions are given for necessity. Interpreting K as the conversational background relative to which epistemic modals are evaluated, applying Kratzer's analysis we get 'possibly ϕ ' evaluated relative to K is true if and only if ϕ is true in some world $w \in K$, i.e. if and only if ϕ is consistent with K . We will give a more formal expression of contextualism below.

Since contextualism treats epistemic modals as indexical expressions, it easily predicts INCONSEQUENTIAL DISAGREEMENT. To see this, note that given two bodies of knowledge K_1 and K_2 , it can easily occur that ϕ is consistent with K_1 and inconsistent with K_2 . Then an utterance of 'it is possible that ϕ ' will be true in a context where K_1 is the contextually salient body of knowledge and false where it's K_2 , explaining the diverging opinions on the acceptability of epistemic possibility statements. Moreover, the belief that ϕ is consistent with K_1 is compatible with the belief that ϕ is inconsistent with K_2 , explaining why these diverging opinions don't give rise to incompatible beliefs.

Contextualism predicts CORRECTABILITY by positing that the content of an epistemic modal is fixed by the knowledge of some *community* rather than the knowledge of an individual. A speaker can be wrong about what the knowledge of the community rules out, and hence can be corrected by other members of the community. But note that CORRECTABILITY would not be predicted if the contextually salient body of knowledge was taken to be the knowledge of the speaker.

Given that a speaker is being cooperative, she should be the best authority on her own knowledge, and this fact would render any correction infelicitous.

2.2.2 Relativism

Relativist theories hold that an epistemic modal has a constant meaning across all contexts of utterance, but varies in truth value based on who is assessing it. To make sense of this claim, relativists utilize the notion of ‘assessment relative truth’, or truth in a context of assessment. Utterances of epistemic modals are then only assigned a truth value relative to a context of assessment, so such utterances are never simply true or false. For purposes of this paper a context of assessment will be thought of as a person, i.e. an assessor.

Using this notion of a context of assessment, relativism can be characterized by the following:

Relativism.

- (i) Epistemic modals are not true or false simpliciter, but, where C is some context of assessment, true-as-assessed-in-C or false-as-assessed-in-C.
- (ii) A statement expressing epistemic modality of the form ‘it is possibly the case that ϕ ’ uttered in context C is true-as-assessed-in-context-C’ if and only if ϕ is consistent with the knowledge of the assessor in C’.
- (iii) A statement expressing epistemic modality of the form ‘it must be the case that ϕ ’ uttered in context C is true-as-assessed-in-context-C’ if and only if ϕ is entailed by the knowledge of the assessor in C’.

Clause (i) of RELATIVISM posits the new notion of assessment relative truth. Clauses (ii) and (iii) at first make relativism and contextualism seem quite similar, but there is a subtle, yet important, difference between the two. Relativists hold that the meaning of an epistemic modal is constant across all contexts, but that its relative truth or falsity will vary with its context of assessment. Contextualists, on the other hand, hold that the meaning of an epistemic modal varies in different contexts, but that a token utterance of an epistemic modal is true or false simpliciter.

Though contextualist theories of epistemic modality have been popular for quite some time, relativism is a very new theory. Relativism about epistemic modality was first articulated in MacFarlane (ms) and Egan, Hawthorne, and Weatherston (2005). Our exposition of relativism draws largely on these two articles.

As with contextualism, it is very easy to see how relativism accounts for IN-CONSEQUENTIAL DISAGREEMENT. An epistemic modal can be true as assessed in one context yet false as assessed in another, explaining the disagreement with respect to the acceptability of the assertion: different assessors only agree with what

is true in *their* context of assessment. But since this disagreement is ultimately traceable to nothing more than the possession of different bodies of knowledge, this disagreement need not lead to contradictory beliefs about the world. Individuals disagree because their knowledge rules out different propositions, not because they believe inconsistent facts about the way the world is.

But relativism's account of correctability differs significantly from contextualism's. For a relativist, a speaker can be corrected even when what she utters is true in her context of assessment. Such a correction takes place when a speaker's utterance is false in her interlocutor's context of assessment, a situation that occurs when an interlocutor has more knowledge than the speaker. The interlocutor then passes knowledge to the speaker that allows her to judge her original utterance to be false, and hence corrects her ignorance about the world, not the absolute truth value of her utterance. In the example 'Boston', A is corrected because B has more knowledge than her concerning Joe's location. When B corrects A, he makes explicit the proposition that rules out Joe being in Boston, and when A adds this proposition to her knowledge she will no longer take it to be epistemically possible that Joe is in Boston.

Thus, correction for a relativist is not about one speaker pointing out that another speaker has made an absolutely 'false' utterance. Rather, correction is about one speaker possessing more information than another, and then making this information explicit to the other party. After this information exchange has taken place, the speaker who has been corrected will no longer accept the content of her corrected utterance. Hence, a form of correction does take place. Crucially, though, this differs greatly from the sort of correction appealed to by contextualism, in which one speaker simply points out that another speaker has said something false.

2.2.3 Formal Articulations of Contextualism and Relativism

Though the informal account of epistemic modality given by contextualism differs greatly from relativism, formally the two accounts are very similar. The main difference between them concerns the definition of truth at a context of utterance. For a contextualist, this is an unproblematic notion to define. Contexts of utterance determine some salient body of knowledge, and the evaluation of an epistemic modal in that context is determined by what is consistent with or entailed by this knowledge. A relativist, however, can only define a partial or 'gappy' notion of truth in a context of utterance. This is because epistemic modals are assigned different truth values relative to different contexts of assessment, and a single utterance can have several contexts of assessment even though it only has one context of utterance. Our formal articulations of the theories is largely original (many contextualist theorists do not formally state their theory), though influenced by MacFarlane (ms). We begin by defining a simple propositional language with an epistemic modality operator.

Definition 2.1. Modal Propositional Language

Fix a set φ of propositional symbols. The modal propositional language MPL is defined as the smallest set such that:

- (i) If $p \in \varphi$, then $p \in \text{MPL}$.
- (ii) If $\phi \in \text{MPL}$, then $\neg\phi \in \text{MPL}$.
- (iii) If $\phi, \psi \in \text{MPL}$, then $\phi \wedge \psi \in \text{MPL}$.
- (iv) If $\phi \in \text{MPL}$, then $\diamond\phi \in \text{MPL}$.

Disjunction, material implication, and epistemic necessity can be defined in the usual ways.

Contextualists and relativists will appeal to a similar notion of satisfaction relative to a body of knowledge, which we define below.

Definition 2.2. (Satisfaction with Respect to a Body of Knowledge)

Let W be the set of all functions from φ to $\{0,1\}$; we call W the set of possible worlds. Call a non-empty subset K of W a body of knowledge. Let $\phi \in \text{MPL}$, let $w \in W$, and let K be a body of knowledge. We recursively define the notion of satisfaction with respect to (w, K) as follows:

- (i) p is satisfied by (w, K) if and only if $w(p) = 1$.
- (ii) $\neg\phi$ is satisfied by (w, K) if and only if ϕ is not satisfied by (w, K) .
- (iii) $\phi \wedge \psi$ is satisfied by (w, K) if and only if ϕ is satisfied by (w, K) and ψ is satisfied by (w, K) .
- (vi) $\diamond\phi$ is satisfied by (w, K) if and only if there is a $v \in K$ such that ϕ is satisfied by (v, K) .

As mentioned above, the two theories use this notion differently to define truth in a context of utterance. For a contextualist, every context of utterance C determines a body of knowledge K_C . Then, truth at a world w and context of utterance C is simply satisfaction by (w, K_C) .

Definition 2.3. (Contextualist Truth)

Let $\phi \in \text{MPL}$, $w \in W$, and let C be a context of utterance. Then ϕ uttered in C is true_{Con} if and only if ϕ is satisfied by (w, K_C) .

Relativists will appeal to a similar notion for the truth of an utterance in a context of assessment.

Definition 2.4. (Assessment Relative Truth)

Let $\phi \in \text{MPL}$, $w \in W$, let C be a context of utterance, and let A be a context of assessment. Then ϕ uttered in C is true-as-assessed-in- A if and only if ϕ is satisfied by (w, K_A) , where K_A is the knowledge of the assessor in A

However, since an utterance made in a context of utterance can have several contexts of assessment, relativists cannot build the notion of truth in a context of utterance from assessment relative truth. However, relativists can recover a gappy notion of truth by quantifying over contexts of assessment. This will create truth-value gaps for epistemic modal formulas, but no gaps for non-modal formulas.

Definition 2.5. (Relativist Truth)

Let $\phi \in \text{MPL}$, and let C be a context of utterance. Then ϕ uttered in C is true_{Rel} if and only if for all contexts of assessment C' , ϕ uttered in C is true-as-assessed-in- C' .

So, the key formal difference between contextualism and relativism is that relativists have a gappy notion of truth for modal formulas, and contextualists have no such gaps. This corresponds to the philosophical difference between the theories: contextualists think of epistemic modals as having their meaning, and hence truth-value, fixed by their context of utterance, whereas relativists think of the assessment relative truth of an epistemic modal as varying with its context of assessment, and hence cannot recover an absolute notion of truth for utterances of epistemic modals.

With this formal understanding of both theories, the differing notions of correction become more transparent. Recall the ‘Boston’ example above, let C be the context for A ’s first utterance, and let K_B be B ’s knowledge, let K_{A_1} be A ’s knowledge before she is corrected by B , and let K_{A_2} be A ’s knowledge after she is corrected by B . For contextualist theories, the body of knowledge K_C will include A and B ’s knowledge. Therefore, A ’s utterance is not true_{Con} because B ’s knowledge is inconsistent with Joe being in Boston, and B ’s correction is felicitous because she points this out. However, on a relativist theory, A ’s utterance is true-as-assessed-in- A_1 , false-as-assessed-in- B , and false-as-assessed-in- A_2 . Here, B ’s correction is not felicitous because she points out that A ’s utterance is ‘false’ in any absolute sense. Instead, her correction is felicitous because it changes A ’s information, bringing her to a context of assessment in which she judges her previous utterance to be false.

2.2.4 The Dynamic Theory

The two accounts mentioned above assume that epistemic modals are, in some sense, truth conditional. We now consider a non-truth conditional alternative: the dynamic theory. Not only does the dynamic theory give a non-truth conditional analysis of epistemic modality, but also utilizes a conception of linguistic

meaning that differs fundamentally from a truth-conditional conception of linguistic meaning. Rather than taking the meaning of a declarative utterance to be its truth conditions, a proponent of the dynamic theory takes its meaning to be its potential to change an agent's cognitive state. Dynamic theorists replace the classic slogan 'To know the meaning of a sentence is to know its truth conditions,' with the slogan 'To know the meaning of a sentence is to know its impact on an agent's cognitive state'. This is called a 'dynamic' conception of meaning because it holds that the meaning of a sentence consists in its ability to *change* something.

Dynamic theories of meaning have been very popular in linguistics and logic since the 1980s, with some classics in the field being Heim (1988), Kamp (1981), Groenendijk and Stokhof (1991), and Veltman (1996). Dynamic theories have yielded insightful analyses of pronouns, presupposition, defaults, modal subordination, and questions. In our exposition of the dynamic theory of epistemic modality, we draw largely on Veltman (1996).

To understand the difference between a truth-conditional conception of meaning and a dynamic conception of meaning, consider the sentence 'snow is white'. A truth-conditional theorist will analyze the meaning of this sentence by devising a theory according to which the sentence is true if and only if snow is white. For a dynamic theorist, the meaning is identified with a function that adds the information that snow is white to an agent's information. Its meaning is not some truth conditions, but its potential to provide the information that snow is white. Truth is not mentioned at a fundamental level in the dynamic conception of meaning.

In this section, since we are only considering certain features of declarative sentences, we will focus on how sentences change an agent's information about the world. We call an agent's information about the world her *data set*. Thus, for the sentences we deal with in this section, we can think of the dynamic theory as claiming that the meaning of a sentence is its impact on a data set. We will call the effect of a sentence ϕ on a body of information its 'update' of that information, and say that sentences perform 'updates' on pieces of information. The intuition behind this terminology is that formulas act on bodies of information by updating them with new facts.

On a dynamic conception of meaning, communication is a process of coordinating distinct bodies of information. Each participant in a conversation possesses a different body of information, and through the dialogue the participants change each others' bodies of information. This updating process brings their bodies of information more in sync with one another, as each participant is expected to only propose updates that would bring her interlocutors' information more in sync with her own. In the formal articulation of our theory, we will make this dynamic conception of conversation more precise.

Given this picture of conversation, epistemic modals play the role of sharing, in a sense, a lack of information. Rather than adding information, epistemic modals act as 'information indicators' that reveal properties of data sets. Both semantic and pragmatic factors play critical roles in the ability of epistemic modals to

function as information indicators. Semantically, epistemic modals ‘test’ data sets for consistency with formulas; they do not perform non-trivial changes to an agent’s information. An epistemic possibility modal $\Diamond p$ maps an information state that is consistent with p to itself, and maps all others to the empty set (which is thought of as the contradictory information state, for reasons explained below). Information states consistent with p ‘pass’ the test, whereas others ‘fail’.

This semantics for epistemic modals, coupled with some pragmatic principles based on Grice’s maxim of quality (see Grice (1975)), allows them to play the role of information indicators. The maxim of quality requires speakers to only say what they believe. Translated into dynamic terminology, the maxim of quality is the following: only assert sentences that perform a non-trivial update on your information state. If a sentence changes an information state into another non-empty information state, then it has added some information to it. However, the sentences an agent believes should not provide her with new information, and hence any sentence she believes should only perform a trivial update on her information.

The maxim of quality governs assertions, but there is a similar and equally plausible pragmatic principle governing the acceptance of assertions: do not accept anything you cannot consistently believe. After all, rational agents must retain consistent sets of beliefs. In dynamic terms this principle becomes: reject a sentence if it maps your information to the empty set. Since any inconsistent set of information can be represented as a null or empty set of information, this dynamic articulation of the principle is very natural.

With these principles, we can now explain how epistemic modals function as information indicators. If an agent S asserts $\Diamond p$, then everyone she is speaking to will know that since she is obeying the maxim of quality, she does not have information that is inconsistent with p . Otherwise, her assertion would map her information state to the empty set, in violation of this maxim. Moreover, if an agent accepts $\Diamond p$, then her information must be consistent with p , because otherwise this assertion would map her information to the empty set. Thus, assertions of epistemic modals indicate the asserters lack of information, and acceptances of epistemic modals indicate other conversational participants’ lack of information, so an asserted and accepted utterance of $\Diamond p$ indicates that no one in the conversation possesses the information that $\neg p$.

Epistemic necessity modals work in a similar manner: $\Box p$ is also a test, but rather than testing for consistency with p , it tests whether or not an information state already contains the information that p . Specifically, $\Box p$ functions as the identity map on information states that contain the information that p , and maps all others to the empty set. Thus, if someone asserts $\Box p$, then by the maxim of quality they must have information that p is the case, and similarly for speakers who accept $\Box p$.

Some readers may be uneasy with this account of epistemic necessity, since it seems to suggest that speakers without the information that p can never accept

$\Box p$. In a sense, this is right: if you lack the information that p is the case, *prima facie* why would you accept that it must be the case? However, there is a pragmatic story to tell here too. Upon hearing a speaker S assert $\Box p$, one can trust that S has accurate information, and choose to include p in one's own body of information. In other words, one can infer that p is the case based on S 's assertion of $\Box p$, the assumption that S is following the maxim of quality, and the assumption that S has accurate information. This inference expands the information state of a speaker who previously lacked the information that p , and hence allows speakers lacking the information that p to ultimately accept assertions of $\Box p$.

Putting all these considerations together, we get the dynamic theory:

The Dynamic Theory.

- (i) Epistemic modals are not truth conditional constructions, but instead act as information indicators.
- (ii) An epistemic modal 'it is possible that ϕ ' maps a body of information I to the empty set if I is inconsistent with ϕ , and maps I to itself otherwise.
- (iii) An epistemic modal 'it must be the case that ϕ ' maps a body of information I to the empty set if ϕ is not entailed by I , and maps I to itself otherwise.
- (iv) An epistemic modal E is assertable or acceptable by a speaker S if E does not map S 's information to the empty set.

The dynamic theory easily accounts for both INCONSEQUENTIAL CONSEQUENT and CORRECTABILITY. As with relativism, since diverging judgments concerning epistemic modals are traceable to different data sets, this disagreement need not lead to contradictory beliefs about the world. In general, agents with less information will think that more is epistemically possible than agents with more information. But while there is a difference in the amount of information possessed by these individuals, their different bodies of information do not need to be in conflict with one another.

A dynamic theorist's account of CORRECTABILITY is also similar to a relativist's account. Since epistemic modals are neither true nor false, a dynamic theorist cannot appeal to the falsity of an epistemic modal to explain correction. Instead, a dynamic theorist thinks of correction in terms of information growth. A speaker's data set changes in the process of correction, and as a result of this change she no longer accepts her original utterance. In 'Boston', for example, A 's data set was changed when she accepted B 's assertion that Joe is in Berkeley, and her original utterance is not acceptable relative to this new information state.

2.2.5 Formal Articulation of the Dynamic Theory

Formally, we express the dynamic theory using the system developed in Veltman (1996). In this system, formulas from MPL are interpreted as functions from data sets to data sets, where a data set is a set of possible worlds. We will also discuss certain notions of consistency and their relation to the dynamic picture of conversation, as well as define a notion of truth within a dynamic theory of meaning and show that under this conception of truth, relativism is formally similar to the dynamic theory.

Recall the language MPL, and recall that possible worlds are functions from propositional symbols to the truth values $\{0,1\}$. As we said above, we think of data sets as sets of possible worlds. Intuitively, the information represented by a set of possible worlds is the information that is consistent with the possible worlds in that set. So, for example $\{w \in W : \text{snow is white in } w \text{ and grass is green in } w\}$ represents the information that snow is white and grass is green. Here, we think of information as something that eliminates possible ways the world could be. Having the information that snow is white and grass is green eliminates possibilities in which snow is not white and grass is not green.

The empty set has a special status given this conception of information: it represents any inconsistent body of information. Since no possible worlds are compatible with contradictory information, any inconsistent body of information corresponds to the empty set. Hence, we think of updates that map a data set to the empty set as somehow problematic for that information state.

The conception of information as eliminating possibilities is crucial for understanding the semantics we provide for MPL. Propositional symbols p eliminate worlds from an information state that falsify p , negations of propositional symbols $\neg p$ eliminate worlds that satisfy p , and conjunctions of propositional symbols $p \wedge q$ eliminate worlds that falsify either p or q . Epistemic modals, however, do not eliminate worlds from an information state, since these formulas do not carry information but rather act as ‘tests’: as explained above, epistemic modals test information states for consistency with some formula. If an information state passes the test, then it is mapped to itself, and if it fails the test, it is mapped to the empty set.

Definition 2.6. (Update Semantics for MPL)

Let W be the set of all possible worlds, and let $I \subseteq W$. We then define the update of I with a formula of MPL as follows:

- (i) $I[p] = \{w \in I : w(p) = 1\}$
- (ii) $I[\neg\phi] = \{w \in I : w \notin I[\phi]\}$
- (iii) $I[\phi \wedge \psi] = I[\phi][\psi]$
- (iv) $I[\diamond\phi] = I$ if $I[\phi] \neq \emptyset$

(v) $I[\diamond\phi] = \emptyset$ if $I[\phi] = \emptyset$

To illustrate this semantics, we use it to represent a simple conversation. Suppose a conversation has two parties, A and B, and A only has the information that p and q , and B only has the information that r and s . Identifying conversational participants with their information, we get $A = \{w \in W : w(p) = w(q) = 1\}$, and $B = \{w \in W : w(r) = w(s) = 1\}$. In the conversation, A and B will coordinate their information states by uttering formulas that express updates containing the information they each possess. If A utters $p \wedge q$ and B accepts, then B's information state is updated with $p \wedge q$. Likewise, if B utters $r \wedge s$, and A accepts, then A's information state is updated with $r \wedge s$. The result is that A and B end up with the same information state, since $A[r \wedge s] = B[p \wedge q]$.

Using this semantics, one can articulate two notions of consistency for formulas of MPL. Under one notion, a formula is consistent if it is possible to update some information state with that formula without mapping it to the empty set. This notion is called *acceptability*. We say that an information state I accepts a formula ϕ if $I[\phi] \neq \emptyset$, and say that a formula is acceptable if it is accepted by some information state. Intuitively, a formula is acceptable if it is possible for someone to have information that is not in conflict with the formula. Similarly, one could think of acceptability as a kind of learnability: if some information state can accept a formula ϕ , then ϕ is learnable. It is easy to see that no contradiction of classical propositional logic is acceptable, and any non-contradiction is acceptable.

On the other notion of consistency, a formula is consistent if there is some information state that is mapped to itself by the formula. This notion is called *supportability*. We say that an information state supports a formula ϕ if $I[\phi] = I$, and that a formula is supportable if there is an information state that supports it. While acceptability somehow involves learnability or lack of inconsistency, supportability involves a sort of entailment. If a formula is supported by an information state, then the information state already 'contains' the information that formula could provide. Again, no contradiction of classical propositional logic is supportable, and any non-contradiction is acceptable.

These notions depart from one another when one considers modal formulas: the formula $\diamond p \wedge \neg p$ is acceptable, but not supportable: clearly this formula is accepted by W , but it can never be supported. To see this, suppose $I[\diamond p \wedge \neg p] = I$. Then $I[\diamond p] = I$, and hence $I[p] \neq \emptyset$, so there is a $w \in I$ such that $w(p) = 1$. However, $w \notin I[\neg p]$, so $I[\diamond p \wedge \neg p] \neq I$, a contradiction.

This should come as no surprise. After all, one can be in a state where one's information makes it possible to accept $\diamond p$, but then later one can unproblematically learn that $\neg p$ is the case. For instance, one can be in a windowless room and have no information concerning whether it's raining, and hence accept the statement "it might be raining", but later leave the room and learn that it is not raining. This corresponds to updating one's information state with $\diamond p \wedge \neg p$. However, one can never have information that already contains the information

that it might be raining and it's not raining. One's information either supports one claim or the other, never both.

The notions of an acceptable formula and a supportable formula play crucial roles in the dynamic picture of conversation. Keeping in mind the discussion above, the maxim of quality states that one may only assert formulas that are supported by one's data set. Moreover, one can only accept assertions that are acceptable in one's data set. Thus, the pragmatic notions that are so crucial to the dynamic theory of epistemic modality have intuitive and natural counterparts in the formal theory.

The dynamic semantics we defined above allow us to define a notion of truth in a possible world for the non-modal fragment of the language.

Definition 2.7. (Dynamic Picture of Truth)

Let ϕ be a formula of MPL in which the symbol \diamond does not occur, and let w be a possible world. Then ϕ is true in w if and only if $\{w\}[\phi] = \{w\}$.

Clearly this notion cannot be expanded to include the modal fragment of the language, or epistemic possibility and truth would simply collapse. However, we can introduce a notion of *relative* truth for the modal fragment of the language.

Definition 2.8. (Dynamic Relative Truth)

Let ϕ be a formula of MPL of the form $\diamond\psi$, or a conjunction or negation of formulas of this form, and let I be an information state. We say that ϕ is true relative to I if and only if $I[\phi] = I$.

This notion of relative truth is extremely similar to relativism's notion of relative truth: for any propositional symbol p and context of assessment A , $\diamond p$ is true-as-assessed-in- A if and only if $K_A[p] \neq \emptyset$. This shows that the dynamic theory and relativism have a deep formal similarity, the only difference being that relativists are committed to a notion of relative truth whereas dynamic theorists are not: it is definable in their theory, but not necessarily useful. We will argue below that this difference counts heavily against relativism.

2.3 Complex Data

The theories mentioned above all have a degree of plausibility due to their ability to account for the core data. We now examine more complex data, and argue that the dynamic theory offers the best account of the issues it raises.

2.3.1 The Objective Pull

In what we have seen so far, it appears that the evaluation of epistemic modals depends on the knowledge of the participants in a conversation. The watershed between the theories mentioned above is exactly how it depends on this knowledge. In this section, we will examine some data that suggests that the knowledge

relevant to the evaluation of an epistemic modal may extend beyond that of the conversational participants. We will conclude that if what this data suggests is actually the case, then this strongly supports the contextualist theories. However, we will raise some difficulties with this data, and conclude that it does not point to the contextualist theory over the others.

Recall DeRose's Cancer Test Cases, but note Jane will not necessarily respond to Bill by claiming that it is possible that John has cancer; this is demonstrated in another case from DeRose (1991):

Cancer Test Case 1B

Bill: I heard that John has cancer. Is that possible?

Jane: I don't know whether it's possible that John has cancer; only the doctors know. I'll find that out tomorrow when the results of the test are revealed.

According to DeRose, the relevance of the data is the following: "The explanation seems to be that in the first case [Cancer Test Case 1A in section 2.1] Jane is (truly) saying that it is possible that John has cancer relative to the epistemic situation of a fairly small relevant community – perhaps John's family – while in the second case Jane, it seems, is professing to be ignorant as to whether it's possible that John has cancer relative to what is known to a somewhat larger group of people that includes John's doctor". After all, if only Jane's knowledge and Bill's knowledge were relevant to the evaluation of her possibility claim, then it is hard to see how she could felicitously claim that she does not know whether it is even *possible* that John has cancer. Therefore, something more than Jane's knowledge and Bill's knowledge must be relevant to the evaluation of her assertion, and the doctor's knowledge is a likely candidate for this 'something more'.

Other data from DeRose (1991) suggests that an epistemic modal may be sensitive to the knowledge that could be gained by performing certain kinds of investigation (similar data is considered in Hacking (1967) as well). In other words, this data suggests that sometimes not only actual, but potential, knowledge can be relevant to the evaluation of an epistemic modal.

Cancer Test Case 2B

In this case, the test has been run, but not even the doctor knows the results of the test. A computer has calculated the results and printed them. A hospital employee has taken the printout and, without reading it, placed it in a sealed envelope. The policy of the hospital is that the patient should be the first to learn the results. Jane has made an appointment to pick up the results tomorrow. She knows that the envelope with the results has been generated and that nobody knows what the results are. Still, if Bill were to call her to find out the latest news, she might very well say, "I don't yet know whether it's possible that John has cancer. I'm going to find that out tomorrow when the results of the test are revealed."

In this case, Jane is still able to felicitously claim ignorance concerning the possibility that John has cancer. But no one has knowledge that is inconsistent with this possibility, and since Jane knows the hospital's policy on handing out the results of the test, Jane knows that no one has knowledge inconsistent with this possibility. Therefore, it appears that the knowledge that will be gained from reading the test results is somehow relevant to evaluating Jane's assertion.

Explaining this data with a contextualist theory is simple. The contextually salient body of knowledge K can potentially come from anywhere. So, to account for pieces of data like Cancer Test Case 1B and 2B, contextualists can adopt the following two principles:

FLEXIBILITY: In a context of utterance, the contextually salient community can include the speaker, other participants in the conversation, and other salient individuals not participating in the conversation.

EXTENDIBILITY: The body of knowledge K is either the body of knowledge currently possessed by some contextually salient community or certain contextually salient extensions of the body of knowledge currently possessed by some contextually salient community.

In Cancer Test Case 1B, the doctor's knowledge is part of K , explaining Jane's ignorance, and in Cancer Test Case 2B, the extension of Jane's knowledge that results from her reading the test results is part of K . Hacking (1967) and DeRose (1991) incorporate these principles explicitly into their analyses.

While a similar move is open to relativism, going down that path would in some ways compromise the theory. A crucial difference between contextualism and relativism lies in their different accounts of correctness: contextualism posits that epistemic modals are sensitive to knowledge beyond what the speaker possesses, whereas relativism appeals to a notion of assessment relative truth and the growth of knowledge that takes place during conversation. Relativists appeal to assessment relative truth so they do not need to appeal to anything but the knowledge of a single speaker to explain the core facts concerning corrections of epistemic modals. In section 2.3.2 below we will argue that the appeal to community knowledge leads to problematic predictions, so this feature of relativism should be conserved. Thus, ideally relativists would not like to adopt these new principles.

The dynamic theory could also appeal to principles like flexibility and extendibility, but then it too would not differ much from contextualism. The most important difference between the two would be disagreement concerning whether or not epistemic modals are true or false. But Cancer Test Cases 1B and 2B, taken at face value, suggest that epistemic modals are truth conditional: 'knows' is factive, so if Jane can felicitously say that she does not know whether something might be the case, this suggests that there is some fact that will decide the

matter.

Thus, if Cancer Test Cases 1B and 2B are solid data about epistemic modals, this heavily favors contextualism. However, it is not entirely clear that the modality expressed in these examples is an epistemic modality. To see this, consider the following variation on Cancer Test Case 1B:

Cancer Test Case 1B'

Bill: I heard that John has cancer. Is that possible?

Jane: I don't know whether it's possible that John has cancer. In fact, no one knows whether it's possible that John has cancer. The doctors are totally stumped, and don't even know what kinds of tests to run on him.

In this revision, Jane's discourse is still perfectly felicitous. But it is implausible that epistemic possibility is expressed in her statements. Against which body of knowledge should Jane's claims be evaluated? It is certainly no actual body of knowledge, and it is not a body of knowledge gained by performing any salient investigation. A contextualist could say that it should be evaluated relative to a body of knowledge that results from the progression of science. But such a reply leads to more trouble. If an epistemic possibility claim can be evaluated relative to any body of knowledge that results from the progression of science, then epistemic possibility starts to look a lot like just plain truth, since little more than the truth is possible given extremely advanced science.

Thus, the modality expressed in Cancer Test Case 1B' is not epistemic modality. The similarity between Cancer Test Cases 1B and 2B and Cancer Test Case 1B' suggests that the modality in the former cases may not be epistemic either. The other candidate for the type of modality expressed in these examples is *circumstantial* modality. Epistemic modals express necessity or possibility relative to a body of knowledge or information, but circumstantial modals express necessity or possibility relative to a set of circumstances.

The following sentences can all be interpreted as expressions of circumstantial modality:

- (1) John cannot play the trombone.
- (2) You can buy a piece of property during your next turn.
- (3) It's possible for this acorn to grow into an oak tree.

In (1), suppose that John has lung problems. Then given the circumstances, i.e. the state of John's body and the physical requirements of playing the trombone, it is not possible that John will play the trombone. In (2), suppose that the speaker is playing monopoly. Then, given the circumstances defined by the present state of the game and the rules of monopoly, (2) states that it is possible for the addressee to buy property during her next turn. Finally, (3) states that

the salient acorn can, given the laws of plant physiology and the properties of the acorn, grow into an oak tree.

A bit more formally, we think of circumstantial modality in terms of what is possible or necessary given two parameters: some state of affairs and some body of rules governing that state of affairs. We think of a state of affairs as a set of possible worlds s , where the state of affairs is what all worlds in the set have in common, and we think of rules as a relation between states of affairs and possible worlds, where sRw if and only if $w \in s$ and w is governed by the set of rules R . One can think of the worlds w s.t. sRw as ‘completions’ of s that are consistent with the rules represented by R .

In our present case, the relevant state of affairs is the state of John’s body, the relevant rules are the laws of human physiology, and Jane is saying that she does not know whether or not the state of John’s body and the laws of physiology exclude the possibility that John has cancer. Furthermore, it is easy to see how everyone can be ignorant as to whether or not it is circumstantially possible for John to have cancer: such a case would obtain if no one had adequate knowledge of either the state of John’s body or the laws of physiology.

In the cancer test cases, it is easy to confuse what is epistemically possible given the doctor’s knowledge and what is circumstantially possible given the state of John’s body and the laws of physiology because the doctor is presumed to know the state of John’s body and the laws of human physiology. Because of the doctor’s knowledge, there’s an extent to which what is circumstantially possible given John’s body and physiology collapses into what is epistemically possible given John’s doctor’s knowledge: if the circumstances rule something out, John’s doctor should know this, and if they do not rule something out, then unless John has some extra knowledge unrelated to the circumstances his knowledge cannot rule it out either. This is why some of the cancer test cases can express circumstantial modality yet seem to express epistemic modality.

Therefore, we conclude that this data (i.e. Cancer Test Cases 1B and 2B) does not count against relativism or the dynamic theory. It brings to light an interesting connection between circumstantial modality and epistemic modality, but since the modality expressed in this data is not epistemic, it simply cannot be used to call theories of epistemic modality into question.

2.3.2 Warranted Assertability

In recent literature on epistemic modality, relativists have proposed a series of puzzles for contextualist theories (see MacFarlane (ms) and Egan, Hawthorne, and Weatherson (2005)). In these puzzles one party utters an epistemic modal and is being eavesdropped upon by another, and the puzzle is to make sense of the judgments made by both the eavesdroppers and the eavesdroppees. We follow Gillies and von Fintel (ms) and see this as a problem involving the warranted assertability conditions for epistemic possibility modals. In this section,

we will explain this issue about the warranted assertability conditions of epistemic modals, how it poses a problem for contextualist theories, and how it is accounted for by relativism and the dynamic theory.

First, we present some data involving eavesdropping that will be helpful in explaining the warranted assertability problem. Suppose Sally and George are climbing a mountain they have never climbed before, and that there is no one around listening to their conversation. Then the following is perfectly felicitous:

George: That peak might be the summit.

Sally: Well, I've never climbed this mountain before. Perhaps it is.

Next, suppose that there is also an expert climber on the mountain who is eavesdropping on Sally and George. Then the following is also perfectly felicitous:

George: That peak might be the summit.

Sally: Well, I've never climbed this mountain before. Perhaps it is.

Climber: They're wrong; the summit is nowhere close to here.

Thus, even though George did not mean to speak to the expert climber, the climber can felicitously reject George's utterance.

What this data strikingly brings out is a gap between the warranted assertability conditions of epistemic possibility modals and the conditions under which they can be felicitously rejected. In order for George to be warranted in asserting that the peak might be the summit, it suffices that he does not know that it is not the summit. However, the climber can reject the utterance as long as he knows that the peak is not the summit. Thus, there is a gap between the sort of information that justifies assertions of epistemic possibility modals and the sort of information that justifies their rejection.

The existence of this gap is fairly unique to assertions of epistemic modals. Typically, information that justifies an assertion of a sentence is similar to information that justifies a rejection of an assertion. For example, an assertion of p is justified by information that suggests that p is the case, whereas a rejection of an assertion of p is justified by information that suggests that p is not the case. In both cases, the relevant information regards the same issue: whether or not p is the case. The data shows that with epistemic modals we have a puzzling difference. An assertion of $\Diamond p$ appears to be justified as long as the speaker does not have the information that p , but in order to reject such an assertion one does not appeal to information regarding the issue of whether or not the speaker has the information that p . One rather appeals to one's own information in this case, information that addresses a completely different issue.

The existence of this gap suggests an important insight into epistemic modality: assertions of epistemic modals are not assertions of certain *facts*. If an assertion of $\Diamond\phi$ were a claim that some fact is the case, then we would not expect this

gap. Assertions would be justified by information that this fact is the case, and rejections would be justified by information that this fact is not the case. Thus, though the warranted assertability data may at first seem to be somewhat peripheral, it points to a very interesting and important feature of epistemic modals.

Relativism and the dynamic theory can easily explain the existence of this gap. On either of these theories, the assertion and rejection of epistemic modals depends only on the asserter's information and the rejector's information respectively. For relativists, this is the case because only a relative notion of truth can be given for epistemic modals, and their acceptance and rejection is governed by this relative notion. A speaker may assert sentences that she thinks are true in her context of assessment, and should reject sentences she thinks are false in her context of assessment. If the notion of relative truth bears any resemblance to absolute truth, these principles about assertion and rejection are natural parts of a relativist theory. Moreover, they clearly suffice to account for the gap mentioned above.

For a dynamic theorist, the notions of acceptability and support govern the assertion and rejection of utterances: as explained above, a speaker can assert a sentence if it is supported by her information, and a speaker can reject an utterance if it is not acceptable by her information. Therefore, a speaker can assert an epistemic modal $\Diamond\phi$ if ϕ is consistent with her information, and a speaker can reject $\Diamond\phi$ if ϕ is inconsistent with her information. Hence, both relativism and the dynamic theory succeed in orienting the assertion of epistemic modals towards the only information of the asserter, and orienting the rejection of epistemic modals only towards the information of the hearer.

Contextualism, however, cannot handle this gap so easily. By appealing to the knowledge of a community of speakers, the contextualist truth conditions can easily explain why speakers can appeal to their own knowledge in the rejection of epistemic modals. However, it is far from clear why speakers do not need evidence concerning the knowledge of their listeners to support their epistemic possibility assertions. On a contextualist theory, these assertions are claims about the knowledge of others, so it is certainly strange that a speaker does not need evidence about the knowledge of others to support them. Unlike relativism and the dynamic theory, contextualism orients the assertion and rejection of epistemic modals towards the same contextually salient body of information, and hence runs into problems from these eavesdropping scenarios.

Contextualism's failure to explain this gap should not be too surprising. After all, for a contextualist an epistemic modal expresses some fact about the world, viz. a fact about the knowledge of some community. Thus, since it is difficult to explain how different kinds of information could justify the assertion of a fact and the rejection of a fact, it is difficult for contextualists to explain the warranted assertability data.

To solve this problem, contextualists need to find some reason that epistemic modals are an exception to standard rules about justified assertion and rejection.

A recent paper, Gillies and von Fintel (ms), attempts to do just that. In their framework, an utterance of an epistemic modal $\Diamond\phi$ is not a normal assertion of just one proposition, but is a different sort of speech act that ‘puts into play’ a set of propositions. Each proposition in this set is of the form ‘the knowledge of x is consistent with ϕ ’, where x is either an individual or a group of individuals participating in the conversation. According to their theory, a speaker is justified in putting such a set of propositions into play as long as she would be justified in asserting at least one of them.

Strictly speaking, this is not a contextualist theory of epistemic modality. Contextualism makes no mention of ‘putting propositions into play’, and moreover it is not clear what sort of truth conditions epistemic modals possess in the Gillies and von Fintel (ms) picture. However, one can think of it as a descendent of contextualism because each of the propositions put into play by an epistemic modal is of the form ‘ ϕ is consistent with K ’ for some body of knowledge K .

Gillies and von Fintel (ms) goes some way towards solving the problem, because on this theory an epistemic modal somehow involves several different propositions each oriented towards a different participant in the conversation. Moreover, since this account does not invoke a relative notion of truth and takes the meaning of an epistemic modal to be fixed by its context of utterance, it definitely resembles a contextualist theory more than a relativist or dynamic theory. But the theory suffers from two fundamental problems. The first is the vague notion of putting a set of propositions ‘into play’. Gillies and von Fintel (ms) does not say nearly enough about what this is supposed to mean, and it is far from being an intuitive notion.

The second, and larger, problem is the principle that a speaker is justified in putting a set of propositions into play as long as she would be justified in asserting at least one of them. It is difficult to see why this should be the case, perhaps because it is difficult to understand what putting propositions into play is supposed to mean in the first place. What is worse is that Gillies and von Fintel (ms) does not offer either conceptual or independent empirical support for this principle. They simply claim that it’s the case:

But perhaps this relies on a mistaken theory of what the norm of assertion comes to. Perhaps, according to the right story, it turns out that when a speaker utters an underspecified sentence she must – deontic *must*, that is – be in a position to assert each of the propositions she puts in play. And perhaps someone would enjoy arguing for all that. But not us. We’ll leave the norm of assertion untouched and we will say that the speaker only has to be in a position to at out assert one of the propositions she puts in play; any one of them will do.

While their principle, if true, would explain the gap, it is not clear what it means and it is not clear why it should be accepted. Thus, the theory presented in Gillies and von Fintel (ms) is far too vague and ad hoc to help contextualism explain

the gap.

The gap between the assertability and rejectability conditions of epistemic modals is a fundamental feature of epistemic modals. It accounts for why speakers can use the constructions so freely yet still be subject to corrections, and suggests that assertions of epistemic modals are not assertions of facts. Both relativism and the dynamic theory properly account for this feature of epistemic modals in a way that follows naturally from their theories. These accounts are so successful because they easily orient assertions of epistemic modals towards one body of information and rejections of epistemic modals towards a different body of information. In contrast, contextualism is bound to orient the assertion and rejection of epistemic modals towards the same body of information. This is because under a contextualist theory epistemic modals express facts, and hence the assertability and rejectability of epistemic modals must be oriented around the same kind of information, i.e. information concerning whether or not some fact is the case. Hence, the data concerning the warranted assertability conditions of epistemic modals suggests that there are deep problems with no clear resolution in any contextualist theory.

2.3.3 Time Lag Problems

So far, the data we have seen does not bear on the question of whether or not epistemic modals are truth conditional. Though the dynamic theory works well for the core data and the data discussed in 2.3.2, relativism works equally well. Given their formal similarity pointed out in section 2.2.5 this should be no surprise. In this section, we will look at some data that illuminates our intuitions concerning whether or not epistemic modals should be thought of as truth conditional constructions. We will conclude that assigning truth values – even relative truth values – to utterance of epistemic modals makes questionable predictions.

In many simple dialogues, a truth conditional theory of epistemic modals makes sense. Simple sentences that are accepted by all parties in a dialogue can usually be unproblematically thought of as true. Moreover, simple cases involving correction can also be unproblematically thought of as involving falsity: since Bill can respond to Jane by saying “No...you’re wrong” in the ‘Boston’ example, it is easy to make sense of what it means for her utterance to be false in this context.

In more complex cases, however, it is far from clear how to assign truth values to utterances of epistemic modals. One class of cases that brings this out well are called ‘time lag’ cases (see Gillies and von Stechow (2008)). In these situations, one speaker utters an epistemic modal to another at a point in time t , and then these speakers evaluate this utterance from a later point t' . The dialogue below is an example:

Infelicitous Correction

A: The Brown family might come over for dinner. Hide the ice cream.

B: Ok.

[Later, after the Browns didn't show up]

#B: You were wrong, the Browns never came for dinner. I hid that ice cream for nothing.

A: Hey, I never said anything misleading. I said that they *might* come, not that they *would* come.

Here, A and B are evaluating A's earlier utterance from a later point in time, and hence it is a time lag case. The key features of this case are that A's first utterance would not be acceptable in A and B's later context, but that A's first utterance was not entirely wrong either. Intuitively, this is why B's attempted correction is somewhat problematic.

The dynamic theory can explain both of these features easily. A's utterance, if made again in the later context, is unacceptable because A and B have information that it is inconsistent with. However, B's correction seems odd because A's original utterance was unproblematic to accept given the information they possessed at the time. It functioned as an information indicator that displayed that neither A nor B knew that the Browns were not going to show up, and performed this function perfectly fine. Thus, B's attempt to correct A's assertion is infelicitous.

Recall that on the dynamic theory a conversation is a process in which discourse participants coordinate their bodies of information, and epistemic modals function in this process as information indicators. With this picture, B's inability to correct A's utterance is very easy to understand. In B's later context, A's utterance is no longer relevant to the coordination of A and B's present bodies of information, and hence is simply not relevant to their conversation. A's utterance served its coordinating function perfectly well in earlier in the evening, so an attempt to 'correct' it is just infelicitous.

This suggests that felicitous corrections of epistemic modals are only possible when the corrected utterance and the correction are part of the same information coordination process. An epistemic modal can be corrected if accepting it will fail to properly coordinate speakers' information. This happens if one speaker asserts $\diamond p$ and another knows $\neg p$. In such a case, if correction does not take place, the speakers will fail to coordinate their information (since they will not share the information that $\neg p$). Thus, corrections are necessary in these circumstances for proper information coordination. But utterances of epistemic modals in past contexts are generally not part of current attempts to coordinate information, and hence corrections of these past utterances are generally infelicitous.

Relativist theories, however, have a problem with time-lag data. Clearly, it makes no sense for A and B to judge A's original utterance to be true: in their later context they know that the Browns did not come for dinner. Therefore, given relativism, the utterance must be false in the later context. But this makes an inaccurate prediction as well, given the following general principle about falsity:

Remark 2.9. Falsity

(*) If utterance U is false in C, and its falsity is relevant in C, then it's felicitous to say that U is false/wrong/incorrect/etc.

Given principle (*), B's attempted correction in 'Infelicitous Correction' above should be felicitous, and hence relativism makes suspect predictions in time lag cases.

Relativists, then, must explain why cases like 'Infelicitous Correction' should be exceptions to this principle. One way of doing this would be to say that B's utterance seems odd because 'you were wrong' in 'Infelicitous Correction' means something like 'what you said earlier was not justified' rather than 'what you said earlier was not true'. But against this, note that a parallel dialogue involving a false yet well justified assertion is completely felicitous:

Felicitous Correction

Suppose the Brown's call A and say they're coming for dinner, and then the following takes place.

A: The Brown family is coming over for dinner. Hide the ice cream.

B. Ok.

[Later]

B: You were wrong, the Brown's never came for dinner. I hid that ice cream for nothing.

#A: Hey, I never said anything misleading.

In this case, A's assertion is clearly justified but false. However, here there is nothing wrong with B's assertion, suggesting that B should be able to correct justified yet false assertions in this way, contrary to the suggested solution to relativism's problem.

Thus, it is difficult to deal with this data in a truth conditional framework: either assignment of a truth value to A's original utterance leads to a problem. Different aspects of our intuitions pull in different ways. When we consider whether or not the speakers should have accepted the original utterance, we are pulled towards considering it to be 'true'. But when we consider whether or not the speakers should currently accept the utterance, we are pulled towards considering it to be 'false'. The dynamic theory provides a much better framework in which we can address our intuitions about correction. According to the dynamic theory, correction is not governed by truth or falsity, but by proper and improper attempts to coordinate bodies of information. Utterances that do not properly coordinate bodies of information can be corrected, but past utterances that are no longer relevant to the coordination of information should be left alone.

Thus, the assignment of truth values to epistemic modals sets up a poor framework for understanding our intuitions about corrections of epistemic modals.

Rather than seeing epistemic modals as truth conditional constructions, it is better to understand them as tools to regulate the coordination of speakers' information in a dialogue. Accepted assertions of the form $\Diamond\phi$ do not add information about the world, but serve as indicators of who has what sort of information. Both relativism and the dynamic theory incorporate this insight into their account of epistemic modality, a fact that can be seen by their formal similarities. But relativism needs claims about truth and falsity to articulate this account, and such claims ultimately make poor predictions. Thus, despite their similarities, the dynamic theory is a stronger implementation of these insights about epistemic modality, and hence a superior theory.

2.4 Conclusion

In this section, we presented and argued for a dynamic theory of epistemic modality. According to this theory, epistemic modals are not truth conditional constructions that express facts about the world, but rather are 'information indicators'. Asserted and accepted utterances of epistemic modals suggest facts about the information possessed by conversational participants, but do so without expressing any truth conditional content. The appeal to truth is not only unnecessary, but as we argued in section 2.3.4, is also problematic.

Here, we end our discussion of the truth conditions of epistemic modals. However, our non-truth conditional theory will play a crucial role in our discussion below. Though below we focus mainly on the non-informative content of epistemic modals, the formal account we develop makes crucial use of the dynamic theory we articulated in this section.

3 Salient Possibilities and Non-Informative Content

If someone says “there might be elephants in South America” and everyone else in the conversation agrees, then this utterance indicates that no one in the conversation knows that there are not elephants in South America. But is that all this utterance accomplishes? Intuitively, it seems that something more occurs: the utterance also focuses everyone’s attention on the possibility that there are elephants in South America. This is a fundamentally non-informative effect of the utterance, but it is an important effect nonetheless.

In this section, we examine the intuitions elicited above, i.e. the intuition that statements expressing epistemic modality have a non-informative dimension to them. The non-informative dimension of such statements ‘raises awareness’ of possibilities, or directs attention in a conversation to previously ignored or overlooked possibilities. The observation that epistemic modals accomplish something like this is nothing new: it has been observed in Gillies and von Stechow (2007), Yalcin (2008), and Groenendijk, Stokhof, and Veltman (1996) among surely many others. What is new is our formal analysis of this aspect of epistemic modality and our use of this insight to account for data regarding epistemic modals used to answer questions.

The non-informative aspects of epistemic modality are generally ignored in formal semantics. This may be because the philosophical tradition in which formal semantics is rooted has generally thought of truth, reference, and information as the fundamental semantic notions. It is difficult to fit in the notions of ‘awareness of’ or ‘attention to’ possibilities into such a conceptual framework: what is true is true regardless of what anyone is aware of, and bringing someone’s attention to a possibility is far different from providing someone with information regarding what is actually the case. Below, however, we will argue that awareness and attention should not be ignored. Not only does an account of these non-informative aspects of meaning produce an elegant formal account of questions and answers, but it is also needed to understand the truth conditions of certain attitude reports (we argue the latter point in section 4).

In our theory of the non-informative content of epistemic modals, we develop the notion of a ‘salient possibility’. Roughly, a salient possibility is a proposition on which the participants in a conversation are focusing. By focusing attention on certain propositions, we say that epistemic modals ‘raise’ or ‘introduce’ salient possibilities into a conversation. The notion of a salient possibility allows us to account for the use of epistemic modals to answer questions, as in the following dialogue.

A: Where’s Joe?

B: He might be in Boston.

Despite the fact that B does not provide any information about Joe’s whereabouts, her answer is still felicitous. This is anomalous on the many popular theories of questions that require answers to questions to provide information. The theory presented in this thesis does not make this requirement. Rather, on our account, B felicitously answers A’s question by making certain factual answers to A’s question salient in their conversation, or introducing certain relevant salient possibilities.

This section is outlined as follows. First, we informally introduce the notion of a salient possibility. Then we discuss the semantics and pragmatics of epistemic modals used to answer questions, using the notion of a salient possibility in our discussion. We then conclude by defining a formal system in which we make our previous analyses precise.

3.1 What is a Salient Possibility?

In the classic picture of communication given in Stalnaker (1978), a conversation is a process of distinguishing between various possibilities, or ways the world might be. Conversational participants share a body of common information called the context set, which is modeled as a set of possible worlds: the possible worlds that could be actual given the shared information. A conversation is, in part, a process of adding information to the context set in order to distinguish the actual world that is (hopefully) in the context set from other merely possible worlds in the context set.

This is a useful and insightful perspective on conversation, but can be refined by taking into account two truisms. The first is that conversational participants do not attend to all the logical possibilities their information fails to rule out. Some are simply not noticed, and others may simply not matter, so they fall by the wayside. The second is that it is wrong to model the attention of conversational participants by thinking of a conversation as a process of distinguishing between *individual* possible worlds. Such a model suggests very precise attention to detail, whereas attention to detail is not nearly that fine-grained even in the most precise contexts. Instead, conversational participants distinguish possibilities based on whether or not they possess some relevant feature. For example, a possible world w in which it’s raining is not distinguished from a possible world w' in which it’s not raining through a process that attends to all of the details of each world; rather, individuals distinguish between the class of worlds in which it’s raining and the class of worlds in which it’s not raining, and by virtue of this distinguish between w and w' .

These two truisms are closely related. Note that if attention were paid to all the possibilities not ruled out by the information in the context set, then attention would have to be fine-grained enough to focus on individual possible worlds. Only by focusing on such fine-grained options could conversational participants be said to truly be taking into account all the possibilities their information fails

to exclude. Ignoring certain possibilities entails, among other things, ignoring possible distinctions between worlds in the context set. Conversely, note that even if attention were focused on propositions rather than possible worlds, this would not be meaningful if conversational participants focused their attention on all propositions. This is the case because given any two worlds, one can distinguish them based on the propositions they satisfy, and hence attention on all propositions provides the ability to make arbitrarily fine-grained distinctions between possible worlds.

In the picture that emerges from taking these truisms into account, a conversation is still a process of distinguishing between various ways the world could be. However, in this altered picture the conversational participants are not interested in every way the world could be, and hence are not interested in every possible distinction between worlds that can be made. What both of these truisms express is the coarse-grained and incomplete nature of the attention with which conversational participants focus on possibilities. Attention is still focused on possibilities, but these are possibilities such as ‘the possibility that chocolate makes cats sick’ or ‘the possibility that it’s going to rain tomorrow’, not individual possible worlds. These large scale possibilities are better conceived of as propositions or sets of possible worlds, e.g. the set of worlds in which chocolate makes cats sick or the set of worlds in which it is going to rain tomorrow.

Crucially, not only do conversational participants focus their attention on these large scale possibilities, but this attention can have a public, shared quality to it. Just as conversational participants can share their information and share their inquiries (through making assertions and asking questions), conversational participants can share attention to possibilities. This is a basic fact about conversation that emerges after one considers the truisms mentioned above. An example will easily bring this to light. Suppose two people are searching for a lost set of keys, and one asks the other “Did you check your car?”. This introduces a shared inquiry concerning whether or not the keys are in the car, but brings about another change as well: now both discourse participants are focused on the possibility that the keys are in the car, and each knows that the other has her attention focused in this way. Thus, it is common for interlocutors to share not only questions and information, but recognition of possibilities as well.

We can now introduce the notion of a salient possibility: the salient possibilities in a conversation are the large scale shared possibilities that the conversational participants are interested in. For example, if detectives are discussing a murder and have established x,y, and z as the prime suspects, then the possibilities that x,y, and z committed the murder are all salient possibilities in their conversation. If a couple is debating whether to go to Paris or London on vacation, the possibility that they go to London and the possibility that they go to Paris are both salient in their discussion.

Several constructions in natural language introduce salient possibilities into a conversation. Polar questions are a clear case: a polar question introduces

salient possibilities corresponding to its affirmative and negative answers into a conversation. This is intuitively obvious: after one conversational participant utters a polar question, the others focus their attention on the two possible answers. Moreover, since asking a question is a public act, discourse participants know that their interlocutors have focused their attention on the possibilities corresponding to the answers to the polar question, and hence the possibilities corresponding to these answers become salient possibilities.

Disjunctions introduce salient possibilities as well: ‘A or B’ makes the possibility that A and the possibility that B salient in a conversation. To see this, consider the following dialogue.

A: Who is going to lead class today?

B: Either John or Jane.

B responds to A by introducing two salient possibilities into their conversation: the possibility that John will lead class and the possibility that Jane will lead class. As H.P. Grice points out (from “Indicative Conditionals” (Grice (1991)):

A standard (if not the standard) employment of “or” is in the specification of possibilities (one of which is supposed by the speaker to be realized, although he does not know which one), each of which is relevant in the same way to a given topic.

The ability of disjunctions to specify, and hence call attention to, possibilities is properly captured in our system: ‘A or B’ draws the attention of the discourse participants to the possibility that A and the possibility that B.

Following Grice’s insight, many theorists have developed formal accounts of disjunction that go beyond the simple truth conditional ‘ \vee ’ (e.g. Simons (2000), Alonso-Ovalle (2006), and Groenendijk (2008)) Hence, by formally analyzing disjunction as something more than a truth function, we do not break new ground in this thesis. What’s new is the use of one notion, i.e. the notion of a salient possibility, to account for the similar non-informative effects that disjunctions, polar questions, and epistemic modals have in conversations.

Several constructions change the salient possibilities in a conversation, but hereafter we will mostly focus on the changes epistemic modals bring about in the salient possibilities. An epistemic modal $\diamond\phi$ introduces a salient possibility constructed from ϕ into a conversation.

Salient possibilities are well tailored for an analysis of epistemic modals. Unlike other kinds of assertions, an assertion of an epistemic modal does not contribute information to a conversation. Instead, its function is to (i) act as an information indicator (as explained in sections 2.2.4 and 2.2.5), and (ii) call attention to certain possibilities that the conversational participants should, for some reason, find interesting. Thus, to analyze epistemic modals one must develop a framework in which assertions can significantly change a context without

providing information. Since this thesis's framework postulates that epistemic modals affect the salient possibilities in a context rather than its information, the non-informative yet non-trivial effects of epistemic modals are properly represented.

3.2 Epistemic Modals as Answers to Questions

While the notion of a salient possibility makes intuitive sense, if it failed to account for important data it would not be a very useful theoretical notion. However, it turns out to be quite useful in developing an account of non-informative answers to questions. Epistemic modals, we will argue, answer questions by making certain informative answers salient possibilities, not by providing information. Below we will discuss the semantics and pragmatics of answering questions with epistemic modals to support this idea.

Epistemic modals seem to be perfectly felicitous as answers to wh-questions. We introduced one example earlier, repeated below:

A: Where's Joe?
B: He might be in Boston.

First, we note that B really does answer A's question here. B's response is far different from avoiding the question or opting out of the discussion. To see this, compare the above dialogue with the following:

A: Where's Joe?
B: I don't know / I don't care / Stop bothering me.

Then, to see that epistemic modals answer questions by introducing salient possibilities, note that there is no other candidate for the mechanism through which an epistemic modal could answer a question. B's assertion has two effects: it introduces a salient possibility and suggests that B does not know that Joe is not in Boston. Taking the latter of these to answer the question is extremely implausible:

A: Where's Joe?
#B: I don't know that he's not in Boston.

Thus, we conclude that epistemic modals both answer questions and do so by introducing salient possibilities.

One potential objection to this account holds that wh-questions themselves raise salient possibilities that correspond to their possible informative answers. If this is the case, then an epistemic modal could not answer a wh-question by making some of its possible informative answers salient possibilities, as these answers

would already be salient in the first place. However, it is a misunderstanding of salient possibilities to hold that wh-questions make all of their possible factual answers salient possibilities. For instance, if one conversational participant asks ‘Where are my keys?’, then the attention of the conversational participants is not focused on all the potential locations of the questioner’s keys. This is a simple, basic, intuitive fact about wh-questions. Moreover, this is why B can answer A’s question by saying that Joe might be in Boston: A doesn’t necessarily focus her attention on this factual answer to her question.

Rather than thinking of wh-questions as raising salient possibilities, it is better to think of them as defining a topic for discussion. Intuitively, a topic for discussion does not itself raise any possibilities, but rather is a constraint on what sorts of possibilities or pieces of information are relevant to the discussion. Given a topic, felicitous uses of disjunctive sentences and epistemic modals focus on certain possibilities that are relevant to that topic. Thus, since focusing attention within a topic is very different from raising a topic, epistemic modals and disjunctions act in a categorically different way than wh-questions.

Another way of seeing that wh-questions do not introduce salient possibilities is by contrasting the use of an epistemic modal after a wh-question with the use of an epistemic modal after an alternative question, a construction that definitely introduces salient possibilities.

A: Will ALF come to the meeting, or will BEA?¹
#B: Alf might come to the meeting.

A’s attention is already clearly focused on the possibility that Alf comes to the meeting, and hence B cannot answer her question by saying that he might come. The situation with wh-questions is far from analogous, suggesting that wh-questions do not introduce salient possibilities.

But given this story, we are left with the confusing fact that epistemic modals seem to also be able to somewhat felicitously answer polar questions.

A: Are John and Bill coming to the party?
B: They might.

Though the intuition is weaker here than it is with wh-questions, it is still clear that B’s response felicitously answers A’s question. But it is wrong to say that B answers A’s question by making the possibility that John and Bill come to the party salient in their discussion, since A clearly already has this possibility in mind. Rather than introducing this salient possibility, B’s response seems to somehow bolster or reinforce it. To address how B’s response accomplishes this,

¹A’s utterance should not be read with rising intonation throughout, but rather with accents on ‘Alf’ and ‘Bea’, and with falling intonation

we now turn to a discussion of the pragmatics of epistemic modals and salient possibilities.

A misconception one could have about epistemic possibility is that a speaker can felicitously assert $\diamond\phi$ whenever she does not know that ϕ is false. To see that this is a misconception, note the infelicity of B's response to A if B has no idea who Joe is at all:

A: Where's Joe?

#B: He might be in Boston.

In a sense, B's utterance is felicitous since B is not saying anything she knows to be wrong. But something has gone wrong here: B has no business saying anything about Joe's possible whereabouts because she knows nothing about Joe. This suggests that in many contexts, a speaker needs something more than ignorance to assert an epistemic possibility statement.

We propose that an assertion of $\diamond\phi$ generally gives rise to a conversational implicature that the speaker has some reason to believe that ϕ is the case.² This explains the infelicity of B's response in the dialogue directly above: B's utterance conversationally implicates that she has some evidence that Joe is in Boston, whereas she has nothing of the sort. Likewise, this explains the felicity of epistemic modals as answers to polar questions: the answerer implicates that she has some reason to believe that one of the answers is the case.

The existence of this conversational implicature also explains the difference between modal answers to alternative questions and modal answers to polar questions. Alternative questions of the form 'A, or B?' suggest that the speaker has some degree of commitment to 'either A or B'. We do not have a formal way of deriving this commitment from alternative questions, but its presence is an uncontroversial empirical fact. Hence, since the alternative question asker already has reason to believe that A is the case (since it would be odd to believe 'either A or B' without any reason to believe that A is the case), simply reiterating this fact cannot answer the question.

This analysis of modal answers to alternative questions only predicts that 'possibly A' and 'possibly B' are infelicitous answers, a prediction that is clearly borne out:

A: Will ALF come to the meeting, or will BEA?

B: They both might come/It might be that neither one comes.

²The notion of a 'reason to believe', like the notions of knowledge and evidence, is a potentially complicated epistemological notion. However, we want to avoid epistemological questions in this thesis, and instead rely on an intuitive understanding of a 'reason to believe'. We take the intuitive notion to be something like the following: A is a reason to believe that ϕ if A provides some amount of (potentially defeasible) evidence that ϕ is the case.

In such a context, A's question introduces the possibility that Alf will come to the meeting as well as the possibility that Bea will come to the meeting, and further suggests a belief that either or them will come. Thus, A does not suggest she has a reason to believe that they will both come, and seems to disregard the possibility that they neither come, and for this reason B's answer is felicitous.

Deriving this implicature from Grice's conversational maxims is straightforward. The most important maxim in the derivation is relevance. Bringing a possibility to the attention of one's interlocutors is, in many cases, only relevant if one has some reason to believe that it is actually the case. To see this, consider the infelicity of the following conversation:

Bank Robbery

Suppose A and B are part of a gang that is hiding out after committing a bank robbery, and the gang is trying to figure out where the police are looking for them.

A: The police might be looking for us in Los Angeles.

B: Why do you say that?

#A: Well, I don't know where they're looking, so for all I know it could be Los Angeles.

Here, A does not violate the maxim of quality: she does not say anything she does not believe. But given that the group of robbers is trying to figure out where the police are looking for them, A's contribution is completely irrelevant. It does nothing to advance the goals of their conversation, with the possible exception of displaying A's lack of relevant information and odd character.

Thus, epistemic modals answer both *wh*-questions and polar questions, but cannot answer alternative questions. Moreover, epistemic modals are perfectly felicitous as answers to *wh*-questions, but seem remotely less felicitous in response to polar questions. This is because epistemic modals accomplish more as a response to *wh*-questions: an assertion of an epistemic modal brings attention to some possibility, and conversationally implicates that there is some reason to believe that it is actual. As an answer to a polar question, only the latter takes place, and hence epistemic modals as answers to polar questions are somewhat less felicitous.

A final question to deal with is whether or not the relation between a question and a modal answer is semantic or pragmatic. In the case of a polar question, it is unquestionably pragmatic: the speaker answers the polar question by conversationally implicating that she has good reason to believe in one of the question's factual answers. However, with *wh*-questions the case is more complicated, since epistemic modals answer *wh*-questions without a conversational implicature. Though it is generally thought that semantic answers to question must provide information, contrary to this view we will argue below that epistemic modals can

provide non-informative semantic answers to questions, and provide an informal semantic view of questions and answers that supports our considerations.

In logical semantics, it is common to think of the logical answers to questions as propositions, as opposed to constituents (see Hamblin (1958)). With this understanding of questions and answers, a statement is commonly thought to answer a question if it asserts one of these propositions. As a more formally explicit example, consider the partition theory of questions (see Groenendijk and Stokhof (1997)). In such a theory, a question is a partition of logical space (i.e. the set of all logically possible worlds), and its answers are propositions that correspond to the elements of that partition or the unions of such elements. A statement is considered to be an answer to a question if it asserts such a proposition.

Clearly, with such a theory of questions epistemic modals do not count as answers. However, we think this is a bad result. After all, epistemic modals intuitively answer questions, and do so without requiring any pragmatic inferences. An assertion of $\Diamond\phi$ directly answers a question: one does not need to ‘find’ an answer by making any inferences concerning the speaker’s meaning rather than the semantic meaning. Though informative answers are pragmatically preferred to non-informative answers by the maxim of quantity, this does not mean that non-informative answers are somehow only ‘pragmatic answers’ or not answers at all. While we should keep in mind this pragmatic preference for information, it should not prevent us from constructing a theory that respects our intuitions concerning epistemic modals as answers to questions.

To make sense of epistemic modals as answers to questions within the logical perspective sketched above, we only need to change our notion of a linguistic answer. Rather than thinking of linguistic answers to questions as assertions of logical answers, we think of linguistic answers as statements that *do* something with a logical answer. In our framework, this can be either an assertion of a logical answer or an utterance that makes a logical answer salient in the conversation. We can still view the answers to a question as a set of propositions; we only need to change our view concerning what a statement must do with these propositions in order to count as a linguistic answer.

Thus, accounting for non-informative semantic answers only requires us to expand our definition of linguistic answers to include non-informative uses of propositions. Our formal theory of answers below is based on this understanding of answerhood.

Our discussion of epistemic modals as answers to questions showed that this topic has important semantic and pragmatic dimensions. Moreover, the notion of a salient possibility was crucial to our account, showing that it is both empirically useful and intuitively compelling. We now turn to a formal theory in which we make our considerations about salient possibilities and epistemic modals as answers to questions precise.

3.3 Formal Theory of Salient Possibilities

In this section, we define a dynamic semantics for a first-order language with an epistemic modality operator. This language can also express questions, but does not need a basic question operator because our semantics is *inquisitive*. In defining our semantics, we need two different notions: the notion of an update of an *information state*, and an update of a *context*; the latter is defined in terms of the former. With our semantics in place, we define two notions: the notion of an informative answer and the notion of a non-informative answer, and show that epistemic modals provide non-informative answers to questions.

3.3.1 Language and Syntax

Our language is a first order language with an epistemic possibility operator.

Definition 3.1. (Formal Language)

- The logical vocabulary of L contains the Boolean connectives \neg , \vee , and \wedge ; the quantifier \forall ; variables x, y, z etc.; and the epistemic possibility operator \diamond .
- The non-logical vocabulary of L contains constants a, b, c , etc. and n -place predicates P .
- A term of L is a variable or a constant.

Definition 3.2. (Syntax)

The set of formulas of L is the smallest set such that:

- (i) If P is an n -place predicate and t_1, \dots, t_n are terms, then $P(t_1, \dots, t_n) \in L$
- (ii) If $\phi \in L$, then so is $\neg\phi$
- (iii) If $\phi, \psi \in L$, then so are $\phi \vee \psi$ and $\phi \wedge \psi$
- (iv) If $\phi \in L$ and x is a variable, then $\forall x\phi \in L$.
- (v) If $\phi \in L$, then so is $\diamond\phi$.

In an inquisitive semantics, taking the existential quantifier as primitive makes quantification inquisitive, i.e. statements of the form $\exists x\phi$ can raise issues. Though this sort of quantification is interesting, implementing it in our system would complicate our exposition and not address the topics we are considering in this thesis. Hence, we take the universal quantifier as basic, and in this system the derived existential quantifier (i.e. $\exists x\phi \equiv \neg\forall x\neg\phi$) is not inquisitive.

Moreover, taking the existential quantifier as basic introduces a novel semantics for wh-questions, which is another reason we take the universal quantifier

as primitive. For purposes of this thesis the classic partition approach to wh-questions suffices, and hence we adopt it for ease of exposition.

The material conditional $\phi \rightarrow \psi$ can be defined as $\neg(\phi \wedge \neg\psi)$. With this definition of the conditional, we can express the conditional question $\phi \rightarrow ?\psi$ as $(\phi \rightarrow \psi) \vee (\phi \rightarrow \neg\psi)$. We define the polar question operator $?$ in the following way: $? \phi \equiv \phi \vee \neg\phi$. In our presentation of the semantics, we explain why this is an appropriate definition for the polar question operator.

3.3.2 Inquisitive Semantics for Information States

In section 3.3.4, we will define a dynamic semantics where formulas can introduce questions and have other informative and non-informative effects on conversations. The semantics we introduce in this section provides the foundations for the semantics in 3.3.4. Here, we define the notion of an *information state* and the manner in which a formula acts on an information state, and we will later define how a formula acts on a *conversation* (or a *context* as we will say later) in terms of how it acts on certain information states. Moreover, our semantics for questions is essentially defined in this section as well; the semantics for questions in 3.3.4 is a trivial generalization of the semantics we define here.

As we mentioned above, the semantics we define here is inquisitive (see Groenendijk (2008)). In an inquisitive semantics, the content of a formula is given by not only the data it contains, but also the *issues* it raises. The data contained in a formula is the information about the world that a formula expresses, whereas the issues raised in a formula are the questions about the world that a formula asks. Some formulas, like atomic formulas, only contain data: Pa only contains the information that a has the property P .¹ Other formulas, like questions, only raise issues: the polar question $?Pa$ only contains a question asking whether or not a has the property P .

So far, this is nothing new: there are many formal systems that can express questions and assertions, where questions raise issues and assertions contain data (e.g. Groenendijk (1999)). What is new in inquisitive semantics is that there is no categorical distinction between issue-raising formulas and data-containing formulas. There are also semantic hybrids that contain data but also raise issues. A simple example of a hybrid is the disjunction $Pa \vee Pb$. This formula provides the data that either a has the property P or b has the property P . But it also raises the issue ‘which one is it?’, i.e. the issue of whether it is a that has the property P or b that has the property P .

This account of disjunction allows us to define polar questions in terms of disjunctions. To reiterate, a polar question $?Pa$ raises the issue of whether or not Pa is the case. Given our semantics, this polar question is equivalent to the disjunction $Pa \vee \neg Pa$: such a disjunction does not provide any data, but it raises

¹This is a bit of a stipulation. It could also be that an atomic formula Pa raises the issue ‘Who is a ?’

the issue of ‘which one?’ for Pa and $\neg Pa$. This latter issue is the same issue that is raised by a polar question, and hence in inquisitive semantics polar questions can be defined in terms of disjunctions.

In section 2, we briefly motivated the view that data is fundamentally something that eliminates possible worlds, those worlds incompatible with the data. Similarly, an issue is fundamentally something that distinguishes between worlds. For example, a polar question distinguishes between worlds that correspond to a ‘yes’ answer to the question from worlds that correspond to a ‘no’ answer to the question. The wh-question ‘Who is coming to the party?’ distinguishes between worlds in which different guests attend the party. A disjunction ‘A or B’ distinguishes the worlds in which A is the case from the worlds in which B is the case. Thus, we can represent both data and issues using a possible worlds semantics.

Since issues are non-informative, inquisitive semantics, like the semantics we define in section 3.3.4 below, does not identify the semantic content of a formula with its informative content. Thus, though the formal semantics developed below is original insofar as it posits that the content of epistemic modals has a non-informative dimension, the idea that formulas have content beyond their informative content is not original. However, the notion of non-informative content developed below, defined in terms of salient possibilities, is new to this thesis.

There are both static and dynamic versions of inquisitive semantics; in this thesis we present a dynamic version. Recall that in a dynamic semantics, the meaning of a formula is its potential to change an agent’s cognitive state. Earlier we (simplistically) modeled an agent’s cognitive state as a set of information about the world, a data set. Since our language is now able to provide information about the world and raise issues, we need to refine our model. We now model an agent’s cognitive state as an *information state*, a formal object that can represent both data and issues.

An information state is a set of pairs of indices $\langle i, j \rangle$, where an index is essentially a possible world (here, a first order model) coupled with an assignment function. We require information states I to be reflexive, and then take the set $\{i : \langle i, i \rangle \in I\}$ to represent the data in I . We think of a set of indices as representing data in the same way that a set of possible worlds represents data. Though we do not need pairs of indices to represent data, we use pairs in a crucial way to represent issues. Call two indices i and j *connected* in an information state I if $\langle i, j \rangle \in I$, and call them *disconnected* in I if $\langle i, i \rangle, \langle j, j \rangle \in I$ but $\langle i, j \rangle \notin I$. The set of connected and disconnected indices in an information state determines the issues that are represented in it. If two indices i and j are connected in I , then there is no issue in I that distinguishes between i and j . However, if i and j are disconnected in I , then there is some question in I that distinguishes between them. For example, suppose I only contains the question $?Pa$: ‘does a have the property P ?’. Then any two indices in I that agree on the truth value of Pa will be connected in I , but indices that disagree on the truth value of Pa will be disconnected in I . Since issues that distinguish i from j also, trivially, distinguish

j from i , we require the connectedness relation to be symmetric.

We formally define indices and information states below. Note that in the definition of indices, we follow Kripke (1980) and treat constants as rigid designators.

Definition 3.3. (Indices)

Fix a non-empty set D , the domain. An index i for the language L is a function such that:

- (i) For all n -place predicates P of L , $i(P) \subseteq D^n$
- (ii) For all terms t of L , $i(t) \in D$.
- (iii) For all constants c of L and all indices i and j , we require that $i(c) = j(c)$.

Definition 3.4. (Information States)

Let T be the set of all indices for the language L . An information state I is a reflexive and symmetric relation on a subset of T .

A key notion in inquisitive semantics is that of a *possibility* in an information state. A possibility in inquisitive semantics is different from a salient possibility. In inquisitive semantics, a possibility is a potential resolution of an issue that is represented in an information state, not a proposition on which conversational participants are focusing. For example, if an information state only contains the issue of whether or not a has the property P , then there are two resolutions of this issue in the information state: one in which a is P , and another in which a is not P . In such a case, the set of indices that satisfy Pa is a possibility, and the set that do not satisfy Pa is a possibility.

Furthermore, possibilities are ‘maximal’ resolutions of issues in an information state: any single possibility corresponds to a resolution of all the issues in an information state. If an information state only contains the issue raised by ‘who is coming to the party’, then each possible list of party guests and non-party guests corresponds to a different possibility in that information state. Partial resolutions, however, can be recovered by considering the union of possibilities: the possibility that a is coming to the party is the union of all possibilities that include a among the party guests.

Definition 3.5. (Possibilities)

Let I be an information state. A set of indices \mathbf{P} is a possibility in I if:

- (i) For all $i, j \in \mathbf{P}$, $\langle i, j \rangle \in I$.
- (ii) There is no set of indices \mathbf{P}' such that $\mathbf{P} \subset \mathbf{P}'$ and \mathbf{P}' satisfies (i).

By clause (i) of this definition, the indices in a possibility are not distinguishable by the issues in an information state. Thus, the data represented by the possibility is sufficient for resolving the issues in an information state, which is why we think of possibilities as resolutions of issues in an information state. Clause (ii) ensures that possibilities do not contain any data irrelevant to resolving the issues in an information state. Any proper subset of a possibility contains data that eliminates indices that the issues in an information state do not distinguish between, and hence contains data that is not necessary to resolving the issues in an information state.

Possibilities for Polar Questions

Consider the following information state:

$$\{\langle i, j \rangle : \text{either } i(a) \in i(P) \text{ and } j(a) \in j(P) \text{ or } i(a) \notin i(P) \text{ and } j(a) \notin j(P)\}.$$

The possibilities in this information state are:

$$\{i : i(a) \in i(P)\} \text{ and } \{i : i(a) \notin i(P)\}$$

Possibilities for Wh-Questions

Consider the following information state:

$$\{\langle i, j \rangle : i(P) = j(P)\}$$

The possibilities in this information state are all sets of indices of the form

$$\{i : i(P) = D'\} \text{ where } D' \subseteq D.$$

In our update semantics, formulas perform two sorts of actions: either they disconnect pairs of indices $\langle i, j \rangle$ by excluding such pairs, or they exclude reflexive pairs of indices $\langle i, i \rangle$ and any non-reflexive pair containing i . The first sort of update corresponds to raising an issue: disconnecting indices corresponds to raising an issue that distinguishes between the disconnected indices. The second sort of update corresponds to contributing data: excluding a reflexive pair of indices corresponds to adding data that is inconsistent with the way that index represents the world to be.

We present one more standard notion relativized to our system, and then state our update semantics for information states.

Definition 3.6. (Changing Assignments)

Let $\langle i, j \rangle$ be a pair of indices, let I be an information state, and let $d \in D$.

- (i) $i_{[x/d]}$ is the index such that for terms $t \neq x$, $i_{[x/d]}(t) = i(t)$, and $i_{[x/d]}(x) = d$.
- (ii) $\langle i, j \rangle_{[x/d]} = \langle i_{[x/d]}, j_{[x/d]} \rangle$
- (iii) $I_{[x/d]} = \{\langle i, j \rangle \in I : \langle i, j \rangle_{[x/d]} = \langle i, j \rangle\}$

Definition 3.7. (Update Semantics for Information States)

Let I be an information state, and let $\phi, \psi \in L$. We define the update of I with a formula of L as follows:

- (i) $I[P(t_1, \dots, t_n)] = \{\langle i, j \rangle : \langle i(t_1), \dots, i(t_n) \rangle \in i(P) \text{ and } \langle j(t_1), \dots, j(t_n) \rangle \in j(P)\}$
- (ii) $I[\neg\phi] = \{\langle i, j \rangle \in I : \langle i, i \rangle \notin I[\phi] \text{ and } \langle j, j \rangle \notin I[\phi]\}$
- (iii) $I[\phi \vee \psi] = I[\phi] \cup I[\psi]$
- (iv) $I[\phi \wedge \psi] = I[\phi][\psi]$
- (v) $I[\diamond\phi] = I$ if $I[\phi] \neq \emptyset$
- (vi) $I[\diamond\phi] = \emptyset$ if $I[\phi] = \emptyset$
- (vii) $I[\forall x\phi] = \{\langle i, j \rangle \in I : \text{for all } d \in D, \langle i, j \rangle_{[x/d]} \in I_{[x/d]}[\phi]\}$

We work through some examples to illustrate this semantics. In these examples, we write \mathbf{T} for the set of all pairs of indices $T \times T$.

Simple Disjunction

We represent the simple disjunction “Alf is coming to the party or Bea is coming to the party” as $Pa \vee Pb$.

$$\begin{aligned} \mathbf{T}[Pa \vee Pb] &= \mathbf{T}[Pa] \cup \mathbf{T}[Pb] \\ &= \{\langle i, j \rangle : i(a) \in i(P), j(a) \in j(P)\} \cup \{\langle i, j \rangle : i(b) \in i(P), j(b) \in j(P)\} \\ &= \{\langle i, j \rangle : \text{either } i(a), j(a) \in i(P) \text{ or } i(b), j(b) \in i(P)\}. \end{aligned}$$

Thus, this simple disjunction provides the information that either Pa or Pb is the case, and raises the issue ‘which one?’.

Polar Question

We represent the polar question “Is Alf coming to the party?” as $Pa \vee \neg Pa$.

$$\begin{aligned} \mathbf{T}[Pa \vee \neg Pa] &= \mathbf{T}[Pa] \cup \mathbf{T}[\neg Pa] \\ &= \{\langle i, j \rangle : i(a) \in i(P), j(a) \in j(P)\} \cup \{\langle i, j \rangle : i(a) \notin i(P), j(a) \notin j(P)\} \\ &= \{\langle i, j \rangle : \text{either } i(a), j(a) \in j(P) \text{ or } i(a) \notin i(P) \text{ and } j(a) \notin j(P)\}. \end{aligned}$$

As we saw above, the possibilities in this information state represent the two resolutions of the polar question, one in which Alf comes and one in which he does not.

Wh-Question

We represent the wh-question “Who is coming to the party?” as $\forall x?Px$.

$$\mathbf{T}[\forall x?Px] = \{\langle i, j \rangle : i(P) = j(P)\}.$$

As we saw above, the possibilities in this information state correspond to all possible denotations for the predicate P . Thus, there is a distinct possibility in this information state for each list of potential party guests.

Relatedness in inquisitive semantics is a notion that governs the resolution of issues. If I is related to I' , then I' contains only issues contained in I , and potentially resolves some of the issues raised in I . Intuitively, if a formula ϕ is an answer to any of the issues raised in an information state I , then I is related to $I[\phi]$. The formal definition runs as follows:

Definition 3.8. (Relatedness)

Let I and I' be information states. Then I is related to I' , $I \propto I'$, if every possibility in I' is the union of some possibilities in I .

It is easy to check that in general $I[?\phi]$ is related to $I[?\phi][\phi]$, and $I[\phi \vee \psi]$ is related to $I[\phi]$ and $I[\psi]$.

3.3.3 Salient Possibilities

We argued above that a conversation takes place not only against a background of shared information (a common ground), but against a background of salient possibilities as well. In this section, we introduce a formal structure that represents conversational backgrounds as containing both information and salient possibilities; we call these structures *contexts*. Since formulas in the update semantics below map contexts to contexts, these formal structures are a crucial part of our theory.

Recall that the salient possibilities in a conversation are the propositions in which the participants in the conversation are interested. Below, we represent salient possibilities as subsets of the common ground. The intuition behind this representation is that for each salient possibility, there is a subset of the common ground that would result from the discourse participants accepting this salient possibility as actual, and we think of this subset as representing the salient possibility. We choose this representation of salient possibilities because when conversational participants consider these propositions, they do not consider them in a vacuum. Rather, they consider the propositions relative to the information they already possess and the issues they are already considering. Thus, we represent salient possibilities as certain continuations of the common ground, not sets of possible worlds that are insensitive to the current state of a conversation's common ground.

To see why this is appropriate, recall the example of detectives solving a murder. Suppose they know that the suspect murdered the victim in the library with a candlestick, and suppose the salient possibilities are the possibility that Green is the murderer, that Jones is the murderer, and that Bill is the murderer. We represent such a situation as follows. The common ground of the detectives contains their facts about the murder location and the weapon. The salient

possibilities in their context are then best thought of as continuations of this common ground: one continuation where Bill commit the murder in the library with the candlestick, one continuation where Jones commit the murder in the library with the candlestick, and a final continuation where Green commit the murder in the library with a candlestick. Intuitively, this represents the manner in which the detectives consider the salient possibilities relative to their current information about the murder.

Moreover, the detectives consider salient possibilities in light of the issues they are currently considering. Suppose the detectives are considering three issues: who commit the murder, which weapon was used, and where the murder took place. When considering salient possibilities, the detectives do not only consider the information represented by the possibilities, but also the manner in which this information is relevant to their current issues. For example, suppose the detectives are considering the possibility that Bill commit the murder, and know that if Bill commit the murder then he used a knife. Intuitively, the detectives regard this possibility as one in which two of their issues are resolved, but one remains: the issue of where the murder was committed. This suggests that when discourse participants consider salient possibilities, they consider possibilities in light of not only their currently shared information, but their currently shared issues as well.

Formally, a context is a set of information states with one maximal information state. The maximal information state represents the common ground in the context, and the other information states are the salient continuations of the common ground.

Definition 3.9. (Contexts)

Let T be the set of all indices for the language L . We define a context to be a set $\Gamma = \{P_1, \dots, P_n, \dots\}$ such that:

- (i) $P_i \subseteq T \times T$ for all n
- (ii) For some i , $P_i = \emptyset$
- (iii) There is an i such that for all j , $P_j \subseteq P_i$. This maximal set P_i is called the common ground.

We write CG (common ground) for the maximal set P_i defined in (iii), and write $\Gamma = \{CG, P_1, \dots, P_n, \dots, \emptyset\}$. We also write Γ_{CG} to refer to the common ground of Γ

We now formalize a few simple conversations to demonstrate this formalism.

Detectives

Recall the example of the detectives solving the murder above. Let the formula Mx represent “x commit the murder”, and let j denote Jones, b

denote Bill, and g denote Green. We then represent the context in their conversation as follows:

$$\Gamma_{Detectives} = \{CG, CG[Mj], CG[Mb], CG[Mg], \emptyset\}$$

Here the common ground represents their shared information, and the other possibilities represent the extensions of that information they find interesting.

Beach or Grandma’s House

Suppose a parent tells their child that today they will either go to the beach or to the child’s grandmother’s house. Let ϕ represent “the parent and the child go to the beach” and let ψ represent “the parent and child go the grandmother’s house”. Then:

$$\Gamma_{Beach/Grandma} = \{CG, CG[\phi], CG[\psi], \emptyset\}$$

Considering a Question

Suppose one participant asks another if it is going to rain tomorrow. Let ϕ represent “it is going to rain tomorrow”. Then:

$$\Gamma_{Rain} = \{CG, CG[\phi], CG[\neg\phi], \emptyset\}$$

Here, CG contains the issue of whether or not it is going to rain tomorrow, and the two salient possibilities correspond to resolutions of this issue.

A final question to address is whether or not we should impose any formal constraints on contexts to distinguish between contexts that corresponded to ‘reasonable’ sets of salient possibilities and contexts that correspond to ‘unreasonable’ sets of salient possibilities. After all, there are certainly infelicitous ways of introducing salient possibilities, for example in the dialogue below:

A: John might be in Paris.

B: Right. In fact, John might be in France.

One might think that these exchanges are infelicitous because they lead to an improper set of salient possibilities, and that we can express a generalization concerning such infelicitous sequences by demarcating a class of ‘reasonable’ or ‘felicitous’ contexts through some formal constraint. However, we are skeptical that such a constraint can be given. Our skepticism derives from the fact that initially plausible candidates for these formal constraints fail. For example, one might try a formal constraint that barred salient possibilities that correspond to formulas of the form $\phi \vee \psi$ if a possibility corresponding to ϕ is already salient in the context. Such a constraint would rule out the manner of introducing possibilities executed in the dialogue above. However, sometimes this manner of considering possibilities is perfectly felicitous:

A: John might be in Paris.

B: Right. In fact, though it's especially likely that he's in Paris, he really might be anywhere in France.

Unlike our first example, this dialogue is totally felicitous. However, the same possibilities were raised, and this suggests to us that formal constraints on salient possibilities are not desirable.

On a more principled level, we find such constraints to be undesirable because of the diverse reasons for which possibilities can be raised. Possibilities can be raised regardless of whether they are likely or unlikely, feared or hoped for, or silly or important. Very complex pragmatic factors concerning when possibilities are relevant to consider interfere with any such overarching formal constraints we would like to place on them. Thus, though there are certainly felicitous and infelicitous ways of introducing possibilities, this does not necessarily translate into formal constraints on contexts that can express why these methods of introducing possibilities are infelicitous.

3.3.4 Semantics for Contexts

In this section, we will use our semantics for information states to define a dynamic semantics in which formulas map contexts to contexts. We then define a notion of an informative answer and a non-informative answer, and show that with these definitions epistemic modals provide non-informative answers to wh-questions. Finally, we define entailment for this semantics, show why universal instantiation fails with this definition of entailment, and then argue that this is a desired consequence of the system.

Earlier, we said that under a dynamic conception of meaning, the meaning of a sentence is its potential to change an agent's cognitive state. Now, however, we are defining a dynamic semantics in which contexts are mapped to contexts. This may seem odd because contexts are not cognitive states of agents. To resolve this difficulty, think of a formula as changing an agent's cognitive state through its effect on a context. After all, the notions we used to construct a context, i.e. the notion of a common ground and the notion of a salient possibility, both make crucial reference to the cognitive states of discourse participants. We represent formulas as changing contexts, but this is really a way of expressing the complex effect a formula has on the cognitive states of discourse participants in conversations.

Simple atomic formulas and conjunctions of such formulas act in a fairly mundane way in our semantics: these formulas add information to the common ground of a context. Quantified formulas do the same. Disjunctions and epistemic modals, however, introduce salient possibilities. These formulas act on the common ground of a context, but also add updates of the common ground to a context, and hence create new salient possibilities. The sorts of salient possibil-

ities added by these constructions are the natural choices: epistemic modals $\diamond\phi$ add $\text{CG}[\phi]$ as a salient possibility, and $\phi \vee \psi$ adds two salient possibilities, $\text{CG}[\phi]$ and $\text{CG}[\psi]$.

Definition 3.10. (Update Semantics for Contexts)

Let $\Gamma = \{\text{CG}, P_1, \dots, P_n, \dots, \emptyset\}$ be a context, and let $\phi, \psi \in L$. We define the update of Γ with a formula of L , written $\Gamma[[\phi]]$ where $\phi \in L$, as follows.

- (i) $\Gamma[[P(t_1, \dots, t_n)]] = \{\text{CG}[P(t_1, \dots, t_n)], P_1[P(t_1, \dots, t_n)], \dots, P_n[P(t_1, \dots, t_n)], \dots, \emptyset\}$
- (ii) $\Gamma[[\neg\phi]] = \{\text{CG}[\neg\phi], P_1[\neg\phi], \dots, P_n[\neg\phi], \dots, \emptyset\}$
- (iii) $\Gamma[[\phi \vee \psi]] = \{\text{CG}[\psi \vee \phi], P_1[\psi \vee \phi], \dots, P_n[\psi \vee \phi], \dots, \emptyset\} \cup \{\text{CG}, \emptyset\}[[\phi]] \cup \{\text{CG}, \emptyset\}[[\psi]]$
- (iv) $\Gamma[[\phi \wedge \psi]] = \Gamma[[\phi]][[\psi]]$
- (v) $\Gamma[[\diamond\phi]] = \{\text{CG}[\diamond\phi], \text{CG}[\phi], P_1, \dots, P_n, \dots, \emptyset\}$
- (vi) $\Gamma[[\forall x\phi]] = \{\text{CG}[\forall x\phi], P_1[\forall x\phi], \dots, P_n[\forall x\phi], \dots, \emptyset\}$

We now present some examples. We use the following notational convention: where $\phi \in L$, we let $\{\phi\} = T \times T[\phi]$. Recall that T is the set of all indices for L .

Polar Question

Let $\Gamma = \{T \times T, \emptyset\}$. Then:

$$\Gamma[[Pa \vee \neg Pa]] = \{T \times T[Pa \vee \neg Pa], \{Pa\}, \{\neg Pa\}, \emptyset\}$$

Thus, the indices in the common ground that disagree on Pa are disconnected, and the possibility that Pa is the case and the possibility that it is not the case are both salient.

Disjunction

Let $\Gamma = \{T \times T, \emptyset\}$. Then:

$$\begin{aligned} \Gamma[[Pa \vee Pb]] &= \{T \times T[Pa \vee Pb], \emptyset\} \cup \{T \times T, \emptyset\}[Pa] \cup \{T \times T, \emptyset\}[Pb] \\ &= \{T \times T[Pa \vee Pb], \{Pa\}, \{Pb\}, \emptyset\} \end{aligned}$$

Disjunction of Epistemic Modals

Let $\Gamma = \{T \times T, \emptyset\}$. Then:

$$\begin{aligned} \Gamma[[\diamond Pa \vee \diamond Pb]] &= \{T \times T, \emptyset\} \cup \{\{Pa\}, \emptyset\} \cup \{\{Pb\}, \emptyset\} \\ &= \{T \times T, \{Pa\}, \{Pb\}, \emptyset\}. \end{aligned}$$

So $\Gamma[[\diamond Pa \vee \diamond Pb]] = \{T \times T, \{Pa\}, \{Pb\}, \emptyset\}$, and hence a disjunction of epistemic modals raises two possibilities but does not provide any data.

An interesting observation made in Gillies and von Fintel (2007) is that disjunctions of epistemic modals seem to entail each of their disjuncts: in many contexts, a speaker who utters (4) seems committed to both (5) and (6):

- (4) Joe might be in Boston or he might be in New York.
- (5) Joe might be in Boston.
- (6) Joe might be in New York.

Since we find ‘ \vee introduction’ to be an extremely plausible logical principle, we do not think that the relationship between (4) and (5)/(6) is strictly speaking one of entailment. Instead, we hold that in contexts in which neither $\neg\phi$ nor $\neg\psi$ has been established, $\diamond\phi \wedge \diamond\psi$ is equivalent to $\diamond\phi \vee \diamond\psi$, explaining the perceived entailment. After all, disjunctions of epistemic modals are only felicitous in such contexts, and hence the entailment holds relative to felicitous uses of disjunctions of epistemic modals. The following example shows that the claimed equivalence holds in simple cases, and highlights why the equivalence holds more generally.

Conjunction of Epistemic Modals

Let $\Gamma = \{T \times T, \emptyset\}$. Then:

$$\begin{aligned} \Gamma[[\diamond Pa \wedge \diamond Pb]] &= \Gamma[\diamond Pa][\diamond Pb] \\ &= \{T \times T, \{Pa\}, \emptyset\}[Pb] = \{T \times T, \{Pa\}, \{Pb\}, \emptyset\}. \end{aligned}$$

Thus, $\diamond Pa \wedge \diamond Pb$ and $\diamond Pa \vee \diamond Pb$ act identically on $\{T \times T, \emptyset\}$

Though disjunctions and epistemic modals raise new salient possibilities in contexts, other constructions only have trivial effects on salient possibilities. While these other formulas update salient possibilities as well as the common ground, they only do so to ensure that salient possibilities remain continuations of the common ground. The following fact makes this clear.

Fact 3.11. (Purely Informative Updates)

Let $\Gamma = \{CG, P_1, \dots, P_n, \dots, \emptyset\}$ be a context, and let ϕ be a formula of L that does not contain ‘ \vee ’ or ‘ \diamond ’. Then: $\Gamma[[\phi]] = \{CG[\phi], P_1 \cap CG[\phi], \dots, P_n \cap CG[\phi], \dots, \emptyset\}$.

Fact 3.11 can easily be proven by induction. According to this fact, updating a context with a purely informative formula amounts to updating the common ground with it and then redefining the salient possibilities to be subsets of the new common ground. Thus, the non-informative level of meaning that can non-trivially alter the salient possibilities does not creep into the purely informative fragment of the language.

A formula ϕ entails ψ with respect to this semantics if after an update with ϕ , and update with ψ is null. So, if ϕ entails ψ , then the update potential of ψ is encompassed within the update potential of ϕ .

Definition 3.12. (Entailment)

Let $\phi, \psi \in L$. Then $\phi \models \psi$ if for all contexts Γ , $\Gamma[[\phi]][[\psi]] = \Gamma[[\phi]]$.

Some may find it a bit odd to call this relationship ‘entailment’. For some, entailment should be fundamentally about information (or perhaps issues if one is dealing with questions), whereas this notion takes into account the ability of formulas to introduce salient possibilities. Such readers could think of the notion defined above, then, as not entailment but ‘redundancy’: given our definition, if ϕ entails ψ then ψ does not contribute anything beyond what ϕ , so ψ is redundant after ϕ . Nothing in the exposition hangs on this distinction, though perhaps the failure of universal instantiation below is less surprising if one opts to interpret the notion defined in definition 3.12 as redundancy.

Our entailment relation shares many of the properties of entailment in Veltman (1996): it is non-monotonic and non-reflexive. A new feature of this entailment relation is the failure of universal instantiation.

Universal Instantiation

To show that universal instantiation fails, consider $\Gamma = \{T \times T, \emptyset\}$. Since anything is possible in Γ , clearly $\Gamma[[\forall x \diamond Px]] = \Gamma$. However:

$\Gamma[[\diamond Pa]] = \{T \times T, \{Pa\}, \emptyset\}$, and hence $\Gamma[[\forall x \diamond Px]][[\diamond Pa]] \neq \Gamma[[\forall x \diamond Px]]$.

Given the intuitions behind our semantics, however, this is to be expected. Universally quantified statements do not introduce salient possibilities, but simple epistemic modals do, and hence we should not expect the former statement to entail the latter. Moreover, this prediction seems to be borne out in natural language.

- (7) The keys might be anywhere.
- (8) The keys might be in the basement.

Whereas the information contained in (8) encompasses the information contained in (7), an utterance of (7) does not introduce the salient possibility that the keys are in the basement. But since (8) does introduce this possibility, the update potential of (8) is not contained in the update potential of (7).

As discussed above, with the theory of salient possibilities one can define informative and non-informative answers to questions. Informative answers change the common ground, whereas non-informative answers add salient possibilities. While informative answers are certainly pragmatically preferred to non-informative answers it is crucial to note that non-informative answers are still a kind of answer. We define the two relations as follows.

Definition 3.13. (Informative Answers)

Let $\phi, \psi \in L$, and let Γ be a context. We say that ψ is an informative answer to ϕ in Γ if $\Gamma_{CG}[\phi] \propto \Gamma_{CG}[\phi][\psi]$

Definition 3.14. (Non-Informative Answers)

Let $\phi, \psi \in L$, and let Γ be a context. We say that ψ is a non-informative answer to ϕ in Γ if there is a $P \in \Gamma[\phi][\psi]$ such that:

- (i) $\Gamma_{CG}[\phi] \propto P$
- (ii) $P \notin \Gamma[\phi]$

We work through the following examples to demonstrate.

Wh-question Answered by an Atomic Sentence

Let $\forall x?Px$ represent the question “Who is coming to the party?”, and let Pa represent “Alf is coming to the party”. Then in the context $\Gamma = \{T \times T, \emptyset\}$, Pa informatively answers $\forall x?Px$:

$$\begin{aligned} \Gamma[[\forall x?Px]][[Pa]] &= \{\{\langle i, j \rangle : i(P) = j(P)\}, \emptyset\}[[Pa]] \\ &= \{\{\langle i, j \rangle : i(P) = j(P) \text{ and } i(a) \in i(P)\}, \emptyset\}. \end{aligned}$$

Thus, every possibility in $\Gamma_{CG}[\forall x?Px][Pa]$ is a possibility in $\Gamma_{CG}[\forall x?Px]$, so $\Gamma_{CG}[\forall x?Px] \propto \Gamma_{CG}[\forall x?Px][Pa]$.

Wh-question Answered by a Modal Sentence

Again, Let $\forall x?Px$ represent the question “Who is coming to the party?”, and let Pa represent “Alf is coming to the party”. Then in the context $\Gamma = \{T \times T, \emptyset\}$, $\diamond Pa$ non-informatively answers $\forall x?Px$:

$$\begin{aligned} \Gamma[[\forall x?Px]][[\diamond Pa]] &= \{\{\langle i, j \rangle : i(P) = j(P)\}, \emptyset\}[[\diamond Pa]] \\ &= \{\{\langle i, j \rangle : i(P) = j(P)\}, \{\langle i, j \rangle : i(P) = j(P) \text{ and } i(a) \in i(P)\}, \emptyset\}. \end{aligned}$$

Thus, every possibility in $\{\langle i, j \rangle : i(P) = j(P) \text{ and } i(a) \in i(P)\}$ is a possibility in $\Gamma_{CG}[\forall x?Px]$, so:

$$\Gamma_{CG}[\forall x?Px] \propto \{\langle i, j \rangle : i(P) = j(P) \text{ and } i(a) \in i(P)\}. \text{ Moreover:}$$

$\{\langle i, j \rangle : i(P) = j(P) \text{ and } i(a) \in i(P)\} \notin \Gamma[\forall x?Px]$, and therefore $\diamond Pa$ non-informatively answers $\forall x?Px$ in Γ .

3.4 Conclusion

In this section of the paper, we introduced a new feature of conversations: salient possibilities. A conversation takes place not only against a background of shared information, but against a background of shared salient possibilities as well. Their existence follows from certain truisms about the manner in which conversational participants focus on possibilities in a conversation. The attention of conversational participants is coarse grained, focused on propositions instead of possible worlds. This is because the conversational participants do not care about every distinction between worlds that can be made, and also because they simply do not think about every distinction between worlds that can be made.

Salient possibilities allowed us to define a non-informative dimension of the meaning of epistemic modals and disjunctions. These constructions do not only provide information in a context, but focus attention on possibilities as well, and hence there is a non-informative dimension to their content. Though many have shared our intuitions about this non-informative dimension of the meaning of epistemic modals, we used these intuitions to understand data concerning modal answers to question and construct a formal theory about the non-informative content of epistemic modals. Frameworks that do not admit something like salient possibilities into their logical ontology will have difficulty accounting for this dimension of meaning.

4 Cognitive Aspects of Non-Informative Content

Salient possibilities are certain propositions on which the participants in a conversation are focusing. Thus, there is a very trivial sense in which salient possibilities are related to some kind of cognition: paying attention to a possibility in a conversation is a cognitive act. Below, we will argue that the cognitive role of salient possibilities is not exhausted by their ability to call attention to certain possibilities in a conversation. Instead, we will argue that salient possibilities are a component of belief states. Rather than modeling an agent's belief state as a body of information (see e.g. Stalnaker (1984)), we model an agent's beliefs as a body of information along with some salient possibilities.

To get a sense of what is at stake in this discussion, consider the alternative: that belief states are not composed of salient possibilities. A proponent of such a view might hold that the only role of salient possibilities is to direct attention in a conversation. Salient possibilities, for such a theorist, are not components of thought, but simply tools that aid in the management of conversation. One can imagine a similar view involving discourse referents (see Karttunen (1971), Kamp (1981)). One might think that the role of discourse referents is to keep track of certain entities in a dialogue, but that discourse referents never make up mental content.

To argue against such a view, we will consider two kinds of data. The first concerns epistemic modals embedded under attitude verbs, and the second concerns the use of epistemic modals to give warnings or pieces of advice. We will argue that belief ascriptions of the form 'X believes that possibly ϕ ' are only true if X has consciously attended to the possibility that ϕ , a fact which suggests that belief states contain salient possibilities. Further, we will argue that epistemic modals used to warn or give advice must effect a change in beliefs, and that this suggests that belief states are composed of salient possibilities. We follow our informal discussion with a brief discussion of another framework (presented in Yalcin (2008)) that addresses our data concerning belief reports, and then conclude with a formal theory of belief based on the framework developed in section 3.

4.1 Why Belief States Contain Salient Possibilities

To show that salient possibilities are parts of belief states, we will show how a toy theory of attitude reports that does not utilize salient possibilities fails to account for data about epistemic modals under attitude reports. Our toy theory of attitude reports is loosely based on the theory of belief in Hintikka (1962), or alternatively can be seen as a Kripke semantics in which we think of 'believes possibly' as equivalent to the dual of the necessity modal 'believes' (see Kripke

(1963)). In our toy theory of attitude reports, we model an individual A's beliefs as a set of possible worlds, B_A . These worlds represent the set of worlds that are consistent with the content of A's beliefs. We then say that A believes $\Diamond\phi$ if and only if there exists a world $w \in B_A$ such that ϕ is true in w . This represents a very simple picture of belief in epistemic modals: a speaker believes that something is possible if there is a world consistent with her beliefs in which it is actual.

Our toy theory seems to work for some basic cases. If someone believes that it might be raining, then it must be consistent with her beliefs that it is raining, and therefore there must be a world in her doxastic space in which it is raining. However, the theory runs into trouble with more complex cases. Suppose that Alf has never given any thought to the city of Amsterdam, and perhaps does not even know that it exists. Then it is certainly consistent with his beliefs that it is raining in Amsterdam, and (9) seems to be true. However, things are not the same with (10):

- (9) Given Alf's beliefs, it might be raining in Amsterdam.
 # (10) Alf believes that it might be raining in Amsterdam.

Sentence (10) claims not only that it is consistent with John's beliefs, but something extra as well. However, according to the toy theory sketched above, (9) and (10) should be equivalent.

A similar problem arises in far more mundane cases.

Jane: Where are my keys?

John: I don't know, they could be anywhere in the house.

Though it would be fine for Jane to infer (11), it would be wrong for Jane to infer (12) from John's response.

- (11) For all John knows, the keys might be in the basement.
 # (12) John thinks that the keys might be in the basement.

The attitude report in (12) wrongly portrays John as having singled out the basement as a possible place for the keys. It would be fine if John had responded to Jane's question by saying 'I don't know, they might be in the basement', but as his response stands in this dialogue the inference to (12) is problematic. However, the toy theory again has problems here: John's response (given that it's felicitous) entails that it is consistent with his information that the keys are in the basement, so our toy theory predicts that (12) would be a proper inference from his response.

The general point of this data is that one cannot infer that X believes $\Diamond\phi$ from X's lack of belief in whether or not ϕ is the case. The toy theory fails to make this prediction because it identifies belief in the possibility that ϕ with lack of belief in $\neg\phi$, hence making insufficient information concerning ϕ sufficient for belief in

$\diamond\phi$.

We take this data to point to the fact that belief states need to be composed of salient possibilities and information. However, that is not the only conclusion that could be made from this data. Another conclusion would be that for an attitude holder to believe that ϕ is possible, she must have some reason to believe that ϕ is the case. This claim would account for the data, and has intuitive plausibility as well: if someone asserts $\diamond\phi$, then as discussed above we generally take that person to have some reason to believe ϕ is the case. Thus, if an assertion of possibly ϕ suggests the possession of reason to believe in ϕ , belief in possibly ϕ could easily suggest the same.

However, on closer inspection the data does not support this latter conclusion. Note that in the dialogue above, John may have good reason to believe that the keys are in the basement. Perhaps he has found them there in the past, or knows that Jane has a tendency to lose them there. However, he may still overlook the possibility that the keys are in the basement, and hence overlook the relevance of this evidence to the current conversation. In such a case, the attitude report (12) is still infelicitous, suggesting that what is at issue is not evidence, but attention to possibilities.

Further evidence that salient possibilities make up belief states comes from the use of epistemic modals to give advice or warnings. For example, in the dialogue below B gives a warning and a piece of advice to A.

A: Katherine and I are going fishing tomorrow in Leiden.

B: You might need a license.

A: Oh thanks, I hadn't thought of that.

After this dialogue, it would be very natural for A to check whether or not a license is needed for fishing in Leiden. Moreover, from A's last statement, it is clear that she would not have done this if B had not warned her. Therefore, in this conversation she must have realized something that led her to check on the need for fishing licenses in Leiden.

However, in this dialogue B does not give her any new information: B does not tell A that she needs a license. B does nothing more than bring A's attention to this possibility. Crucially, this is enough to lead A to change her behavior: A's new attention to the possibility that she needs a license to fish in Leiden gives her a reason to act differently from how she would have acted without this attention to the possibility. Therefore, attention to new possibilities can ground changes in action, or give agents new reasons to act.

Since attention to a possibility can provide an agent with a reason to act, this suggests that salient possibilities are part of belief states. This is because reasons to act are either beliefs or desires of some sort, and since an attention to a possibility is clearly not a desire it must be a kind of belief. Summarizing this argument, we have: (i) reasons for actions are beliefs and desires, but (ii)

an attention to a possibility can be a reason for an action, and it's not a kind of desire, so therefore (iii) an attention to a possibility must be a kind of belief. Since salient possibilities are possibilities that some kind of attention is focused on, we take this argument to be strong support that salient possibilities make up belief states.

To resist this argument, one could interpret the dialogue above quite differently. Instead of holding that A is just ignoring the possibility that she does not need a fishing license, one could hold that A is actually assuming that this is not the case. B, then, rather than focusing her attention on a possibility she had been ignoring, instead brings her attention to this assumption. A then revises her beliefs by giving up this assumption, and this change in beliefs is what explains her modified behavior.

The key problem with this interpretation of the dialogue is that it is too strong to say that A believes that she doesn't need a license to fish in Leiden. To see this, note that A may have simply never heard of fishing licenses in the first place. Perhaps she has never been fishing before and has no idea that these sorts of activities are regulated. In such a case, it is wrong to say that she really *believes* that she does not need a license. It is far more accurate to say that she just overlooks this possibility. After all, how could she believe that she does not need a license when she has never even considered the matter?

Thus, the plausibility of the reinterpretation of the dialogue hinges on a certain ambiguity in the notion of an assumption: sometimes assuming that ϕ means overlooking the possibility that $\neg\phi$, but other times it means believing without warrant that ϕ . So, it sounds intuitive to say that A is assuming that she does not need a license to go fishing. However, this is not because she believes that she does not need a license; it is simply because she overlooks the possibility that she needs one. And overlooking a possibility is very different from believing that it is not the case. The reinterpretation attempts to show that by giving up her assumption, A must change her beliefs. But giving up the assumption is not about revising beliefs, it is about attending to overlooked possibilities, and hence the reinterpretation of what takes place in the dialogue does not threaten our argument.

If, as we claim, salient possibilities are a feature of belief states, one would expect that some other attitudes are sensitive to salient possibilities. This is exactly what we find:

- (13) John forgot that the keys might be in the basement.
- (14) John remembered that the keys might be in the basement.
- (15) John wonders whether the keys might be in the basement.

When John forgets that the keys might be in the basement, he does not simply forget that his information fails to rule out this possibility, but also forgets to focus his attention on this possibility. The facts are similar in sentence (14).

Finally, in order for (15) to be true it must be the case that John is consciously attending to the possibility that keys are in the basement, i.e. it must be the case that the possibility that the keys are in the basement is salient in his wonderings.

Thus, we have two very different arguments for considering salient possibilities to be parts of belief states. First, it seems that attitude reports of the form ‘X believes $\diamond\phi$ ’ are sensitive to whether or not the attitude holder X has consciously attended to the possibility that ϕ . This suggests that in order for an agent X to believe that ϕ is epistemically possible, ϕ must be a salient possibility for X, and hence salient possibilities must be part of X’s belief state. Second, we argued that an attention to a possibility can be a reason for an action, and that since these sorts of reasons for actions are beliefs, salient possibilities must compose belief states. Below, we will sketch a formal theory of attitude reports that allows us to represent these facts.

4.2 Yalcin’s Theory

This data concerning belief reports has also been considered in Yalcin (2008). Following a strategy similar to the one used in this thesis, Yalcin (2008) incorporates the coarse-grainedness of attention to possibilities into his analysis of epistemic modality, conversational contexts, and belief states, and uses this analysis to account for the data concerning belief reports above. We will briefly present his account of coarse-grainedness and epistemic modals in attitude reports, but argue that it does not properly account for all of the data above.

Yalcin accounts for the coarse grained nature of attention with a concept called the *modal acuity* of an agent. The modal acuity of an agent is the specificity with which an agent distinguishes possible worlds from one another. As we mentioned above, possible worlds are not distinguished from one another individually, but based on whether or not the worlds satisfy certain propositions. Agents with greater modal acuity use several propositions to make fine distinctions between possible worlds, but agents with less modal acuity use only a few propositions to make heavy handed distinctions.

Since issues (in the sense of 3.3.2) distinguish worlds based on whether or not they satisfy propositions, the modal acuity of an agent is determined by the issues that the agent considers. Yalcin (2008) does not consider the sorts of issues raised by disjunctions, so the modal acuity of an agent is given by a partition of logical space. The partition corresponds to the questions an agent is considering, what Yalcin (2008) calls her inquiries. Such a partition is called a *modal resolution*. This terminology suggests a visual metaphor: if we think of logical space as a picture, we can view it in ‘high resolution’ or ‘low resolution’. In high resolution, we think of each distinct possible world as a pixel in the picture, producing a picture with very high resolution. But in low resolution, we think of the blocks of a modal resolution as pixels, producing something of much lower resolution.

Relative to a modal resolution, certain propositions are ‘visible’ and others

are ‘invisible’. A proposition (considered as a set of possible worlds) is visible in a modal resolution if it is the union of elements of the modal resolution. Intuitively, visible propositions are those that address the issues an agent is considering, and hence distinguish between worlds in a way that is relevant to the agent. If a modal resolution is thought of as a question, the visible propositions are the informative answers to that question. Since these propositions are relevant in the modal resolution, visible propositions ‘show up’ by being constructible from blocks of the modal resolution. Alternatively, invisible propositions are those that are not visible. These propositions provide information that does not directly address issues the agent is considering, and hence they do not ‘show up’ in an agent’s modal resolution.

At this point, it may seem that salient possibilities and modal resolutions amount to the same thing. If we think of agents and conversations as possessing modal resolutions, then since a modal resolution determines certain visible propositions, one might think that we could simply identify salient possibilities with visible propositions. However, this is not the case. If an agent is considering a question, then all possible answers to that question are visible in the agent’s modal resolution. Similarly, if participants in a conversation are considering a question, then all answers to that question will be visible in that conversation’s modal resolution. But as we saw in section 3, similar facts do not hold for salient possibilities. After a question is asked, each possible answer to that question does not correspond to a salient possibility. This is why epistemic modals can answer questions: they make certain answers to these questions salient in the conversation. Thus, there are generally far more visible propositions for an agent or a conversation than salient possibilities.

Using this notion of a modal resolution, Yalcin (2008) proposes the following analysis of epistemic modals under belief reports: ‘X believes $\diamond\phi$ ’ is true if and only if (i) ϕ is consistent with X’s information and (ii) ϕ is visible in X’s modal resolution. To believe an epistemic modal, the embedded proposition of the modal must be both consistent with an agent’s information and relevant to her inquiries. This accounts nicely for the data in (9) and (10) above. Since Alf has never considered whether it is raining in Amsterdam, the proposition that it is raining in Amsterdam is not visible on his modal resolution, explaining why (10) seems wrong.

However, this theory cannot account for cases like (11) and (12). Clearly, the location of Jane’s keys is one of John’s inquiries, and hence each possible location of the keys is visible in John’s modal resolution. Therefore, the theory proposed in Yalcin (2008) predicts that (9) should be perfectly acceptable, contrary to the facts we observed above. However, on our theory, there is something wrong with (9): since John never consciously thought of the basement as a possible location for the keys, this possibility is not salient in his belief state and hence the belief report is false.

This problematic prediction of Yalcin (2008) is related to the difference be-

tween visible propositions and salient possibilities. The notions of a modal resolution and a visible proposition do not invoke the idea of awareness or attention to possibilities. A modal resolution marks off the propositions that an agent *could* be interested in given her inquiries, not the ones she is *actually* interested in. This is why the proposal in Yalcin (2008) does not distinguish between the situations in (9) - (10) and (11) - (12), and hence makes problematic predictions.

Thus, though the proposal in Yalcin (2008) shares a great deal with the one in this thesis, there is a very important difference: awareness of possibilities is crucial here, but does not play a central role in Yalcin (2008). We now turn to a formal articulation of our own theory of epistemic modals in attitude reports.

4.3 Formal Theory of Attitude Reports

In this section, we use the formal system developed in section 3 to develop a semantics for attitude reports. More specifically, we enrich the language L defined above with a belief operator, and define the semantics for this operator with respect to information states and contexts. The crucial feature of our account is the analysis of belief states as contexts (as defined in section 3). Since an attitudes holder's beliefs are modeled by contexts, they contain both information (the maximal set in the context) and salient possibilities (the other information states in the context). With this formal system, we will show how one can account for the data presented above.

4.3.1 Belief States as Contexts

Recall that a context is a set of information states that contains a maximal information state. Above, we interpreted a context in the following way: the maximal information state was thought of as the common ground in a conversation, and the other information states were extensions of the common ground that interested the conversational participants. We now want to think about contexts in a different way: rather than thinking about contexts as linguistic contexts, we want to think of them as belief states. To do so, we now interpret the maximal information state as the information contained in an agent's beliefs. Following Stalnaker, we think of the information believed by an agent as a set of possible worlds, i.e. the set of worlds that are compatible with the agent's information.

The new feature of our account is the interpretation of the other information states in a context. These do not correspond to any information in an attitude holder's belief state, but rather correspond to salient extensions of the information in an attitude holder's belief state. These are extensions of an agent's believed information that come about if the attitude holder were to accept certain propositions that she is currently attending to. Suppose, for example, that an agent does not have any beliefs concerning whether or not it will rain tomorrow, but is actively considering both possibilities. Then her belief state will not contain

the information that it rains tomorrow, but it will contain an extension of her information in which it rains, and an extension in which it does not rain.

While it is relatively uncontroversial that agents do pay more attention to certain possibilities than to others, some may find it odd that we choose to include this in our representation of belief states. Are the possibilities an agent is focused on really part of her beliefs? Aren't beliefs just about information? But recall the data above. The contrasting minimal pairs in section 4.1.1 suggest that beliefs are not just about information. If a possibility ϕ is consistent with the information believed by an attitude holder x , an attitude report claiming that x believes might ϕ does not always follow. Moreover, the data about changing behavior based on utterances of epistemic modals suggest that changes in beliefs are not always changes in information believed. While it is unconventional to represent belief states as containing salient possibilities, we believe this is motivated by the data.

4.3.2 Semantics for Attitude Reports

We augment the language L defined in section 3 with a belief operator. To define the syntax of this operator, we define a special class of constants that name potential attitude holders. Note that we do not allow stacked belief operators.

Definition 4.1. (Belief Operator)

Fix a set Att of constants of L , the set of attitude holders. We then define the language L' as follows:

- (i) If $\phi \in L$, then $\phi \in L'$
- (ii) If $\phi \in L$, and $t \in \text{Att}$, then $B(t, \phi) \in L'$.

To define the semantics for belief reports, we slightly modify the definition of an index. Previously, an index was essentially a first order model combined with an assignment function. We now also let indices assign belief states to individuals.

Definition 4.2. (Belief Function)

Fix a subset A of the domain D , where A represents the set of attitude holders. We require that for all indices i and constants a , $i(a) \in A$ if and only if $a \in \text{Att}$. Each index i assigns a context Γ to each attitude holder index pair $\langle a, i \rangle$, where $a \in A$.

This is admittedly an overly simplistic way of formally treating attitude reports. However, nothing about our system precludes it from being implemented in a way that allows stacked operators and respects the principles of dynamic epistemic logic. We do not implement the system in this way because it is complex and not necessary for the purposes of arguing that belief states are composed of salient possibilities. We do not intend to define a full-blown dynamic epistemic logic in this thesis, but rather only intend to argue that beliefs states possess

both informative and non-informative components. Those readers interested in a more sophisticated method of handling belief reports are encouraged to consult Gerbrandy (1999).²

With this understanding of belief states, one can recover a classical model of belief states as sets of possible worlds. Given an attitude holder X , $i(X)_{CG}$ is the set of information believed by X at index i . This corresponds to a belief state in a traditional possible worlds semantics (e.g. in Hintikka (1962) or Stalnaker (1984)), and in our semantics informative beliefs will only be sensitive to this aspect of an attitude holder's belief state.

Updating an information state with a belief report $B(a, \phi)$ eliminates the indices that assign a context to a that is non-trivially updated by ϕ . For a wholly informative formula ϕ , this corresponds to adding the information that A 's beliefs already contain the information provided by ϕ . But for formulas that introduce salient possibilities, this corresponds to adding the information that A 's belief state must also contain any salient possibilities introduced by ϕ . Since belief reports do not introduce salient possibilities, the update that belief reports perform on contexts does not require much explaining. We define both kinds of update below.

Definition 4.3. (Updating Information States with Attitude Reports)

Let I be an information state, and let $\phi \in L'$. We define $I[\phi]$, the update of I with ϕ , as follows:

- Clauses (i) - (vii) are unchanged from the definition in section 3 above.
- (viii) $I[B(a, \phi)] = \{ \langle i, j \rangle \in I : i(i(a))[\phi] = i(i(a)) \text{, and } j(j(a))[\phi] = j(j(a)) \}$

We note that for any constant $a \in \text{Att}$, $i(i(a))$ is a context because $i(a)$ is an attitude holder.

Definition 4.4. (Updating Contexts with Attitude Reports)

Let $\Gamma = \{CG, P_1, \dots, P_n, \dots, \emptyset\}$ be a context, and let $\phi \in L'$. We define $\Gamma[[\phi]]$, the update of Γ with ϕ , as follows:

- Clauses (i) - (vi) are unchanged from the definition in section 3 above.

(vii) $\Gamma[[B(a, \phi)]] = \{CG[B(a, \phi)], P_1[B(a, \phi)], \dots, P_n[B(a, \phi)], \dots, \emptyset\}$

With this semantics, we can show the difference between believing that something might be the case and not believing anything inconsistent with it.

²However, we do not avoid all the complications encountered in dynamic epistemic logic with our simple model. Specifically, given our definition indices are elements of their own domains, which would create foundational problems if implemented in classical set theory. However, the definition is not problematic if one implements our system using non-well-founded set theory (see Gerbrandy (1999)). We do not explicitly introduce non-well-founded set theory in this thesis because an intuitive understanding of definition 4.2 will suffice for our purposes.

Believing Simple Sentences

Let $\Gamma = \{T \times T, \emptyset\}$. Then $\Gamma[[B(a, Pb)]] = \{\{i, j\} : i(i(a))_{CG}[Pb] = i(i(a))_{CG}$
and $j(j(a))_{CG}[Pb] = j(j(a))_{CG}\}, \emptyset\}$.

Thus, the result is a common ground that only contains indices i such that the attitude holder a 's believed information already contains the information provided by Pb . In other words, one removes indices which do not depict the attitude holder a as believing the information Pb .

We have the following fact regarding belief states and the informative fragment of our language.

Fact 4.5. (Purely Informative Attitudes)

Let ϕ be a formula of L that does not contain ' \diamond ' or ' \vee ', let $a \in \text{Att}$, and let i be an index. Then $i(i(a))[[\phi]] = i(i(a))$ if and only if $i(i(a))_{CG}[\phi] = i(i(a))_{CG}$

This follows directly from Fact 3.11. Thus, whether or not an attitude holder believes a purely informative formula is determined wholly by the informative component of an attitude holder's belief state.

Believing Epistemic Modals

Recall examples (9) and (10) above, let i be an index representing the actual world, let ϕ represent "it's raining in Amsterdam", and let a be Alf. Then $i(i(a))_{CG}[\diamond\phi] = i(i(a))_{CG}$, since the information Alf believes does not exclude the possibility of it raining in Amsterdam. However, since this is not a salient possibility for Alf, $i(i(a))_{CG}[\phi] \notin i(i(a))$, and hence $i(i(a))[[\diamond\phi]] \neq i(i(a))$. Therefore, for any information state I , $\langle i, i \rangle \notin I[B(a, \phi)]$.

Just as our framework makes room for two ways of answering questions, we also make room for two kind of beliefs: informative beliefs and non-informative beliefs. The information in a belief state represents an agent's informative beliefs, and the salient possibilities in a belief state represent the agents non-informative beliefs. Crucially, natural language belief reports are sensitive to this distinction, and can provide information about both an agent's informative beliefs and an agent's non-informative beliefs.

5 Conclusion and Further Issues

5.1 General Account of Indirectly Relevant Answers

In this thesis, we expanded our notion of a linguistic answer to allow non-informative statements to count as answers to questions. This allowed us to make sense of the fact that epistemic modals answer questions without asserting one of the question's logical answers. But there are several other constructions that answer questions without asserting logical answers:

A: Is John coming to the party?

B: I doubt it./Bill thinks so./Probably./He's supposed to come./He said he would.

Answers of this form differ from epistemic modals because they give information, but are similar in that they answer a question without providing the information it asks for. We call such answers 'indirectly relevant answers'.

As our theory stands, clearly we cannot account for indirectly relevant answers. However, our framework could be modified to provide an account of indirectly relevant answers. The key change we have to make is in our notion of a salient possibility in a conversation. Recall that in our framework we think of salient possibilities as certain propositions on which the discourse participants focusing. We could enrich this notion and think of salient possibilities as coming with a certain *force*, or a certain *character*. The force or character of a salient possibility would describe why the discourse participants are focusing on that possibility. Referencing the example above, a possibility could be salient because someone doubts it, because someone believes it, because it is epistemically probable, because someone is deontically obliged to make it the case, or because someone said it was actual.³

Under this enriched notion of a salient possibility, several natural language constructions introduce salient possibilities with different characters. A belief report 'x believes ϕ ' introduces a salient possibility with the force '*x believes*' constructed from ϕ . Statements of the form 'x is obliged to do ϕ ' introduce a salient possibility with the force '*x is obliged – to*' constructed from ϕ . Developing a compositional formal system in which such constructions act this way would likely be difficult, but the intuition behind the informal idea is very clear. We could then generalize the theory given above to account for all indirectly relevant answers: belief reports, statements of doubt, etc. raise salient possibilities with various characters, and they count as an answer to a question if the salient possibility they raise corresponds to a logical answer to the question. All indirectly

³Within this taxonomy, the salient possibilities introduced by epistemic modals are somewhat characterless: they can be interesting for a variety of reasons. The crucial feature of an epistemic possibility is that no one in the discourse knows it is false, but it can be interesting for various other reasons.

relevant answers function just like epistemic modals do as answers to question; the main difference is in the character of the possibility they introduce.

5.2 Epistemic Necessity

Aside from a few very basic remarks in section 2, this thesis has said very little about epistemic necessity. Since the bulk of the thesis focused on salient possibilities and non-informative content, this is not surprising. Epistemic necessity modals can be used to provide others with information, and do not introduce salient possibilities into a conversation.

While we could try to derive an account of epistemic necessity from our account of epistemic possibility by defining ' $\Box\phi$ ' as ' $\neg\Diamond\neg\phi$ ', the resulting account would not be very interesting. Essentially, we would have the account given in section 2.2.4. Though this theory can account for certain basic features of epistemic necessity, it makes ' $\Box\phi$ ' very similar to ' ϕ '. This similarity prevents our theory from accounting for our intuition that epistemic necessity modals express something about a speaker's (or attitude holder's) degree of commitment to a proposition and something about the sort of evidence a speaker (or attitude holder) has for a proposition.

(16) It must be raining.

(17) John believes that it must be raining.

An assertion of (16) suggests that the speaker arrived at the conclusion that it is raining through some kind of 'indirect evidence'. After all, (16) would be quite odd to assert if one is outside in the rain and seeing and feeling the water falling from the sky. Moreover, a speaker asserting (16) expresses a somewhat tenuous commitment to the proposition that it is raining. Though such a speaker surely commits to the proposition, the commitment is of a hypothetical nature: the speaker believes that it is raining on the basis of some other proposition, and this is signaled by the epistemic necessity modal. Similar facts hold for (17), only it suggests facts about the attitude holder's evidence for and degree of commitment to the proposition that it is raining.

Thus, it appears that in our theory 'must' is not the dual of 'might', and some may find this unappealing. But we do not think it is a bad result. Epistemic possibility modals introduce salient possibilities and pragmatically suggest a reason to believe that the possibility is actual. How could the dual transformation change such constructions into statements about indirect evidence and hypothetical commitment? It would be a very strange sort of negation that could accomplish such a feat.

To summarize, we conclude that there is much more work to be done on epistemic necessity, and that although this thesis's framework is not right for addressing it, this should not count against our analysis of epistemic possibility.

5.3 Conclusion

The analysis of epistemic modals given in this thesis is very dynamic. Epistemic modals play an important role in conversation as information indicators, but do not need truth conditional content to do so. Instead, their dynamic content, i.e. their functions as ‘tests’, allows epistemic modals to act as information indicators. Moreover, their non-informative content is very naturally described in a dynamic theory: the function of this content is, fundamentally, to *change* the salient possibilities in a conversation.

A crucial theme that emerged from this thesis is that there are important non-informative dimensions of belief and meaning. We have strong intuitions about the ability of epistemic possibility modals to express this non-informative content. In this paper, we converted these intuitions into linguistic data concerning attitude reports and question-answer pairs, and developed a formal theory about non-informative content. This theory of non-informative content has some surprising results (e.g. the failure of universal instantiation), but allows us to formally account for our important intuitions about epistemic modality. It is our conjecture that other facets of natural language have empirically important non-informative dimension to their content. These non-informative features may be very unlike the ones possessed by epistemic modals, but by formally analyzing the non-informative content of epistemic modals, we hoped to show in this thesis that these elements of linguistic meaning are not beyond the scope of formal semantics.

Natural language and belief are far too complex to be described with a simple truth conditional theory of meaning. An important feature of natural language is its ability to call attention to possibilities, and this feature is intractable on a truth-conditional theory. Thus, while the concepts of truth, reference, and information are crucial to understanding the semantics of natural language, it must be understood that they are not sufficient for the job.

References

- [1] Luis Alonso-Ovalle. *Disjunction in Alternative Semantics*. PhD thesis, University of Massachusetts at Amherst, Amherst, MA, September 2006.
- [2] Keith DeRose. Epistemic possibilities. *The Philosophical Review*, 100(4):581–605, 1991.
- [3] Andy Egan, John Hawthorne, and Brian Weatherson. Epistemic modals in context. In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, pages 131–168. Oxford University Press, 2005.
- [4] Jelle Gerbrandy. *Bisimulations on Planet Kripke*. Ph. d. dissertation, Universiteit van Amsterdam, 1999.
- [5] Anthony Gillies and Kai von Fintel. Might made right. Unpublished Manuscript, 2005.
- [6] Anthony Gillies and Kai von Fintel. An opinionated guide to epistemic modality. In Tamar Szabo Gendler and John Hawthorne, editors, *Oxford Studies in Epistemology: Volume 2*, volume 2, pages 32 – 62. Oxford University Press, 2007.
- [7] Anthony Gillies and Kai von Fintel. CIA Leaks. *The Philosophical Review*, 117(1):77 – 98, 2008.
- [8] H. P. Grice. Logic and conversation. In Peter Cole and Jerry L. Morgan, editors, *Syntax and semantics*, volume 3. New York: Academic Press, 1975.
- [9] H.P. Grice. Indicative conditionals. In *Studies in the Ways of Words*. Harvard University Press, 1991.
- [10] Jeroen Groenendijk. The logic of interrogation: classical version. In *Proceedings of the Ninth Conference on Semantics and Linguistics Theory (SALT-9)*. Publications, 1999.
- [11] Jeroen Groenendijk. Inquisitive semantics: Two possibilities for disjunction. In *Proceedings of Tbilisi 2007: Seventh International Tbilisi Symposium on Language, Logic and Computation*, 2008.
- [12] Jeroen Groenendijk and Martin Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14(1):39–100, February 1991.
- [13] Jeroen Groenendijk and Martin Stokhof. Questions. In Johan van Benthem and Alice T. Meulen, editors, *Handbook of Logic & Language*, chapter 19, pages 1055–1124. MIT Press - North Holland, 1997.

- [14] Jeroen Groenendijk, Martin Stokhof, and Frank Veltman. Coreference and Modality. In S. Lappin, editor, *The Handbook of Contemporary Semantic Theory*, pages 179–216. Blackwell, 1996.
- [15] C.L. Hamblin. Questions. *Australasian Journal of Philosophy*, 36:159 – 168, 1958.
- [16] Irene Heim. *The Semantics of Definite and Indefinite Noun Phrases*. Ph. d. dissertation, New York and London, 1988.
- [17] Jaakko Hintikka. *Knowledge and belief : an Introduction to the Logic of the Two Notions*. Cornell University Press, Ithaca, N.Y., 1962.
- [18] Hans Kamp. A theory of truth and semantic representation. In J. A. G. Groenendijk, T. M. V. Janssen, and M. B. J. Stokhof, editors, *Formal Methods in the Study of Language*, volume 1, pages 277–322. Mathematisch Centrum, Amsterdam, 1981.
- [19] Lauri Karttunen. *Discourse Referents*. Indiana University Linguistics Club, 310 Lindley Hall, Bloomington, Indiana, 1971.
- [20] Angelika Kratzer. What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy*, 1(3):337–355, January 1977.
- [21] Saul Kripke. *Naming and Necessity*. Harvard University Press, 1980.
- [22] Saul A. Kripke. Semantical considerations on modal logic. *Acta Philosophica Fennica*, 16:83–94, 1963.
- [23] John MacFarlane. Epistemic modals are assessment sensitive. Unpublished Manuscript, 2005.
- [24] G.E. Moore. *Commonplace Book, 1919–53*. Allen and Unwin, London :, 1962.
- [25] Mandy Simons. *Issues in the Semantics and Pragmatics of Disjunction*. Garland Publishers, 2000.
- [26] Robert Stalnaker. Assertion. *Context and Content*, pages 78–96, 1978.
- [27] Robert Stalnaker. *Inquiry*. MIT Press, March 1987.
- [28] Paul Teller. Possibility. *Philosophia*, (2):303 – 320, 1972.
- [29] Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25:221–61, 1996.
- [30] Seth Yalcin. *Modality and Inquiry*. Ph. D. dissertation, Massachusetts Institute of Technology, 2008.