

“Might” Made Right

Kai von Fintel
MIT

Anthony S. Gillies
University of Michigan

Colloquium at the Institute of Cognitive Science, Osnabrück

December 16, 2004

Abstract

In recent work, MacFarlane [14] and Egan et al. [3] have used new observations about epistemic modals to motivate radical new semantic analyses. MacFarlane introduces a triple-indexed semantics which makes sentences with epistemic modals not just sensitive to the context of utterance and an index of evaluation but also to a context of assessment. Egan et al. propose that epistemically modalized sentences do not express standard propositions but functions from individuals to propositions, something like centered propositions. We will argue that these moves are not only problematic but also unnecessary.

I The Standard Theory (Kratzer)

- (1) a. Every car must be inspected.
b. In view of what the laws are, every car must be inspected.
- (2) a. That coin might have landed heads.
b. In view of what the laws of physics provide, that coin might have landed heads.
- (3) a. Jackl might be the murderer.
b. In view of what is known, Jackl might be the murderer.

Kratzer’s picture [9, 10, 12]:

- Modals are quantifiers over possible worlds.
- Possibility modals are existential quantifiers, necessity modals are universal quantifiers. They are *duals* of each other.
- Modals quantify over the worlds given by a background such as “in view of what the laws are” (i.e. the worlds compatible with what the laws are in the evaluation world).
- The “ambiguity” between epistemic modals and deontic modals (and others) is really context-dependency.

- (4) $\llbracket \text{might } \phi \rrbracket^{c,i} = \text{I}$ iff $\exists w' \in f_c(i) : \llbracket \phi \rrbracket^{c, \langle w', t_i \rangle} = \text{I}$.
 $\llbracket \text{must } \phi \rrbracket^{c,i} = \text{I}$ iff $\forall w' \in f_c(i) : \llbracket \phi \rrbracket^{c, \langle w', t_i \rangle} = \text{I}$.

Comments: We work with a standard doubly-indexed semantics where we assign semantic values relative to a context of utterance c and an index of evaluation i . We assume that the index consists at least of a pair of an evaluation world and an evaluation time. For the most part, we will ignore time-dependency. f_c is a function supplied by the context assigning to the evaluation world w_i and the evaluation time t_i a set of accessible worlds. We get an epistemic reading of the modal when f_c is that function that assigns to any index the set of worlds compatible with “what is known” in w_i at t_i . But what does “what is known” mean?

Under one interpretation of “what is known”, we get *solipsistic* readings of epistemic modals. Kratzer [11]:

If epistemic interpretations of modals are relativized to the evidence available in the utterance situation, different utterances of one and the same sentence involving such a modal might express different propositions. Let us look at an example:

Suppose a man is approaching both of us. You are standing over there. I am further away. I can only see the bare outlines of the man, in view of *my* evidence, the person approaching may be Fred. You know better. In view of *your* evidence, it cannot possibly be Fred, it must be Martin. If this is so, *my* utterance of (5) and *your* utterance of (6) are both true.

- (5) The person approaching might be Fred.
 (6) The person approaching cannot be Fred.

Had *I* uttered (6) and *you* (5), both our utterances would have been false.

We would get this solipsistic interpretation if f_c assigns the worlds compatible with what the speaker of c knows in w_i . We will use s to refer to the speaker of c and will write f_s for a solipsistic knowledge base.

2 The Extended Standard Theory

Kratzer is not committed to the claim that epistemic modals are always to be understood solipsistically. There is a lot of suitable vagueness and flexibility in the expression “what is known”. There is in fact an entirely independent tradition of exploring what the range of epistemic readings is.

2.1 Hacking, Teller, DeRose

DeRose [1] presents a number of variations built around the following scenario:

- (7) John has some symptoms indicative of cancer. John’s doctor decides to run a “filtering” test, which has two possible results: If the results are “negative”, then cancer is conclusively ruled out; if the results are “positive”, then John might, but also might not, have cancer and further tests will have to be run.
- (8) DEROSE’S CANCER TEST CASE 1A (CTC-1A)
- a. Bill: “I have heard 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.”
 - b. John’s doctor, to a colleague: “It’s impossible that John has cancer, so we should start planning tests for other diseases.”
- (9) DEROSE’S CANCER TEST CASE 1B (CTC-1B)
- Bill: “I’ve heard that John may have cancer. Is that possible?”
Jane: “I don’t know whether it is possible that John has cancer; only the doctors know. I’ll find that out tomorrow when the results of the test are revealed.”

Crucial: Jane can truthfully say that she doesn’t know whether it is possible that John has cancer, which entails that for all she knows, it might not be possible, even though clearly it is compatible with everything that *she* knows that he has cancer. So, this use of *possible* must be sensitive to more than Jane’s knowledge state.

- (10) DEROSE’S PROPOSAL – FIRST VERSION
it is possible that p is true in world w iff *no member of the relevant community* knows in w that p is false. (DeRose borrowed this idea from Hacking [6]).

Flexibility: in CTC-1A, the relevant community is Jane (and possibly Bill); in CTC-1B, it is a somewhat larger group of people that includes John’s doctors. (DeRose attributes the idea of FLEXIBILITY OF THE RELEVANT COMMUNITY to Teller [22]).

We will reformulate the proposal as one about *might* as follows:

- (II) $\llbracket \text{might } \phi \rrbracket^{c,i} = \mathbf{I}$ iff $\forall x \in G_c : \exists w' \in f_x(i) : \llbracket \phi \rrbracket^{c, \langle w', t_i \rangle} = \mathbf{I}$.

Here, G_c is the contextually supplied group (typically including the speaker). f_x delivers the set of worlds compatible with what x knows.

Note: This semantics does not reduce to the Kratzer-schema: there is not a single modal base which is checked for compatibility with ϕ . The semantics iterates over the knowledge bases of the individual members of G_c .

2.2 More Objectivity

A wrinkle: sometimes it seems that facts not in the possession of anyone in the relevant community count for the truth of an epistemic modal.

- (I2) HACKING’S SALVAGE SHIP CASE
 Imagine a salvage crew searching for a ship that sank a long time ago. The mate of the salvage ship works from an old log, makes some mistakes in his calculations, and concludes that the wreck may be in a certain bay. It is possible, he says, that the hulk is in these waters. No one knows anything to the contrary. But in fact, as it turns out later, it simply was not possible for the vessel to be in that bay; more careful examination of the log shows that the boat must have gone down at least thirty miles further south. The mate said something false when he said, “It is possible that we shall find the treasure here,” but the falsehood did not arise from what anyone actually knew at the time. [6].
- (I3) “Why was the mate speaking falsely when he said, “It is possible that the hulk is here”? Because one could have found out from the data that the wreck took place a good deal further south.” (Hacking)
- (I4) DEROSE’S CANCER TEST CASE 2B (CTC-2B)
- a. 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.
 - b. Jane to Bill: “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.”
- (I5) DEROSE’S PROPOSAL – SECOND AND FINAL VERSION
- it is possible that p is true in world w iff*
- (i) no member of the relevant community knows in *w* that *p* is false,
 - (ii) there is no relevant way in which members of the relevant community can come to know that *p* is false.

2.3 Complaints

We present two ways in which DeRose’s proposal still seems to fall short of perfection.

2.3.1 Expectant Fathers, Drawn Curtains, etc.

(16) TELLER’S EXPECTANT FATHER CASE

At the time of writing I am an expectant father, and one will grant that I speak truly when I say, “It is possible that my child will be a boy,” and “It is possible that my child will be a girl.” Most readers are probably informed here for the first time that there is a practicable, in fact quite easy test which will establish the sex of my expected child. . . . On Hacking’s definition one of my two statements must be false. But I submit that they are both true, and that our judgment that they are true is not affected by knowledge that a test for sex is available before birth. The doting grandmother who agonizes, “It’s possible it will be a boy, it’s possible it will be a girl. Should I buy blue or pink?” is not shown by the availability of the test to have said anything false. Nor will grandmothers who learn about the test abandon such locutions. [22]

(17) GILLIES’ RAIN BEHIND THE CURTAIN CASE

You come to my office one afternoon. The curtains are drawn. We have a nice chat, and you are about to head home. I have not been outside since early morning, but the forecast was for a 50% chance of rain. In such a situation, it seems perfectly acceptable for me to say before you leave *It might be raining out*. What does [DeRose] predict? Well, that all depends on condition (ii). What counts as a relevant way is supposed to be supplied by context. But it is hard to imagine that the context just described is a context in which opening the curtains does not count as a relevant way of finding out the fact of the matter about the current weather. If that is so, then I have not only said something inappropriate by saying *It might be raining*. According to [DeRose], I have said something false. And this is counter-intuitive. [5].

Problem #1: In these cases, there seem to be “relevant ways in which members of the relevant community can come to know that p is false”, but nevertheless we can truthfully say *might p*. In other words, unless the application of “relevant way of coming to know” gets tailor-made, the analysis is wrong.

2.3.2 Distributed Knowledge

Consider a variant of DeRose’s Cancer Test Case. Here, two doctors have run two different tests, but they have not communicated with each other. Everyone, the doctors and John and his wife, will come together on Monday to compare results and decide on what to do next. As it happens, the two tests together rule out cancer.

Even though Jane knows that nobody yet knows whether John has cancer, she can say:

- (18) We don't know whether it is possible that John has cancer. The doctors haven't compared their results yet.

Problem #2: Again it seems that knowledge not in the possession of any agent matters to truth of an epistemic possibility claim. But here it is knowledge that is already present in the group, but in a “distributed” way. Perhaps, we should make space for that in the analysis.

2.4 What Do We Know?

Our proposal to amend the Standard Theory has two parts:

1. base the semantics of *might*-statements in the DISTRIBUTED KNOWLEDGE of a group G (solving Problem #2)
2. treat certain coherent bodies of evidence (computer printouts, ship logs, ...) as potentially part of G (solving Problem #1)

2.4.1 Notions of Group Knowledge

We can distinguish several different notions of group knowledge.

1. A very strict notion is COMMON KNOWLEDGE which requires that something (p) is known by G only if everyone in G (i) knows p and (ii) knows that everyone else in G knows p too and (iii) knows that everyone else in G knows that everyone else in G knows p and
2. A somewhat looser notion is FIRM KNOWLEDGE which merely requires that everyone in G knows it, without the additional mutual knowledge conditions.
3. A yet looser notion is DISTRIBUTED KNOWLEDGE which just requires that if all the knowledge in the group is pooled, the group would know.

At the level of knowledge sets (sets of worlds compatible with one's knowledge), we can define:

- (19) For each member x of a group G, we have the set of worlds $f_x(w)$ compatible with what x knows in w .
- (20) $E_G(w) = \bigcup_{x \in G} f_x(w)$ (the union of what the x 's consider possible)
 \Rightarrow firm knowledge in G

- (21) $D_G(w) = \bigcap_{x \in G} f_x(w)$ (the intersection of what the x 's consider possible)
 \Rightarrow distributed knowledge in G

Note: the weaker a notion of group knowledge we employ, the more things a group knows, the smaller the set of worlds compatible with the group's knowledge gets, and thus the stronger an epistemic possibility statement becomes! [PICTURE ON BLACKBOARD]

- (22) [Beginning of a brainstorm session among some dedicated members of MIT's Modality Lab:] What do we know about epistemic modals?
- (23) [Comment of a 9/11 widow on NPR, 4/15/2004:] We knew more than is being owned up to. But nobody put the pieces together.

Our proposal: epistemic *might* is interpreted with respect to the distributed knowledge of the relevant group.

- (24) OUR SEMANTICS FOR 'MIGHT'
 $\llbracket \text{might } \phi \rrbracket^{c,i} = 1$ iff $\exists w' \in \bigcap_{x \in G_c} f_x(w) : \llbracket \phi \rrbracket^{c, \langle w', t_i \rangle} = 1$

Note: the solipsistic reading is the special case where $G_c = \{s\}$. So, our semantics covers both the solipsistic reading and the group-based readings. Context determines the extent of G_c .

2.4.2 Bodies of Evidence

What about Problem #1? What is the difference between sealed computer printouts or ship's logs on the one hand and possible intra-uterine tests or looking behind the curtains on the other hand?

Our proposal: the evidence contained in computer printouts or ship's logs are treated as honorary individual members of the group G . Directly looking at the facts is a way of finding out what's going on but is not thought of as a body of evidence.¹

¹ Compare this to the conceptualization behind systems of EVIDENTIAL MARKING in certain languages [4, 17, 18]: direct visual experience is distinguished from information encoded in documents or reported by other people. The latter are marked by INFERENCEAL or INDIRECT markers.

Our idea connects to work where epistemic modals are claimed to be evidential markers, for example Westmoreland [23]. It seems clear to us that epistemic modals do not behave exactly like evidential markers, see Faller [4] for arguments, but at the same time they seem to combine a standard modal semantics with constraints about what kind evidence they are sensitive to.

- (i) [Looking out the window at the pouring rain:]
 ??It must be raining.
- (ii) [Looking at people coming in with wet umbrellas:]
 It must be raining.

3 The Relativists

3.1 Wrong!

Hawthorne [7]:

[A]s far as I can tell, ordinary people evaluate present tense claims of epistemic modality as true or false by testing the claim against their own perspective. So, for example suppose Angela doesn't know whether Bill is alive or dead. Angela says *Bill might be dead*. Cornelius knows Bill is alive. There is a tendency for Cornelius to say Angela is wrong. Yet, given Angela's perspective, wasn't it correct to say what she did? After all, when I say *It might be that P and it might be that not P*, knowing that Cornelius knows whether P, I do not naturally think that Cornelius knows that I said something false. There is a real puzzle here, I think, but this is not the place to pursue it further.

MacFarlane [14]:

- (25)
- a. Sally: Joe might be in Boston.
 - b. George: He can't be in Boston. I saw him in the hall five minutes ago.
 - c. (i) Sally: Oh, then I guess I was wrong.
 - (ii) Sally: Oh, OK. So he can't be in Boston. Nonetheless, when I said "Joe might be in Boston," what I said was true, and I stand by that claim.

On a solipsistic understanding, (25-a) would be true iff Joe being in Boston is compatible with what Sally knows at the time of utterance. Finding out *later* that Joe isn't in Boston can't change that. But then what she says in (25-c-i) is surprising, and we should expect that (25-c-ii) should be a fine thing to say.

An immediately obvious response would be to say that George's knowledge counts too because as Sally's addressee he is part of the relevant group whose knowledge counts. But against this, MacFarlane presents this variant:

- (26)
- a. Sally: Joe might be in Boston.
 - b. George: Oh, really? I didn't know that.
 - c. (i) Jane (overhearing and not party to the conversation): Sally is wrong. I saw Joe just a few minutes ago.
 - (ii) Jane (overhearing and not party to the conversation): Joe can't be in Boston. I saw him just a few minutes ago. Nonetheless, what Sally said is true.

On a more complex, but still subjective/contextualist semantics, (26-c-i) would be unexpected, and (26-c-ii) should be a fine thing to say, contrary to fact, or so the relativists claim.

3.2 CIA Theories

The remedy recommended by MacFarlane [14] and Egan et al. [3], see also Egan [2]: semantic interpretation is relative not just to a context of utterance and an index of evaluation but also a context of assessment. That is, we need a triple-indexed semantics.

- (27) ASSESSMENT-RELATIVE MEANING FOR ‘MIGHT’
 $\llbracket \text{might } \phi \rrbracket^{c,i,a} = \mathbf{I}$ iff $\exists w' \in f_a : \llbracket \phi \rrbracket^{c,<w',t_i>,a} = \mathbf{I}$.

where f_a is the set of worlds compatible with what the assessor knows at the time of the assessment.

Note: we are glossing over details of implementation that differ among the various relativist proposals. But (27) is the core of the relativist idea.

So, what goes on in a sequence like

- (28) Sally: Joe might be in Boston.
 Jane (hidden in the bushes): Sally is wrong. I saw Joe just a few minutes ago.

is that when Jane is assessing the truth of Sally’s statement, the modal base automatically shifts to *Jane’s* knowledge base, because *might* is assessment-sensitive.

[**Note:** a similar semantics has recently been proposed by Lasersohn [13] for predicates of personal taste. We can’t comment on his analysis here, but suspect that a non-relativistic semantics may also work for his cases.]

3.3 Trouble in Relativity Town, or Down with the CIA!

1. We have some worries about the central “I was wrong” kind of data. It needs to be carefully argued that this really assesses the truth-value of the original statement. An obvious alternative diagnosis might be that Sally is admitting that she was wrong to consider it possible that Joe was in Boston, i.e. that her cognitive state wasn’t what it should have been – which would still leave her original statement as a true description of her cognitive state as it was. . . . We propose to suspend our disbelief for the remainder of this talk and proceed as if the data are as claimed by the relativists.
2. Another worry is that the assessment-sensitive semantics seems to wreak havoc with the standard picture of communication as developed by Stalnaker [20], where propositions are at the center of the notion of dialogue. It is unclear what will happen to that picture when we have assessment-sensitive sentences – but see Macfarlane [14] and Egan [2] for some ideas, which we are quite skeptical about. Again, we will suspend our disbelief.

3. Having said that, we do want to present some troublesome data. Consider this dialogue:

- (29) a. Kai: I am looking for my keys.
b. Thony: Might they be in your desk?

Why would Thony ask what he asks if the truth-value of *The keys might be in Kai's desk* for him entirely depends on his (the assessor's) knowledge state.

4. We have actually already two instances of a kind of example that seems very problematic for the relativists:

- (30) It's possible it will be a boy, it's possible it will be a girl. Should I buy blue or pink? [from Teller]
- (31) “After all, when I say *It might be that P and it might be that not P*, knowing that Cornelius knows whether P, I do not naturally think that Cornelius knows that I said something false.” [Hawthorne]

Or consider Sally:

- (32) Joe might be in Boston and he might be in New York. I just don't know.²

She would not say “I was wrong” when she finds out that Joe is in Boston, even though her statement entailed that Joe might be in New York, which according to the relativists is false from her new state of evidence.

Similarly, the relativists predict that an assessor who knows where Joe is (in New York, say) would find Sally's utterance false, which seems an absurd prediction to us.

A quantificational variant

- (33) There are many places the keys might be.

should be – but clearly isn't – a statement doomed to quick failure as soon as the assessor knows where the keys are or even figures out that there are only two places that they could be.

² It is often more natural to give lists of epistemic possibilities by using the connective *or*. But this does not affect the point made here. The sentence *Joe might be in Boston or he might be in New York* clearly entails that Joe might be in Boston. Why that entailment holds is a puzzle addressed in quite a bit of recent work on the interaction of disjunction and modality.

4 The Revised Extended Standard Theory

Our main point: once we allow for the flexibility of the relevant group whose distributed knowledge underlies a *might*-claim, we can explain the new data presented by the relativists. At the same time, we don't have the conceptual troubles they have and we don't have a problem with the data we just presented. In other words, the relativistic analyses are not only problematic but also unnecessary.

The main point we need to address, we feel, is the case of eaves-dropping, repeated here:

- (34)
- a. Sally: Joe might be in Boston.
 - b. George: Oh, really? I didn't know that.
 - c. (i) Jane (overhearing and not party to the conversation): Sally is wrong. I saw Joe just a few minutes ago.
 - (ii) Jane (overhearing and not party to the conversation): Joe can't be in Boston. I saw him just a few minutes ago. Nonetheless, what Sally said is true.

We claim that Jane can be seen as part of the community with respect to which Sally's utterance is evaluated, without having to resort to an assessment-sensitive semantics.

When the interpretation of a sentence is sensitive to how the context resolves an indeterminacy in its logical form – and if the speaker chooses not to reduce the indeterminacy –, the speaker exposes themselves to challenges. As we have seen, epistemic modals need as one contextual parameter the group G whose knowledge is relevant to the epistemic claim. Unless there is a very specific *in view of* phrase or very strong contextual clues, there will be indeterminacy, that is there will be a non-trivial set of possible values of G . In such a case, the speaker is responsible for the fact that their statement could be interpreted with respect to a group G that is much bigger – and thus makes the statement much stronger – than it would be under a solipsistic interpretation or even under an interpretation where the speaker and the addressee are the only relevant members of G .³

In general, we would like to model contextual indeterminacy by saying that there is a set C of admissible contexts. Roughly speaking, C contains those contexts which the speaker should reasonably expect to be possible contexts with respect to which the utterance might be evaluated. What distinguishes items like epistemic modals from items like the indexical I is that all of the admissible contexts will have the speaker be the referent of I , while admissible contexts may vary quite a bit as to the value the group G parameter.

³ We owe the germ of this idea to a comment made by Angelika Kratzer when Kai von Stechow gave a colloquium about epistemic modals and conditionals at UMass on December 12, 2003.

The behavior of epistemic modals is found with other “contextually restricted” quantificational expressions as well, but with illuminating differences:

- (35) Sally: Every student was at the meeting.
 George: What, even those that are on leave in Nicaragua?
 Sally: No, what I meant was every student in residence.

In (35), Sally retreats to a more specific resolution of her initial statement. This kind of retreat seems unnatural with *might*:⁴

- (36) (continuing the dialogue in (25)):
 Sally: Well, what I meant was that given what I knew then, Joe might be in Boston.

The explanation is straightforward: Sally’s prior epistemic state is of little interest once George’s information is revealed.

The speaker’s responsibility should not reach beyond a reasonable cloud of contexts:

- (37) *Detective Parker is reading court transcripts from the 1920s where Capone is on the stand being asked about where some money is in relation to a particular safe*
 ...
 Capone: The money might be in the safe.
 Parker: ??Al was wrong/That’s false. The safe was opened in 1956 and found to be empty.

It is hard to imagine a sense in which Capone and Parker form a relevant community whose knowledge Capone’s statement could be sensitive to. In contrast, we argue that the “stranger in the bushes” scenario in (34) is one where Sally’s utterance is transported into a context that she may well not have anticipated, but where it could still be felt that Jane-in-the-bushes and Sally are engaged in the same “investigation” and that therefore there is an admissible context assigning to the modal the evidence shared between Sally and Jane.

The crucial point of our analysis is that when faced with indeterminacy it is a mistake to assume that “the context” will resolve the indeterminacy. Instead, participants in a conversation need to take into account a cloud of admissible contexts within which their conversation is situated. We anticipate that our analysis may also serve as an alternative to other proposed applications of MacFarlane’s notion of assessment relativity, for example the contextual relativity of knowledge attributions [15].

⁴ Actually, it doesn’t sound so odd for Sally to demur: “Look, I said ‘might be’ – I didn’t claim that he was definitely in Boston.” But we had agreed not to challenge the relativistic view of the data too much.

5 The Dynamic Future

Some issues that we probably don’t have time to talk about today but will address in our paper:

- What about epistemic *must*? Does it also refer to the distributed knowledge of the relevant group? Note that if so, the resolution of G that makes a *must*-statement strongest is the solipsistic one.
- Our “objective”, non-solipsistic analysis of many *might*-statements entails that often speakers who assert these statements cannot possibly know whether they are true, since that depends on what other people know. So, when Thony says:

(38) The keys might be in your desk.

his statement depends on what I know, something he can’t fully know. This is a problem if for a speaker to assert a sentence correctly, they should believe that the sentence is true. Thony can’t really be in such a position. Our guess: the proper view is that non-solipsistic *might*-statements are more like conjectures and as such are not subject to the belief-condition.

- Embedding of epistemic modals raises many interesting issues. We suspect that the relativistic analyses have not much to offer here. We do anticipate that the behavior of the G-parameter is peculiar in that it does not behave like a run-of-the-mill indexical. In fact, one possibility is that instead of the *we*-type meaning we’ve been assuming, the G-parameter behaves more like the “generic pronoun” *one* discussed by Moltmann [16] and Safir [19], among others.⁵
- Epistemic modals have been one of the show-pieces of dynamic semantics and we intend to investigate how our proposal could be integrated into a dynamic semantics or into a formal theory of discourse dynamics.

So, stay tuned – and please send us any comments you might have.

⁵ Keshet [8] in fact explores such an analysis for the taste predicates for which Lasersohn had presented a relativistic account.

Bibliography

- [1] DEROSE, Keith: 1991. “Epistemic Possibilities.” *The Philosophical Review*, 100(4): 581–605.
- [2] EGAN, Andy: 2004. “Epistemic Modals, Relativism and Assertion.” URL <http://www.geocities.com/eganamit/might.doc>. Ms.
- [3] EGAN, Andy, HAWTHORNE, John, & WEATHERSON, Brian: 2003. “Epistemic Modals in Context.” URL <http://brian.weatherson.net/em.pdf>.
- [4] FALLER, Martina: 2002. *Semantics and Pragmatics of Evidentials in Cuzco Quechua*. Ph.D. thesis, Stanford. URL <http://lings.ln.man.ac.uk/Info/staff/MTF/publications.htm>.
- [5] GILLIES, Anthony S.: 2001. “A New Solution to Moore’s Paradox.” *Philosophical Studies*, 105: 237–250. doi:10.1023/A:1010361708803. URL <http://www-personal.umich.edu/~thony/moore.pdf>.
- [6] HACKING, Ian: 1967. “Possibility.” *The Philosophical Review*, 76(2): 143–168.
- [7] HAWTHORNE, John: 2004. *Knowledge and Lotteries*. Oxford University Press. Preprint of Chapter 1 available at <http://www.sims.berkeley.edu/~kahern/hawthorne/IntroducingThePuzzle.htm>.
- [8] KESHET, Ezra: 2004. “A Matter of Taste.” Unpublished term paper, MIT.
- [9] KRATZER, Angelika: 1977. “What *Must* and *Can Must* and *Can Mean*.” *Linguistics & Philosophy*, 1: 337–355.
- [10] KRATZER, Angelika: 1981. “The Notional Category of Modality.” In H. J. Eikmeyer & H. Rieser (Editors) *Words, Worlds, and Contexts. New Approaches in Word Semantics*, pages 38–74. de Gruyter.
- [11] KRATZER, Angelika: 1986. “Conditionals.” *Chicago Linguistics Society*, 22(2): 1–15. Reprinted in [21], pages 651–656.
- [12] KRATZER, Angelika: 1991. “Modality.” In [21], pages 639–650.
- [13] LASERSOHN, Peter: 2004. “Context Dependence, Disagreement, and Predicates of Personal Taste.” URL <http://semanticsarchive.net/Archive/mNjMjlkZ/taste.pdf>. Ms.
- [14] MACFARLANE, John: 2003. “Epistemic Modalities and Relative Truth.” URL <http://socrates.berkeley.edu/~jmacf/epistmod.pdf>.
- [15] MACFARLANE, John: 2004. “The Assessment Sensitivity of Knowledge Attributions.” URL <http://philosophy.berkeley.edu/macfarlane/relnow.pdf>. To appear in *Oxford Studies in Epistemology*.

- [I6] MOLTSMANN, Friederike: 2003. “Generic *One* and the First Person.” URL http://www.philosophy.stir.ac.uk/staff/Friederike_Moltmann_files/generic-one.doc. Ms.
- [I7] ROORYCK, Johan: 2001. “Evidentiality, Part I.” *GLOT International*, 5(4): 125–133.
- [I8] ROORYCK, Johan: 2001. “Evidentiality, Part II.” *GLOT International*, 5(5): 161–168.
- [I9] SAFIR, Ken: 2004. “What Does *One* Mean When One Says It?” URL <http://www.fas.harvard.edu/~lingdept/IndexicalityWorkshop/safir-one.pdf>. Notes for a talk given at Harvard, November 19, 2004.
- [20] STALNAKER, Robert: 1978. “Assertion.” In Peter Cole (Editor) *Syntax and Semantics, Vol. 9: Pragmatics*, pages 315–322. New York: Academic Press.
- [21] VON STECHOW, Arnim & WUNDERLICH, Dieter (Editors) : 1991. *Semantics: An International Handbook of Contemporary Research*. de Gruyter.
- [22] TELLER, Paul: 1972. “Epistemic Possibility.” *Philosophia*, 2: 302–320.
- [23] WESTMORELAND, Robert R.: 1995. “Epistemic *must* as Evidential.” In Paul Dekker & Martin Stokhof (Editors) *Proceedings of the Tenth Amsterdam Colloquium. December 18-21, 1995. Part III*, pages 683–702. ILLC, University of Amsterdam.

Authors’ Addresses

Kai von Fintel
 Department of Linguistics & Philosophy
 Massachusetts Institute of Technology
 Room 32-D808
 77 Massachusetts Avenue
 Cambridge, MA 02139-4307
fintel@mit.edu
<http://mit.edu/fintel>

Anthony S. Gillies
 Department of Philosophy
 University of Michigan
 Ann Arbor, MI 48104
thony@umich.edu
<http://www-personal.umich.edu/~thony/>