Deriving a Variable-Strength *Might*[1]

Deniz Rudin — University of California, Santa Cruz

Abstract. This paper combines an empirical argument about the lexical semantics of *might* with a preliminary description and theoretical account of a novel variety of implicatures. Empirically, I introduce the DISMISSIVE AGREEMENT paradigm, which shows that *might* semantically encodes nothing stronger than nonzero probability. Theoretically, I derive the fact that *might* often seems to suggest something stronger from the pragmatic norm that cooperative speakers will make claims that are strong enough to be relevant to the Question Under Discussion, which gives rise to LOWER BOUND IMPLICATURES.

Keywords: epistemic modality, implicature, QUD, relevance.

1. What's In This Paper

This is a paper about *might*. It begins with the observation that all theories of the semantics of *might* must assign to it either a WEAK meaning (*might* entails only that its prejacent is not strictly impossible) or a STRONG meaning (*might* can entail that its prejacent is more than merely non-impossible—that its prejacent is fairly likely, or is worth devoting attention to, or is especially plausible/normal/stereotypical). Though a broad variety of both weak and strong semantics for *might* have been proposed, no investigation of which I am aware systematically examines the empirical evidence supporting each view. This paper fills that gap.

The major empirical contribution of this paper is the DISMISSIVE AGREEMENT paradigm, which has not been discussed previously in the literature:

(1) DISMISSIVE AGREEMENT:
A: Paul might come to the party.
B: Yeah, he might, but it's extremely unlikely.

That it's possible to agree with a *might*-claim while simultaneously dismissing it as extremely unlikely is very difficult to reconcile with many theories of *might*. In §2 I present the relevant theoretical arguments in tandem with an investigation of novel data and conclude that *might* can only be WEAK. I briefly present a standard weak semantics for *might* in §3.

[1]This work benefitted immeasurably from the intellectual generosity of Pranav Anand, Adrian Brasoveanu, Ivano Caponigro, Karl DeVries, Donka Farkas, Valentine Hacquard, Margaret Kroll, Dan Lassiter, Chris Potts, Erik Zyman, everyone involved in the 2014-2015 graduate research seminar sequence at UCSC, and audiences at LASC 2015, the 2015 UCSC Graduate Research Symposium, CUSP 8 and Sinn und Bedeutung 20. Though the abovementioned deserve partial credit for whatever may be good about this paper, I greedily retain sole ownership of any mistakes, unclarities or inconsistencies herein.

However, the intuitions underlying the STRONG family of theories of *might* are palpably clear, and deserve explanation. In §4 I propose that strengthening of *might* can be derived from standard Gricean reasoning. I propose that QUDs are accompanied by PROBABILITY GRAINS that legislate the grain size of probability relevant to answering that QUD; the assumption that cooperative speakers will only point out differences in probability that are large enough to be relevant causes *might*-claims to generate a LOWER BOUND IMPLICATURE, which pragmatically strengthens their meaning. In §5 I elaborate on the nature of Probability Grains, in §6 I explore some differences between the behavior of lower bound implicatures and the behavior of scalar implicatures, and in §7 I argue that lower bound implicatures are not particular to *might*, and in fact can be seen occurring with a broader variety of existentials.

2. The Empirical Terrain: Weak or Strong?

I'll call sentences like (2) *might*-claims:

(2) Paul might weigh 180 pounds.

In sentences like this, *might* takes scope over its prejacent (*Paul weighs 180 pounds* in the sentence above) and, roughly speaking, converts it from an assertion that the prejacent is true to an assertion that the prejacent could possibly be true. I will refer to a *might*-claim with prejacent p as *might-*p.

What precisely does a *might*-claim entail of its prejacent? A WEAK theory of the semantics of *might* takes *might-*p to entail only that p is not strictly impossible. For instance, Kratzer (1977) proposes that *might-*p is true iff there is at least one epistemically accessible world in which p is true, and Veltman (1996) proposes that an update with *might-*p doesn't alter any context that already contains at least one world in which p is true. In contrast, a STRONG theory of *might* allows *might-*p to entail something more of p. For instance, Kratzer (1981) associates *might* with 'human possibility', on which account *might-*p is true iff there is at least one world in a special subset of especially likely/plausible/normal worlds in which p is true, and Willer (2013) associates *might* with 'live possibility', such that an update with *might-*p has the effect of establishing p as a possibility that should be taken seriously. Some theories, like those of Swanson (2006) and Lassiter (2011), take *might-*p to entail that the likelihood of p is greater than a contextually specified threshold value; I lump these theories in with the strong theories, because they allow for *might-*p to entail something stronger than that p is not impossible, though they do not necessitate that it always entails something stronger.

The intuition behind strong theories is clear. Consider the following *might*-claim in its given context:

(3) *Context: Your friend Paul lives on the East Coast. You haven't heard from Paul in a while, and know nothing of his plans or specific whereabouts. I assert the following to you:*

Paul might come to our party in Santa Cruz next weekend.

A natural response to my assertion would be for you to feel surprised, and to assume that I have access to some evidence or information about Paul's plans and whereabouts. If all that my *might*-claim communicated was that its prejacent is not impossible, it would be puzzling for you to be surprised by my statement, as you know nothing that rules out the possibility of Paul making his way to Santa Cruz by next week; likewise, there would be no reason to assume that I have access to any particular information about the prejacent, because it does not require any special knowledge to have realized that Paul coming to Santa Cruz next week is not strictly impossible.

That *might*-claims often seem strong is demonstrated even more clearly by considering dialogues like the following:

(4) **A:** Paul might come to the party.
   a. **B:** Oh really? I didn't know that!
   b. **B:** I guess we should buy some more snacks!

B reacts to **A**'s *might*-claim as though **A** has said that it's fairly likely that Paul will come to the party, or that Paul coming to the party is a possibility that deserves serious attention. Presumably, **B** did not think it was strictly impossible that Paul would come, and yet she can announce that she was unaware of the contents of the *might* claim, as in (4a); (4b) seems like a sensible response to **A**'s assertion, even though buying extra snacks to accommodate Paul only makes sense if there's a fairly large chance that he'll come. **B**'s response to **A**'s *might*-claim is perfectly congruent with a strong theory of *might*, but is *prima facie* mysterious from the perspective of a weak *might*.

The following dialogue makes a similar point:

(5) **A:** Will Paul come to the party?
   B: He might.

A's question indicates that she isn't ruling out the possibility that Paul will come, but that she isn't certain that he will either. If **B**'s *might*-claim communicated only that it's not impossible that Paul will come to the party, it should be a strictly uncooperative response, as it merely reiterates a possibility that **A**'s question indicates that she is not ruling out. However, **B**'s response feels like an informative contribution. Again, this makes perfect sense from the perspective of a strong theory of *might*, but is somewhat mysterious on a weak theory.

Data like the above make it quite clear that *might*-claims can communicate something stronger than that their prejacent is merely non-impossible. However, I've chosen the word 'communicate' carefully. That a sentence has the effect of communicating some information in context

does not necessitate that the semantic content of that sentence entails that information. If we use entailment-sensitive tools to probe *might*-claims, it's easy to see that the strength observed above is not semantic. Consider the following contradiction tests:

(6) a. An asteroid might wipe out all life on Earth tomorrow, but it's a truly remote possibility, and it's not worth worrying about.
b. Running the Large Hadron Collider might destroy the universe, but it's so astronomically unlikely that it's not a relevant consideration.

If the intuitive strength *might*-claims was hard-coded into their semantics, we would expect conjunction of *might*-p with an expression that p is implausible or extremely unlikely or irrelevant or not worth paying attention to result in contradiction. As the data in (6) demonstrate, this is not the case: *might*-p conjoined with a sentence dismissive of p's likelihood or relevance is perfectly coherent. This data is on its own perhaps insurmountably problematic for theories in which *might*-claims always entail something stronger than the non-impossibility of their prejacent. However, they are not necessarily problematic for a theory in which *might*-p entails that p is more likely than a contextual threshold value. Consider this denotation for *might*-p, taken from Swanson (2006) with minor notational modifications:

(7) $[\![\textit{might-}p]\!]^{M,w,g} = 1$ iff $\textit{prob}([\![p]\!]^{M,w,g}) > \alpha$
Where *prob* is a function from propositions to probabilities, and α is a contextually determined threshold.

A theory based on such a denotation for *might*-p could deal with this data by proposing that conjoining *might*-p with an assertion that p is very unlikely simply forces accommodation of a very low value for α, such that α lies beneath the cutoff point for unlikeliness, so as to avoid contradiction.

At this point, I'll introduce novel data that shows that even the flexibility provided by a contextual threshold theory of *might* cannot account for the full range of empirical facts:

(8) **A:** Paul might come to the party.
a. **B:** Yeah, he might, but it's extremely unlikely.
b. **B:** #Well, though it's not impossible that he'll come, you're wrong that he might, because it's so unlikely.

(8a) is an example of DISMISSIVE AGREEMENT: **B** agrees with the *might*-claim, but her response nonetheless feels dismissive, as it goes on to characterize the prejacent as extremely unlikely. The possibility of dismissive agreement makes the same point as the contradiction tests above.

However, crucially for the feasibility of strong theories of *might* that involve contextual thresholds, (8b) shows that the inverse of dismissive agreement is impossible: it is bizarre and contradictory to explicitly acknowledge that p is not impossible but go on to reject a *might*-claim on the basis of the implausibility of its prejacent. This should be acceptable on a threshold-based theory, because for any nonzero valuation of α it is perfectly possible for the likelihood of p to fall beneath α without p being impossible.

Dan Lassiter (p.c.) points out corpus data like the following, in which a *might*-claim is rejected in a context that suggests that what is being rejected is that its prejacent is likely, not that its prejacent is possible:

(9) Bats are very good at flying—they have to be if they want to fly around in the dark! So it's just not true that a bat might get tangled in your hair.
(retrieved from http://rtejr.rte.ie/shutterbugs/shutterbugs-blog/2014/10/bats/ on 11/6/2015)

In this example, it's not the case that there's truly no chance whatsoever that a bat will get tangled in your hair, it's just unlikely enough that it seems reasonable to dismiss the possibility. In §4 I develop a pragmatic account of strengthening that predicts that in most contexts *might*-claims will strengthen; (9) shows that *might*-claims can be rejected/dismissed on the basis of their strengthened meanings. This is a pragmatic phenomenon; the crucial *semantic* fact demonstrated by (8b) is that rejecting a *might*-claim gives rise to contradiction if its prejacent has been *explicitly* acknowledged to be possible earlier in the utterance.

I conclude on the basis of the novel data examined in this section that only a weak semantics for *might* is fully compatible with the empirical landscape.

3. The Semantics of *Might*

I assume the following weak semantics for *might*:

(10) $[\![\textit{might-}p]\!]^{M,w,g} = 1$ iff $P_w([\![p]\!]^{M,w,g}) > 0$
Where P_w is a function from propositions to degrees of epistemic likelihood at w such that for all propositions ϕ, ψ, if $\phi \subseteq \psi$ then $P_w(\phi) \leq P_w(\psi)$.[2]

(11) cf. the denotation given in Kratzer (1977) (with slight notational modifications):
$[\![\textit{might-}p]\!]^{M,w,g} = 1$ iff $\exists w \in \text{EPIST-WORLDS}_w$ s.t. $w \in [\![p]\!]^{M,w,g}$
Where EPIST-WORLDS_w is the set of worlds epistemically accessible from w.

[2]Whether the function P_w represents a finitely additive probability measure, as in Yalcin (2010), or maps to a cruder, rougher-grained scale of intuitive likelihood is irrelevant for the proposal at hand. What is crucial is that it represents epistemic likelihood, not objective likelihood.

The denotation given in (10) is a notational variant of Kratzer (1977)'s seminal account presented in (11), given the assumption that a world w' is epistemically accessible from w iff $P_w(\{w'\})$ > 0. Lassiter (2011 a.o.), Yalcin (2010 a.o.), Swanson (2006 a.o.), and Moss (2015) argue that the semantics of epistemic modals should make direct reference to probabilities; extensions of the Kratzerian theory of modality into degree-based frameworks have been developed by Klecha (2014), Grosz (2009) and Katz (2015). The account given below works equally well in probabilistic, classical Kratzerian and degree-based semantics of epistemic modals. The only crucial assumption I make is that the pragmatics is sensitive to degrees of likelihood; I remain agnostic about whether direct reference to probabilities or degrees is necessary in the semantics of epistemics. I've chosen the notation in (10) over Kratzer's simply for notational parsimony, because I'm going to be dealing with likelihood scales in the pragmatic account of strengthening below.

4. Lower Bound Implicatures

Given a weak semantics for *might*, how can we explain the intuition that many *might*-claims communicate something strong? There is a simple, intuitive reason why a claim that some proposition p is not impossible should tend to convey that p is fairly likely, or that p is a possibility worth considering: trivially small probabilities are seldom relevant, and so for an assertion that the probability of p is nonzero to be a relevant contribution to an average conversation, it must be taken to mean that the probability of p is fairly substantially above zero. In this section, I'll develop an implementation of that intuition in a formal pragmatics that is almost completely standard.

The only non-standard piece of the pragmatics that I assume comprises my novel theoretical contribution: the idea that Questions Under Discussion (QUDs—q.v. Roberts 1996 & Ginzburg 1996) come packaged with a specification of the grain size of probability that is relevant with respect to their answers, called PROBABILITY GRAINS. Given a scale of degrees of epistemic likelihood running from 0, indicating epistemic impossibility, to 1, indicating complete epistemic certainty, I define PROBABILITY GRAINS (PGs) like so:

(12) PROBABILITY GRAINS:
A Probability Grain PG_n ($n \geq 2$) is the unique tuple of n threshold values in ascending order $< t_1, \ldots, t_n >$ that comprises a uniform partition of [0,1] (the unit interval)

(13) a. A COARSE Probability Grain: $PG_4 = < 0, .33, .66, 1 >$

b. A FINE Probability Grain: $PG_{11} = < 0, .1, .2, .3, .4, .5, .6, .7, .8, .9, 1 >$

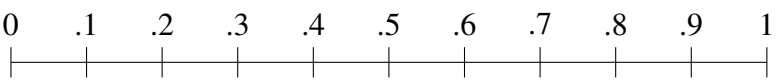

A Probability Grain partitions the epistemic likelihood scale into a set of equivalence classes; members of an equivalence class are not considered to be RELEVANTLY DIFFERENT from each other.[3]

(14) RELEVANT DIFFERENCES:
Relative to a Probability Grain $PG_n = < t_1, \ldots, t_n >$, any number $n' \in [0,1]$ is not RELEVANTLY DIFFERENT from a threshold value t_i unless $n' \leq t_{i-1}$ or $n' \geq t_{i+1}$

For any Probability Grain PG_n, $t_1 = 0$ and $t_n = 1$. As n grows higher, PG_n partitions the unit interval into more and more equivalence classes, and therefore more and more fine-grained distinctions in probability become relevant relative to the Probability Grain. A Probability Grain PG_i is COARSER than a Probability Grain PG_j iff $i < j$; if PG_i is COARSER than PG_j, then PG_j is FINER than PG_i. If a PG partitions the likelihood scale into very few equivalence classes, I will abuse terminology by referring to it as COARSE, and if a PG partitions the likelihood scale into many equivalence classes, I will abuse terminology by referring to it as FINE.

Probability Grains are a way of formally representing the fact that very fine-grained probabilistic distinctions are irrelevant to most conversations. If we're talking, for example, about whether Paul will come to the party, we're not likely to care overmuch about whether there is a 55% chance versus a 56% chance that Paul will come—we care whether he's definitely coming or definitely not coming, and we probably also care about a couple of intermediate probabilities as well: for example, we probably also care to know whether he's probably coming, or probably not coming. In the system I'm proposing, this is cashed out formally by saying that (in most default contexts) the QUD *Will Paul come to the party?* is associated with a Coarse Probability Grain.

The PG associated with a QUD is taken into account by Gricean reasoning about cooperativity when an answer to that QUD expresses a range of probabilities. Consider the dialogue in (5), repeated here as (15):

(15) **A:** Will Paul come to the party?
B: He might.

Let's assume for the sake of argument that **A**'s question is associated with the Coarse PG given in (13a). The semantic contribution of **B**'s response is simply that Paul coming to the party is not epistemically ruled out:

[3]The representations in (13) are quite similar to representations deployed by Krifka (2006) to account for (im)precision in the use of number words. For Krifka, however, such coarse- and fine-grained tuples represent coarser- and finer-grained scales, not coarser- and finer-grained ways of partitioning an underlying continuous scale, and there's no element of pragmatic enrichment via implicature to the way he deploys them.

(16)

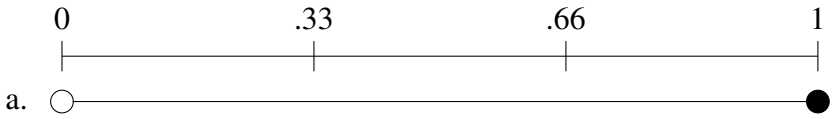

16a = semantic contribution of *might-p*

A white circle represents an exclusive bound, and a black circle represents an inclusive bound. The semantic contribution of *might-p*, given in (10) and depicted visually in (16), is merely that the epistemic probability of *might-p* is not 0. Any degree of likelihood other than 0 is compatible with the truth conditions of *might-p*. However, 0 is t_1 in the QUD's PG; a probability is only *relevantly* different than 0 if it is at least as high t_2, which, because the QUD's PG is very rough, is substantially higher than 0. This is where Gricean reasoning comes into the picture.

The guiding assumption of Gricean reasoning (Grice 1975) is that inferences beyond the literal meaning of a statement can be derived in non-adversarial conversations from the assumption that the speaker is being cooperative, with cooperativity defined (at least in part) in terms of:

(I) giving only accurate and well-substantiated information	(QUALITY)
(II) giving as much information as is necessary	(QUANTITY)
(III) presenting information parsimoniously	(MANNER)
(IV) keeping contributions relevant to the discussion at hand	(RELATION)

Gricean reasoning about the strength of *might-p* relative to a PG proceeds in the following way: the speaker's semantic contribution is simply that $P(p) \neq 0$, or, equivalently, that $P(p) \neq t_1$.[4] However, if the speaker believed that $P(p) < t_2$, her meaning would not be strong enough to be a relevant answer to the QUD; $P(p)$, though technically different from 0, would not be relevantly different. If the speaker was being cooperative, she must mean to communicate that $P(p)$ is *relevantly* larger than 0, i.e. that $P(p) \geq t_2$:

[4] I omit the w subscript from P_w here and throughout the rest of the paper—P should always be taken to represent epistemic likelihood at the world of evaluation.

(17)

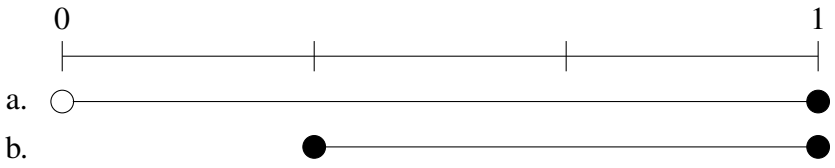

17a = semantic contribution of *might-p*
17b = pragmatically enriched lower bound

The intuitive strength of *might*-claims is an implicature generated by a conspiracy of QUANTITY and RELATION: it follows from the assumption that the speaker's meaning is strong enough to be relevant to the QUD. I'll refer to these implicatures as LOWER BOUND IMPLICATURES, as the implicature has the result of raising the lower bound of the range of probabilities the *might*-claim communicates.[5]

Gricean reasoning also gives us an explanation for why, as remarked on in the discussion of (3), *might*-claims often give rise to the inference that the speaker has access to information or evidence about the prejacent: it follows from QUALITY that an assertion that the epistemic probability of the prejacent is (substantially) nonzero must be justified by good evidence if the assertion is cooperative.

Thinking about things in these terms also gives us a clear explanation for how dismissive agreement works. The dismissive agreement example in (8a) is repeated in (18):

(18) **A:** Paul might come to the party.
B: Yeah, he might, but it's extremely unlikely.

In this example, **B** first agrees with **A**'s *might*-claim, and then goes on to dismiss the prejacent as extremely unlikely. Assuming that the range of probabilities that *extremely unlikely* denotes falls to the left of t_2 in the QUD's PG (i.e., that we're in a context where extremely unlikely possibilities are not relevant), consider the effect of agreeing with a *might*-claim while also asserting that its prejacent is extremely unlikely:

[5]Horn (1984) calls the class of implicatures that have the property of 'inducing lower-bounding implicata' R-IMPLICATURES. He does not discuss implicatures of the precise kind that I've called lower bound implicatures here, focusing instead on phenomena like *I broke a finger yesterday* implicating *I broke my own finger yesterday*. Though there are some similarities between the R-Implicatures Horn discusses and the lower bound implicatures I discuss here, there are also significant differences.

(19)

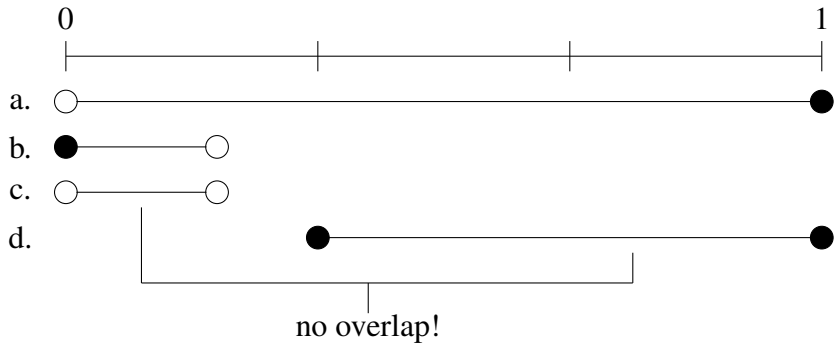

19a = semantic contribution of *might-p*
19b = semantic contribution of *extremely-unlikely-p*
19c = full semantic package (*might-p* + *extremely-unlikely-p*)
19d = relevantly nonzero portion of the scale

In dismissive agreement, an agent agrees with a *might*-claim while also asserting that its prejacent is extremely unlikely; if the portion of the likelihood scale picked out by *extremely unlikely* is a subset of the portion that is not relevantly different from 0 for the purposes of the current QUD, then dismissive agreement is pragmatically identical to asserting that the *might*-claim was an irrelevant contribution. The pragmatic account developed above gives us an explanation for why **B**'s response in (8a) appears superficially to be agreement while still feeling like a rejection of the original *might*-claim: **B**'s response is only non-contradictory if one cancels the implicature that the *might*-claim is strong enough to be relevant. **B** agrees with **A**'s statement, but goes on to (implicitly) reject the implicature that the prejacent is likely enough to be relevant.

It's worth pointing out that *might* interacts with various operators that appear to legislate its relation to the QUD's PG:

(20) a. **A:** Paul might come to the party.
B: No, that's extremely unlikely.

b. **A:** Paul {technically might, might in principle} come to the party.
B: #No, that's extremely unlikely.

c. **A:** Paul very well might come to the party.
B: #Yeah, he very well might, but it's extremely unlikely.

(20a) shows that it's possible to disagree directly with the pragmatically enriched form of a *might*-claim; the *might*-claim can be rejected because its prejacent isn't likely enough. In (20b), however, that is no longer possible: when *technically* or *in principle* is added to the *might*-claim it becomes

infelicitous to reject it on the grounds that the prejacent, though not impossible, is very unlikely. In **B**'s response in (20c) we see that dismissive agreement is rendered infelicitous if the dismissive agreer adds *very well* to a *might*-claim before going on to dismiss it.

My interpretation of these facts, informally, is as follows: *technically* and *in principle* preclude the generation of lower bound implicatures. They signal that the *might*-claim should not necessarily be taken to communicate RELEVANTLY nonzero probability. **B**'s response in (20b) is infelicitous because **A** indicated that a lower bound implicature should not be generated. However, *very well* does the opposite: it strenthens a *might*-claim's communication of RELEVANTLY nonzero probability from an implicature to an entailment. *very-well-might-*p entails that $P(p) \geq t_2$ in the QUD's PG.

I'll call such operators RELEVANCE OPERATORS because of the way they appear to affect the status of the assumption that an assertion is strong enough to be relevant to the current QUD (either calling it off, or strengthening it into an entailment). A formal theoretical model of the semantico-pragmatics of such operators lies far outside the scope of this paper, but strikes me as a very exciting avenue for future work.

5. More About Probability Grains

The discussion of strengthening implicatures in the previous section used a Coarse Probability Grain to show how a weak semantics can implicate a stronger interpretation if small distinctions in probability are irrelevant to the conversation. Although in many conversations such small distinctions are irrelevant, there are conversations in which participants care quite a bit about very small distinctions in probability; my account predicts that in such situations *might*-claims will tend to be interpreted more weakly. One example of such a conversation would be a conversation about particle physics among a group of expert scientists. Because very fine-grained differences in probability could matter a great deal in such a conversation, we would expect some QUDs arising in the course of the conversation to be associated with quite Fine Probability Grains, relative to which only small amounts of strengthening will occur. The prediction that my theory makes about such a conversation is that *might*-claims would generally not be taken by participants to communicate substantially nonzero probability, because very small probabilities are not irrelevant to the conversation. That prediction accords with my intuition.

One crucial distinction between a conversation about who is going to come to a party and a conversation about how subatomic particles interact is the fine-grainedness of the probabilistic information available in principle about each question. It is difficult to see how one would obtain information that would differentiate between a 55% and a 56% chance that someone will attend a party; however, such information is obtainable about many physical interactions. I assume that the fine-grainedness of probabilistic information obtainable in principle about the answers to a question acts as an upper limit on the Fineness of the PG associated with that question, though sometimes what is relevant to a QUD may be Coarser than the fineness of probabilitic information

available in principle.

It should be noted that the full machinery of Probability Grains is not necessary to generate the strengthening effects I've used them to model. A system in which a QUD specifies a minimum threshold that probabilities must reach before they become relevant would accomplish the same effects for *might*-claims. Such an instantiation would be a very simple variant on the semantic threshold account endorsed by Swanson (2006) and Lassiter (2011); the only difference would be locating the threshold in the pragmatics instead of in the semantics of *might*. I've chosen to present the PG system above instead of a simpler threshold-based formulation because the PG system allows us to make principled predictions about which contexts will be the most likely to provoke the most strengthening of *might*-claims, and because it makes principled predictions about interactions with the upper end of the scale as well—namely that contexts in which we expect *might*-claims to be the strongest should also be contexts in which the most skepticism is expressed by an assertion that a proposition p is not certain.

In most cases, PGs are implicit—it is rarely explicitly stated that a probability is only relevantly different from 0 if it is at least .05, for instance. For this reason, we would expect that participants in a conversation will interpret *might*-claims not as strengthening to a particular degree of likelihood as their lower bound; instead, we would expect strengthening to a somewhat vague and nebulous value, in view of listeners' uncertainty about the PG their interlocutors are assuming. However, this is not the case for all conversations. As an example of a QUD accompanied by an explicit relevance threshold, consider the following:

(21) *Context: **A** is teaching a probability class, working through a story problem about stocks. **B** is her student.*
A: Which stocks have at least a 5% chance of rising today?
B: Apple, Facebook and Google stock all might rise in value today.

A's question makes explicit that she is only interested in stocks with at least a 5% chance of rising; **B**'s *might*-claim in this context communicates (defeasibly) that there is at least a 5% chance that Apple, Facebook and Google stock will rise, which is exactly what my theory predicts.

6. The Typology of Implicatures

In this paper I've introduced the novel empirical paradigm of dismissive agreement. In this section, I'll explore that paradigm a little more deeply, and use it to identify differences in the behavior of lower bound implicatures and scalar implicatures. Consider the following facts:

(22) **A:** Paul might come to the party.
*Context: **B** believes that the prejacent is possible but very unlikely.*
   a. **B:** Yeah, he might, but it's extremely unlikely.
   b. **B:** No, that's extremely unlikely.
   c. **B:** #No, you're wrong that he might come, because it's extremely unlikely.

The case of dismissive agreement in (8a) is repeated in (22a). (22b) repeats the observation in (20a) that it is possible to disagree directly with the pragmatically enriched meaning of the *might*-claim; **B** rejects **A**'s assertion on the basis of the prejacent not being relevantly likely. However, (22c) shows that in the same context, disagreement is infelicitous if it overtly targets the *might*-claim. If **B** believes the prejacent to be possible, just unlikely, she can't explicitly target the *might*-claim for disagreement in order to reject only its pragmatically enriched meaning.

This is somewhat surprising, because it is well known that negation can be used metalinguistically with some kinds of implicatures to reject only the implicated content, without rejecting the literal meaning of the expression. For example, consider the following scalar implicatures plugged into the paradigm above:

(23) **A:** Paul might come to the party.
   a. **B:** Yeah, he might come—in fact, he'll definitely come.
   b. **B:** No, he'll definitely come.
   c. **B:** No, you're wrong that he might come—he'll definitely come.

Existential meanings tend to implicate the negation of related universal meanings (Horn 1972, Gazdar 1979). As a special case of such scalar implicatures, *might*-claims tend to implicate that the prejacent is not definitely true. (23) demonstrates how these implicatures pull apart from lower bound implicatures in terms of their interaction with disagreement. Scalar implicatures behave the same as lower bound implicatures in terms of dismissive agreement (23a)—it is coherent to agree with the *might*-claim before going on to reject the scalar implicature. (23b) shows that it is possible to disagree directly with the implicated content—**B** doesn't disagree with the semantic contribution of **A**'s utterance, she disagrees only with its pragmatically enriched meaning. (23c) is where the two types of implicatures pull apart: unlike in (22c) we see here that the *might*-claim itself can be targeted for disagreement when what is being rejected is not the semantics of the *might*-claim, but its scalar implicature.

It appears that the ability to target only the implicature with metalinguistic negation is not a general property of implicated content. Why would we only see metalinguistic negation with scalar implicatures, not lower bound implicatures?

I believe that the explanation for this distinction can be traced to a difference in what drives the computation of each type of implicature. Since Horn (1972), scalar implicatures have been understood to be triggered by the presence of a scalar element: *might* is on a scale with *definitely*, and its presence in a sentence implicates the negation of a sentence in which it has been replaced with its stronger scalemate. Lower bound implicatures, as discussed above, are not triggered by the mere fact of the presence of *might* in the sentence—they result from the evaluation of the semantics of the full sentence relative to the QUD. To sloganeer: scalar implicatures are LEXICAL, while lower bound implicatures are CONTEXTUAL.

This distinction explains the metalinguistic negation asymmetry above if we assume that metalinguistic negation targets some aspect of the *form* of an utterance. Because scalar implicatures are lexical, it makes sense that the use of the word *might* can be targeted for metalinguistic negation when what is being rejected is only the implicature—after all, it was the use of the word *might* that gave rise to the implicature. This explains the fact that (23c) sounds best with heavy emphasis on *might*, which serves to highlight which aspect of the form of the utterance is being targeted by metalinguistic negation. However, because lower bound implicatures are about the interaction between sentence meanings and QUDs, it is not the form of the utterance that gave rise to the implicature, and so it doesn't make sense to target the *might*-claim with metalinguistic negation.

Targetability by metalinguistic negation is not the only difference between lower bound implicatures and scalar implicatures. They also respond differently to focus. It has been widely noted that scalar implicatures are foregrounded or strengthened when the existential element bears focus. However, lower bound implicatures are not foregrounded or strengthened when the existential element bears focus:

(24) Paul MIGHT come to the party.
   a. *Strongly implicates: It is not certain that Paul will come to the party.*
   b. *Does not strongly implicate: There is a large chance that Paul will come to the party.*

This may also be traceable to the fact that scalar implicatures are triggered by the lexical item, but lower bound implicatures are not: maybe focusing the existential foregrounds the scalar implicature by drawing attention to the fact that the existential was chosen instead of one of its stronger scalemates.

I leave a fuller investigation of the empirical facts about the interaction between lower bound implicatures and focus to future work.

7. Extension To Other Existentials

Above, I've described lower bound implicatures as arising from an interaction between the weak semantics of *might* and the notion that very small probabilities are usually not relevant to QUDs.

That the formalization pursued above is instantiated in terms of degrees of likelihood may have suggested that lower bound implicatures are specific to epistemic claims that can be construed as making reference to probability. However, the same crucial phenomena occur for other existentials. In this section, I will focus exclusively on *some*. Consider the following sentence:

(25) Paul read some of the assigned article (but didn't finish it).

The literal meaning of this sentence is quite weak: there is some portion of the assigned article (perhaps trivially small, like two sentences) that Paul read. However, just like with *might*-claims, what gets communicated is somewhat stronger. In a normal context, this sentence communicates that Paul read a relevantly large portion of the assigned article; perhaps the introduction. The line of reasoning is the same as the reasoning detailed above for *might*-claims: trivially small portions of the article are not relevant to the discussion at hand. It seems sensible to assume that it is not cooperative to respond to a QUD like *Who read some of the assigned article?* by pointing out that Paul read the first two sentences.

Strengthening inferences for *some* behave just like strengthening inferences for *might*:

(26) **A:** Paul read some of the assigned article.
   a. **B:** Yeah, but he only read two sentences.
   b. **B:** No, he only read two sentences.
   c. **B:** #It's false that Paul read some of the article, because he only read two sentences.

Dismissive agreement (26a), disagreement with the strengthening implicature (26b), and the infelicity of targeting the existential claim for disagreement while acknowledging that a nonzero amount of cake was consumed (26c) all support the hypothesis that the inference that Paul read a relevantly large portion of the article is a lower bound implicature.

Though the formalism proposed above for lower bound implicatures with *might* makes specific reference to degrees of likelihood, it can be treated as a formula for explaining lower bound implicatures with a broader variety of existentials. The necessary machinery to explain the lower bound implicature that (25) gives rise to is quite comparable to the machinery necessary to explain lower bound implicatures with *might*-claims: assume that QUDs specify the grain size of quantity that is relevant to their answers; the *some*-claim will pragmatically strengthen such that the quantity communicated is relevantly different from 0.

I leave a fuller investigation of lower bound implicatures with the complete range of existential operators (as well as an investigation of the prospects for a unified account of such implicatures for different kinds of existentials) to future work.

8. Conclusion

I've made two major arguments in the course of this paper. The first is an empirical argument about the semantics of *might*. I've argued on the basis of contradiction tests and the dismissive agreement paradigm that *might* is semantically weak. The fact that many theorists of epistemic modality have proposed a strong semantics for *might* can be taken as a methodological parable: it is well known that expressions of natural language can communicate more than their literal semantic meaning; therefore, one must make sure that an aspect of the meaning of an expression in context is truly semantic before hard-wiring that aspect of meaning into the denotation of the expression.

The second argument I've made here is a theoretical argument about the pragmatics of *might*-claims. I've argued that the apparent strength of *might*-claims falls out of a standard formal pragmatics enriched with specifications of the relevant grain size of probability relative to the QUD. If very small distinctions in probability are irrelevant to the QUD, then listeners who assume their interlocutors are being cooperative will strengthen their interpretation of the lower bound communicated by a *might*-claim in accordance with the assumption that the *might*-claim contributes relevant information. It's not surprising that existential claims would strengthen in the way described in this paper: existential meanings can be extremely weak, and extremely weak claims are rarely relevant. I hope that the strengthening mechanism explored here for *might*-claims will become a special case of a more general formula for using the standard techniques of Gricean pragmatics to derive stronger, more relevant meanings from weak claims.

References

Gazdar, G. (1979). *Pragmatics: Implicature, Presupposition, and Logical Form*. Academic Press.

Ginzburg, J. (1996). Dynamics and the semantics of dialogue. In J. Seligman and D. Westerståhl (Eds.), *Language, Logic, and Computation*, Volume 1, pp. 1–16. CSLI.

Grice, P. (1975). Logic and conversation. In P. Cole and J. Morgan (Eds.), *Syntax and Semantics 3: Speech Acts*, pp. 41–58. Academic Press.

Grosz, P. (2009). Grading modality: A new approach to modal concord and its relatives. In *Proceedings of Sinn und Bedeutung 14*, pp. 185–201.

Horn, L. (1972). *On the semantic properties of logical operators in English*. Ph. D. thesis, UCLA.

Horn, L. (1984). Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In D. Schiffrin (Ed.), *Meaning, Form, and Use in Context: Linguistic Applications*, pp. 11–42. Georgetown University Press.

Katz, G. (2015). On the origins of metric scales for gradable modal expressions. Poster presented at *Sinn und Bedeutung 20*.

Klecha, P. (2014). *Bridging the Divide: Scalarity and Modality*. Ph. D. thesis, University of Chicago.

Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and Philosophy 1*, 337–355.

Kratzer, A. (1981). The notional category of modality. In *Words, worlds and contexts: New*

approaches in world semantics, pp. 38–74. De Gruyter.

Krifka, M. (2006). Approximate interpretations of number words: A case for strategic communication. In G. Bouma, I. Krämer, and J. Zwarts (Eds.), *Cognitive foundations of interpretations*, pp. 111–126. Royal Netherlands Academy of Arts and Sciences.

Lassiter, D. (2011). *Measurement and modality: The scalar basis of modal semantics*. Ph. D. thesis, NYU.

Moss, S. (2015). On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics 8*, 1–81.

Roberts, C. (1996). Information structure in discourse. In J. Yoon and A. Kathol (Eds.), *OSU Working Papers in Linguistics 49: Papers in Semantics*, pp. 1–53. Ohio State University.

Swanson, E. (2006). *Interactions with context*. Ph. D. thesis, MIT.

Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic 25*(3), 221–261.

Willer, M. (2013). Dynamics of epistemic modality. *Philosophical Review 122*(1), 45–92.

Yalcin, S. (2010). Probability operators. *Philosophy Compass 5*(11), 916–937.