Obligatory irrelevance and the computation of ignorance inferences*

Brian Buccola
The Hebrew University of Jerusalem
brian.buccola@gmail.com

Andreas Haida
The Hebrew University of Jerusalem
andreas.haida@gmail.com

Abstract

The standard grammatical theory of scalar implicature, as envisioned by Chierchia (2004), Fox (2007), and Chierchia, Fox, and Spector (2012), posits that scalar implicatures are derived in grammar, as a matter semantics, rather than pragmatically, as an implicature rooted in Grice's maxim of quantity. Ignorance inferences, by contrast, e.g. those associated with plain disjunctive sentences, are derived pragmatically, as quantity implicatures. More generally, the standard theory predicts that for any utterance S and any relevant proposition ϕ which isn't entailed, and whose negation isn't entailed, by S, S gives rise to an inference of speaker ignorance about ϕ. We argue that this prediction is wrong: it fails to explain the contrast in ignorance inferences associated with *at least* (which obligatorily implies ignorance) vs. *more than* (which doesn't) (Geurts and Nouwen 2007; Nouwen 2010, 2015). The problem is that, without stipulating restrictions on which propositions are relevant, the theory overgenerates ignorance inferences across the board. We argue that the solution is to close relevance under belief (if ϕ is relevant, then it's also relevant whether the speaker believes ϕ). This move has the effect that ignorance inferences, like scalar implicatures, can only be derived in grammar, via a covert belief operator of the sort proposed by Meyer (2013) and discussed further by Fox (2016). The maxim of quantity, we show, then no longer enriches the meaning of an utterance, per se, but rather acts as a filter on what can be relevant in an utterance context. In particular, certain alternatives (of certain utterances) are shown to be incapable of being relevant in any context where the maxim of quantity is active — a property we dub *obligatory irrelevance*. We argue that obligatory irrelevance provides the key to understanding the contrast in ignorance inferences exhibited by *at least* vs. *more than*. We also argue that translating our proposal into neo-Gricean terms, if at all possible, would yield a conceptually less appealing and empirically less adequate theory.

Keywords: ignorance inferences, scalar implicatures, exhaustivity, Gricean maxims, modified numerals, universal density of measurement

1 Introduction

The last few decades of research in semantics and pragmatics have brought forth a great debate concerning the status of scalar implicatures, namely whether they're derived on the basis of pragmatic reasoning or rather in grammar, as a matter of semantics proper. The debate revolves chiefly around the nature of Grice's maxim of quantity in confrontation with the so-called *symmetry problem*. Roughly speaking, the symmetry problem is that one cannot both maintain a basic, virtually truistic version of the maxim of quantity and also derive scalar implicatures pragmatically: given an utterance S, for any relevant alternative S' of S that one would like to pragmatically derive a scalar inference about (namely, *not S'*), S' has a so-called symmetric partner (*S and not S'*), presumably also relevant, which preempts deriving a scalar inference about S'; one cannot derive a scalar inference about both alternatives without contradicting S, and so instead, one predicts speaker ignorance about each alternative. This problem has led to two broad strands of research, which solve the problem in different ways. The neo-Gricean approach rejects the basic formulation of the maxim of quantity, replacing it with a version that makes reference to formally defined alternatives, e.g. Horn scales (Horn 1972) or structurally defined alternatives (Katzir 2007), so that *S and not S'* is no longer an alternative of S, thus breaking symmetry. The grammatical approach maintains the basic maxim of quantity by concluding that scalar implicatures are computed in grammar, with a covert grammatical exhaustification device, *exh*, which (much like the non-basic maxim of quantity on the neo-Gricean approach) is sensitive to formal alternatives.

Focusing just on the standard grammatical approach as developed by Chierchia (2004), Fox (2007), Chierchia, Fox, and Spector (2012), and Fox (2014), while scalar implicatures are only derived in grammar, ignorance inferences, e.g. those associated with plain (unembedded) disjunctive sentences, are predicted to arise only pragmatically. Not only that, but ignorance inferences are expected to be quite pervasive, for the following reason: given an utterance S, if S' is a relevant alternative of S such that S doesn't entail S' nor the negation of S' (that is, S doesn't *settle S'*), then the maxim of quantity licenses the inference that the speaker doesn't have the belief that S' is true; but on standard assumptions about relevance, *not S'* is also relevant, and since S doesn't settle S', S doesn't settle *not S'* either, so S also licenses the inference that the speaker doesn't have the belief that *not S'* is true (i.e. that S' is false). In other words, the speaker doesn't have a belief one way or the other about the truth of S': she's ignorant about S'. The result is that, for any relevant alternative S' of an utterance S which S doesn't settle, the maxim of quantity licenses an inference of speaker ignorance about S'. This result is essentially a variant of the symmetry problem, though it's not as clear that it's a

problem in this context; without evidence to the contrary, it doesn't seem unnatural or undesirable.

However, a consequence of this theory is the following: two semantically equivalent sentences ought to give rise to exactly the same ignorance inferences, since they entail (hence settle) all the same propositions. In the wake of work by Geurts and Nouwen (2007), Büring (2008), Nouwen (2010), Mayr (2013), and Schwarz (2016), there's quite compelling evidence that this prediction is wrong: numerals modified by *more than* and *at least* appear to be semantically interdefinable (*more than n* seems semantically equivalent to *at least n + 1*), yet they give rise to distinct ignorance inferences. Specifically, sentences with (unembedded) *at least n* obligatorily imply speaker ignorance about the exact number in question, while sentences with *more than n* don't. More generally, there's a class of expressions (dubbed class B by Nouwen (2010)) that obligatorily imply ignorance (when unembedded), and another class (class A) that don't. Prima facie, this class A/B distinction is unexpected under the grammatical view sketched above. Or, put more positively, this paradigm provides the perfect testing ground for formally developing and, when necessary, amending this view.

Our goal is threefold: to develop a formally precise version of the grammatical account that successfully captures the class A/B distinction, to argue that it provides important insight into this distinction and into the nature of the computation of ignorance inferences more generally, and to show that translating our proposal into neo-Gricean terms, if at all possible, leads to a conceptually less appealing and empirically less adequate theory.

We begin by laying out the empirical data we wish to capture (§2). We then develop a formally explicit semantic-pragmatic framework (§3) that has as a consequence the generalization that for any relevant proposition ϕ that an uttered sentence S doesn't settle, the maxim of quantity licenses the inference that the speaker is ignorant about ϕ (§4). (As far as we know, ours is the first formally explicit theory of this kind.)

We then show why this account fails to predict the class A/B distinction (§5). In short, because the set of relevant propositions is assumed to be closed under negation, the theory overgenerates ignorance inferences across the board, for the reasons already sketched above. Intuitively, the solution is to somehow restrict what propositions are relevant, but on first glance, the theory appears to provide no such way beyond pure stipulation about what's relevant.

We then slightly extend the theory (§6): relevance is now closed under belief (if ϕ is relevant, then it's also relevant whether the speaker believes ϕ), an amendment already motivated on independent grounds by Fox (2016). The effect of this amendment is that ignorance inferences (like scalar implicatures) can now only be derived in grammar, by way of a covert belief operator (Chierchia 2006; Meyer

[2013](#)), not pragmatically. We show that, as a result, the maxim of quantity no longer derives any additional inferences, i.e. doesn't enrich the meaning of an uttered sentence, per se, but rather acts as a filter on which propositions can be relevant in an utterance context. More specifically, closing relevance under belief induces a precise condition on when the maxim of quantity can be active: the maxim of quantity can be active for the interpretation of an uttered sentence S only if, for every relevant proposition ϕ, S either settles or entails ignorance about ϕ; otherwise, the maxim of quantity yields a contradiction. Thus, if the maxim of quantity is taken to be active for the interpretation of an uttered sentence S, then for any proposition ϕ that S doesn't settle or entail ignorance about (regardless of what other propositions are relevant), ϕ can't be relevant. We refer to this as *obligatory irrelevance*.

What we ultimately derive is that the alternative propositions that previously led to too many ignorance inferences are now obligatorily irrelevant (given reasonable assumptions about the formal alternatives that feed exhaustification), while those that aren't (those that can be relevant) yield exactly the attested ignorance inferences (§7).

We end by discussing a couple of issues our theory faces (and how to deal with them), a number of correct predictions our theory makes, and some advantages our theory has over a neo-Gricean approach to explaining the full range of data that we examine (§8).

2 Empirical background

A sentence like (1) is normally understood to convey that Ann owns *exactly* two dogs. The classic view of how (1) acquires this meaning, due to Horn ([1972](#)), states that (1) literally means that Ann owns two or more dogs (a 'one-sided' meaning, notated henceforth as "$[\geq 2]$"), and implicates that it's not the case that Ann owns three or more dogs ($\neg[\geq 3]$), which together mean that Ann owns exactly two dogs (a 'two-sided' meaning, $[= 2]$).

(1) Ann owns two dogs. $\hspace{4cm} (\rightsquigarrow \neg[\geq 3])$

By contrast, as Krifka ([1999](#)) points out, (2) and (3) don't have two-sided meanings, even though, on a naive view of how the upper-bound implicature for (1) arises, one might expect them to: (2) doesn't implicate that it's not the case that Ann owns at least three dogs (hence that she owns exactly two dogs), nor does (3) implicate that it's not the case that Ann owns more than three dogs (hence that she owns exactly three dogs).

(2) Ann owns at least two dogs. $\hspace{3cm} (\not\rightsquigarrow \neg[\geq 3])$

(3)   Ann owns more than two dogs. $(\not\rightsquigarrow \neg[> 3])$

In addition, while (2) obligatorily entails speaker ignorance, (3) does not (Nouwen 2010, 2015). Specifically, (2) implies that the speaker is ignorant about whether Ann owns exactly two dogs, and ignorant about whether Ann owns more than two dogs (Büring 2008; Schwarz 2016). To see this more clearly, consider the contrast in (4): (4a) is odd because it implies that the speaker doesn't know how many dogs she herself owns (a very unlikely epistemic state), while (4b) implies no such thing.

(4)   a.   #I own at least two dogs.

       b.   I own more than two dogs.

Not only do sentences like (3) with *more than two* not obligatorily imply ignorance—we'd like to argue that they *cannot* be used to imply ignorance. As evidence for our claim, consider the dialog in (5). A's final utterance A_1 is perfectly felicitous, while A_2 is not, and this can be explained on the basis that B's response prompts A to draw the inference that B is ignorant about how many dogs Ann owns: A_1 makes explicit reference to this inference, while A_2 conflicts with it.

(5)   A.   How many dogs does Ann own?

       B.   Ann owns at least two dogs.

       A_1.   How do you not know? You're her roommate!

       A_2.   #Why won't you say exactly how many?

By contrast, the same dialog with *more than*, (6), behaves in the opposite fashion. In this case, A's final response A_2 is perfectly felicitous, while A_1 is not.[1] Here, B's response provides only a partial answer to A's question, without signaling why a complete, or more informative, answer cannot be provided, and in particular without signaling ignorance. It feels as though B is dodging the question, or purposefully being under-informative. For this reason, A can explicitly ask B why B won't answer more informatively, as in A_2. A_1, on the other hand, indicates that A drew an ignorance inference from B's response, but since no such an inference was licensed, A_2 is infelicitous.[2]

---

1 With the right prosody, the response A_1 might be able to convey ignorance, but this would be due to the combination of an incomplete answer plus special prosody. Our point here is that with normal, falling intonation, a response with *at least* implies ignorance, while a response with *more than* doesn't.
2 Note in this regard that a response by B that explicitly conveys ignorance, such as *Ann owns more than two dogs, but I'm not sure exactly how many*, or *I don't know, but she owns more than two dogs*, or *Well, Ann owns more than two dogs*, with special prosody (cf. fn. 1), rescues the dialog with A_1.

(6) A. How many dogs does Ann own?

B. Ann owns more than two dogs.

A_1. #How do you not know? You're her roommate!

A_2. Why won't you say exactly how many?

To be clear, (3) is certainly *compatible* with speaker ignorance; we're simply claiming that it cannot *imply* ignorance.[3]

Finally, when any of these expressions is embedded under a universal quantifier, e.g. a universal modal (notated here by "□"), the resulting sentence has a reading which involves (what we may describe as) a matrix-level implicature and which doesn't imply any ignorance (Fox and Hackl 2006; Geurts and Nouwen 2007; Mayr 2013). For instance, (7a) and (7b) each have a reading that conveys that Bill must take two or more courses, but need not take more than two courses (i.e. he's allowed to take just two). On this, so-called "authoritative reading" (Büring 2008), the speaker is not ignorant in any way about Bill's requirements. Similarly, (7c) has a reading that conveys that Bill must take more than two courses, but need not take more than three courses (i.e. he's allowed to take just three).

(7) a. Bill is required to take two courses. $(\rightsquigarrow \neg\square[\geq 3])$

b. Bill is required to take at least two courses. $(\rightsquigarrow \neg\square[\geq 3])$

c. Bill is required to take more than two courses. $(\rightsquigarrow \neg\square[> 3])$

We won't have anything new to say about these embedded cases. (Our ultimate theory will predict these inferences for roughly the same reasons that the theories that precede ours predict them.) The reason we introduce them here is to bolster the intuition that *two*, *more than two*, and *at least two* can, and do, give rise to scalar implicatures (hence trigger alternatives), even though we don't observe any for unembedded uses of modified numerals like in (2) and (3).

The observations we wish to explain are summarized in Table 1. (We omit the embedding data because all three expressions behave the same way, and because, as we'll point out, our theory describes these facts as well.) Parts of this paradigm have been successfully explained already. Specifically, Fox and Hackl (2006) provide an analysis of the scalar implicature contrast between *two* and *more than two*, while Mayr (2013), Kennedy (2013, 2015), and Schwarz (2016), building on insights by

---

3 We're not aware that anyone has made this stronger claim before. However, Cummins, Sauerland, and Solt (2012) do observe that a sentence like *More than 100 people got married today* doesn't have the "weak implicature" that the speaker considers it possible that exactly 101 people got married today. Hence, it doesn't imply speaker ignorance about whether exactly 101 people got married today. However, this is a weaker observation than what we observe here; namely, we claim that such a sentence also fails to implicate that the speaker is ignorant about whether more than 101 people got married today. In short, such a sentence implies no ignorance (hence, nor any possibility inference) at all.

	scalar implicature	ignorance inference
two	✓	✗
at least two	✗	✓
more than two	✗	✗

Table 1: Summary of facts for (unembedded) bare numerals, comparative modified numerals, and superlative modified numerals.

Büring (2008), develop an explanation of the ignorance inference contrast between *at least two* and *more than two*.[4] The goal of this article is to ask, and answer, the following three questions:

1. Can a unified theory be developed that captures the full paradigm in Table 1?

2. Do we learn anything new that we didn't know before?

3. Does the investigation add anything to the debate concerning the pragmatic vs. grammatical status of scalar implicatures?

We answer all three questions affirmatively. Specifically, we develop a minimal extension of the standard grammatical theory of scalar implicature that captures the full paradigm in Table 1. The minimal extension is that relevance is closed under belief (if ϕ is relevant, then it's also relevant whether the speaker believes ϕ), which has the effect that not only scalar implicatures, but also ignorance inferences are derived only in grammar, as a part of semantics proper. The new thing we learn is that this theory leads to a principled restriction on what propositions can be relevant in an utterance context. We finally argue that translating our proposal into neo-Gricean terms, if possible at all, would yield a conceptually less appealing and empirically less adequate theory.

3   The standard grammatical theory of scalar implicature

We now lay out the ingredients that form the basis of our theory (in §6, we'll extend it in a minimal way), and we illustrate how it works for a textbook example involving disjunction.

---

4 Our ultimate account of *at least* and *at most* will have much in common with these authors. For two quite different approaches to the analysis of superlative modifiers, see Coppock and Brochhagen 2013; Cohen and Krifka 2014.

3.1 The ingredients

We assume a grammatical view of scalar implicature, in which a sentence S may be parsed as *exh S*, with a grammaticalized exhaustivity operator, *exh*. What *exh S* means is that S is true and that every *innocently excludable* (IE) member of alt(S), the set of relevant formal alternatives of S, is false. Intuitively, a proposition $q \in$ alt(S) is IE (given $[\![S]\!]$) just in case the negation of q doesn't contradict $[\![S]\!]$ and doesn't force any disjunction of members of alt(S) to be true, unless that disjunction is already entailed by $[\![S]\!]$. The formal semantics of *exh* is provided in (8).

(8) **Definition of *exh***   (Chierchia, Fox, and Spector 2012)

     a. $[\![\text{exh } S]\!] = [\![\text{exh}]\!](\text{alt}(S))([\![S]\!])$.

     b. $[\![\text{exh}]\!](A)(p) = p \wedge \bigwedge\{\neg q : q \in \text{IE}(p, A)\}$.

     c. $\text{IE}(p, A) =$
                $\bigcap\{M : M$ is a maximal set of excludable alternatives of A given $p\}$.

     d. $M \subseteq A$ is a set of excludable alternatives of A given p iff $p \wedge \bigwedge\{\neg q : q \in M\}$ is consistent.

     e. M is a maximal set of excludable alternatives of A given p iff M is a set of excludable alternatives of A given p, and there is no M' such that $M' \subseteq A$, $M \subset M'$, and M' is a set of excludable alternatives of A given p.

(9) **Exhaustification**
exh may attach to any propositional constituent (i.e. any node of type st).

     For any context c, the function alt_c, applied to a sentence S, returns the set of all propositions that are both relevant in c and denoted by formally derivable alternatives (in c) of S. More precisely, let $\text{alt}_c^{FA}(S)$ be the set of formally derivable alternatives (in c) of S (e.g. as proposed in Katzir 2007), and let *relevant$_c$* be the set of all relevant propositions in c. Then:

(10) **Alternatives of a sentence (as used by *exh*)**
     For any sentence S, $\text{alt}_c(S) = \{[\![S']\!] : S' \in \text{alt}_c^{FA}(S)\} \cap \text{relevant}_c$.

We almost always omit the subscript c, like in the definition of *exh* in (8).
     In addition, we assume the following *maxim of quality* and basic *maxim of quantity*. This version of the maxim of quantity is "basic" in the sense that it doesn't refer to

formal properties of linguistic expressions, e.g. Horn scales (cf. Fox 2007, 2016).[5, 6] Fox (2016) refers to the two maxims together as "virtual truisms", because we seem to be aware of them without explicitly being taught them.[7]

(11) **Maxim of Quality**
   The speaker should only utter sentences that the speaker believes to be true. ("Tell nothing but the truth.")

(12) **Maxim of Quantity**
   The speaker should utter a sentence S such that, for every relevant proposition ϕ that the speaker believes to be true, S entails ϕ. ("Tell the whole truth.")

We make the assumptions in (13) about relevance, which the maxim of quantity makes reference to. Following standard procedure, we will (at first) simply stipulate which propositions, together with the closure conditions, generate the full set of relevant propositions. (Later, we'll derive on principled grounds that certain propositions can't be relevant in a given utterance context.)

(13) **Closure conditions of relevance**   (cf. Fox 2007, endnote 5)
   a.   If ϕ and ψ are both relevant, then so is $\phi \wedge \psi$.
   b.   If ϕ is relevant, then so is $\neg\phi$.

Quantity reasoning will (on the first incarnation of our theory) be tied to the derivation of ignorance inferences. We characterize ignorance about a proposition ϕ as in (14).

(14) **Characterization of ignorance**
   To be ignorant about a proposition ϕ means to not have the belief that ϕ and to not have the belief that $\neg\phi$. (To not have the belief that $\neg\phi$ is to consider ϕ possible.)

---

5 This formulation of the maxim of quantity departs slightly from the one in Fox 2007, which only states what a speaker should do when confronted with two relevant alternatives. (If S_1 and S_2 are both relevant and S_1 is more informative than S_2, then, if the speaker believes that both are true, the speaker should utter S_1 rather than S_2.) We think that our formulation is a faithful generalization of this version to the full set of relevant alternatives and that it corresponds even more closely to the "oath" (cf. fn. 7). Indeed, it's very close to what Fox (2014, p. 3) takes to be "the simplest possible formulation [of the maxim of quantity], namely that it requires a speaker to convey all the relevant information that she has available."

6 As stated, this formulation of the maxim of quantity doesn't require the uttered sentence to be relevant itself. For instance, in principle it could be a sentence that entails everything relevant that the speaker believes, as well as some irrelevant things that the speaker believes. We think that such a sentence would be independently ruled out by the *maxim of relation*, "Be relevant". Nevertheless, one could also slightly revise our definition of the maxim of quantity to say: "The speaker should utter a sentence S such that S is relevant and ... "

7 Fox (2016) points out that the oath one takes in court ("I swear to tell the truth, the whole truth, and nothing but the truth") is essentially a reaffirmation of these virtual truisms.

Draft: July 31, 2017

We adopt an epistemic logic containing the modal operator K, where $K\phi$ means that the speaker believes the proposition ϕ. Speaker ignorance about ϕ therefore amounts to the conjunction of $\neg K\phi$ and $\neg K\neg \phi$. We sometimes write "$I\phi$" to mean $\neg K\phi \wedge \neg K\neg \phi$.

Formally, we assume an S5 modal logic, which contains the axiom schemata in (15). The importance of each of these axioms will become clear as the development of our theory unfolds.[8,9]

(15)   **S5 epistemic modal logic**

  a. $(K\phi \wedge K(\phi \to \psi)) \to K\psi$ *distribution axiom* (**K**)

  b. $K\phi \to \phi$ *truth axiom* (**T**)

  c. $K\phi \to KK\phi$ *positive introspection axiom* (**4**)

  d. $\neg K\phi \to K\neg K\phi$ *negative introspection axiom* (**5**)

Finally, following Fox and Hackl 2006, we assume that the scales needed for natural language semantics are universally dense. This hypothesis isn't necessary for the example that follows, but it's a core ingredient of our overall theory, so for this reason we introduce it immediately.

(16)   **Universal density of measurement (UDM)**   (Fox and Hackl 2006)
       Measurement scales needed for natural language semantics are always dense.

3.2   Example: Disjunction

To illustrate, let's take the classic case of disjunction, as in (17). Intuitively, (17) licenses both a scalar inference, that Ann doesn't own both a Boxer and a Collie, and a pair of ignorance inferences, that the speaker is uncertain whether Ann owns a Boxer and uncertain whether Ann owns a Collie.

(17)   Ann owns a Boxer or a Collie.
       LF: exh [$_S$ Ann owns a Boxer or (Ann owns) a Collie]

What we wish to show is that parsing (17) with *exh* leads to just the right scalar and ignorance inferences. For the sake of clarity, when we write "(17)", we refer to the LF provided below (17), containing *exh*, and we write "S" to refer to the complement of *exh*, also known as the *prejacent*. The meaning of S is thus $b \vee c$,

---

8 In the context of the initial theory we develop (§3), we take it that this logic appropriately models the way quantity reasoning works, while for our slightly amended theory (§6), which is radically grammatical, we take it that this logic appropriately models that component of the mind that governs grammar and semantics. We don't necessarily take this logic to be an appropriate model of human belief or knowledge more generally.

9 Regarding the **K** axiom, we assume that speakers know certain basic logical validities like $(\phi \wedge \psi) \to \phi$. This means that from $K(\phi \wedge \psi)$ we may conclude $K\phi$ on the basis of **K**.

the proposition that Ann owns a Boxer or a Collie (or both). We take alt(S) to be $\{b \vee c, b, c, b \wedge c\}$, since each of these is (the meaning of) a formal alternative of S (Sauerland 2004) and is presumably relevant. Then the meaning of (17) is the meaning of S, $b \vee c$, plus the conjunction of the negation of all IE alternatives of S. The conjunctive alternative $b \wedge c$ is IE, but the individual disjunct alternatives b and c are not IE. Intuitively, this is because negating b would, together with the meaning of S, entail c; and conversely, negating c would, together with the meaning of S, entail b. So obviously we cannot negate both b and c (doing so would contradict the meaning of S), and instead of making an arbitrary choice of which one to negate, we negate neither. Thus, the overall semantic meaning of (17) is $(b \vee c) \wedge \neg(b \wedge c)$.

(18) **Semantic meaning of (17):**
$(b \vee c) \wedge \neg(b \wedge c)$

Next, based on the maxim of quality, a hearer is licensed to infer that the speaker of (17) believes the semantic content of (17). Thus, (17) licenses the quality inference $K((b \vee c) \wedge \neg(b \wedge c))$.

(19) **Quality inference of (17):**
$K((b \vee c) \wedge \neg(b \wedge c))$

Finally, the maxim of quantity allows us to infer that, for any proposition ϕ that is both relevant and not entailed by S, it's not the case that the speaker believes ϕ, i.e. $\neg K\phi$. Since b and c are each relevant, and since neither of these is entailed by the semantic meaning of (17), the maxim of quantity licenses the inferences $\neg Kb$ and $\neg Kc$. Moreover, the closure conditions on relevance ensure that $\neg b$ and $\neg c$ are each relevant as well, and since neither of these is entailed by the semantic meaning of (17), the maxim of quantity also licenses the inferences $\neg K\neg b$ and $\neg K\neg c$.

(20) **Quantity inferences of (17):**
$\neg Kb \wedge \neg K\neg b \wedge \neg Kc \wedge \neg K\neg c$

Importantly, the pair of inferences $\neg Kb$ and $\neg K\neg b$ together amount to speaker ignorance about b, as characterized by (14). Likewise, $\neg Kc$ and $\neg K\neg c$ characterize speaker ignorance about c.

Putting together the quality and quantity inferences, we get the complete semantic and pragmatic meaning of (17) given below.

(21) **Semantic and pragmatic meaning of (17):**
$$\underbrace{K((b \vee c) \wedge \neg(b \wedge c))}_{\text{Quality inference}} \wedge \underbrace{\underbrace{\neg Kb \wedge \neg K\neg b}_{\text{Ignorance about } b} \wedge \underbrace{\neg Kc \wedge \neg K\neg c}_{\text{Ignorance about } c}}_{\text{Quantity inferences}}$$

4 A generalization about ignorance as quantity inferences in the standard theory

The disjunction example is instructive because it allows us to get a sense of how the various ingredients we assume work together to derive scalar inferences (or not) and/or to derive ignorance inferences (or not). We saw that the disjunctive sentence (17), parsed with *exh*, gives rise to a scalar inference about the conjunctive alternative, namely $\neg(b \wedge c)$, hence doesn't imply any speaker ignorance about that alternative. By contrast, (17) doesn't give rise to a scalar inference about either disjunct alternative (b or c), but does give rise to ignorance inferences about them, due to a certain logical relation between the individual disjunct alternatives and the disjunction as a whole (namely, they aren't entailed, nor are their negations), and to the closure of relevance under negation.

We now characterize more precisely how our procedure of pragmatic reasoning works, so that we can provide an explicit generalization about exactly when (about exactly which propositions) we predict ignorance. We'll see in §5 that the generalization fails to predict the basic contrast in ignorance between *more than* and *at least* (or the class A/B distinction more generally), which will prompt us to slightly amend part of our theory in §6.

The deductive procedure in (22) shows that if ϕ is relevant (in a given context) and S doesn't entail ϕ nor $\neg\phi$, then the maxim of quantity licenses the inference that the speaker is ignorant about ϕ.

(22) **Pragmatic reasoning procedure (as applied to a negated alternative)**
   For any context c, uttered sentence S, and proposition ϕ:

   a. relevant$_c(\phi)$, $[\![S]\!] \not\models \phi$, $[\![S]\!] \not\models \neg\phi$ (assumption)

   b. $\neg K\phi$ ((a) + maxim of quantity)

   c. relevant$_c(\neg\phi)$ ((a) + closure conditions of relevance)

   d. $\neg K\neg\phi$ ((a) + (c) + maxim of quantity)

   e. $I\phi$ ((b) + (d) + def. of ignorance)

If, following Fox 2016, we introduce the notion of *settle*, given in (23), then we can restate our generalization as in (24).

(23) **Definition of *settle***
   S settles ϕ just in case $[\![S]\!] \models \phi$ or $[\![S]\!] \models \neg\phi$.

(24) **Generalization about our pragmatic reasoning procedure**
   For any context c, uttered sentence S, and proposition ϕ, if the maxim of quantity is active in c, and if ϕ is relevant in c and S doesn't settle ϕ, then S gives rise to an inference of speaker ignorance about ϕ.

Thus, for example, (17) gives rise to inferences of speaker ignorance about the individual disjunct alternatives *b* and *c* because (17) doesn't settle either one of them.

This view of the interplay between semantic content and pragmatic reasoning is completely in line with the intuitions of Fox (2007) regarding when and how ignorance arises. However, as we'll see in the next section, this view, at least given just our current assumptions, doesn't seem to have any hope of explaining the contrast in ignorance between *more than* and *at least*. Intuitively, this is because sentences containing (unembedded) occurrences of these quantifiers, like (2) and (3), trigger many alternatives that they don't settle, and yet we don't perceive ignorance about all such alternatives.

5   Applying the standard theory to modified numerals: Too many ignorance inferences

We now apply the theory developed above to sentences with superlative and comparative modified numerals, like (2) and (3). We find that the theory correctly handles the lack of scalar implicatures for both cases—something that was explicitly doubted by Fox and Hackl (2006) and Mayr (2013). However, the theory also overgenerates ignorance inferences; specifically, it doesn't predict the difference in obligatory ignorance between *at least* and *more than* (it predicts ignorance for *more than*), and even for *at least*, it predicts total (rather than partial) ignorance, which Schwarz (2016) has convincingly argued is incorrect. We discuss some potential fixes, and why they don't work. This discussion helps motivate the conservative amendment that we introduce in §6.

5.1   Comparative modified numerals: *More than two*

Consider again (3), repeated below, with the assumed LF containing *exh*.

(3)   Ann owns more than two dogs.
      LF: exh [$_S$ Ann owns more than two dogs]

Following Fox and Hackl (2006), assume that alt(S) = $\{[> n] : n \in \mathbb{Q}^+\}$. Then the semantic meaning of (3) is the meaning of S, $[> 2]$, plus the conjunction of the negation of all IE alternatives of S. It turns out that *every* alternative is IE, and the reason for this is rather trivial: while there are many sets of excludable alternatives, there is no *maximal* set of excludable alternatives, due to density.[10] For

---

10 The adoption of the UDM hypothesis raises the question of what it means for a predicate that intuitively holds only of integers to apply to non-integers, e.g. what it means to say that Ann owns (more than) 2.5 dogs. What is important for us is that *Ann owns more than 2.5 dogs* asymmetrically

instance, the set of all $[> n]$ alternatives for $n \geq 2.1$ is a set of excludable alternatives, but it's not maximal because it's missing, e.g., $[> 2.05]$, and this reasoning can be applied to any other set of excludable alternatives.[11] Since there is no maximal set of excludable alternatives, we have $\text{IE}(S, \text{alt}(S)) = \bigcap \varnothing$, which is equal to the universal set, hence contains *all* alternatives.[12]

Since every alternative is IE, *exh* negates every alternative, and this leads to a contradiction: for example, following the insight of Fox and Hackl (2006), the conjunction of the negation of all alternatives in the set $\{[> n] : n \in \mathbb{Q}^+ \wedge n > 2\}$ is equivalent to $\neg[> 2]$, which contradicts the prejacent of *exh*. (Even more strikingly, the prejacent itself is IE, hence is excluded, which immediately derives a contradiction.)

Since parsing (3) with *exh* leads to a contradiction, a natural assumption to make is that this parse is simply unavailable, hence that (3) has a simple LF without any *exh*. This would then explain why (3) doesn't give rise to any scalar implicatures.[13]

(3) Ann owns more than two dogs.
     LF: [Ann owns more than two dogs]

Unfortunately, however, this account, together with our assumptions about the maxims of quality and quantity and about relevance, predicts that (3) licenses ignorance inferences, which we argued it can't. For example, it's natural to assume that the proposition that Ann owns more than three dogs, $[> 3]$, is relevant in a context where (3) is uttered. Since $[> 3]$ is not entailed by (3), the maxim of quantity licenses the inference $\neg K[> 3]$. By the closure conditions of relevance, the proposition $\neg[> 3]$ is relevant, and it's likewise not entailed by (3). As a result, we derive $\neg K \neg[> 3]$. The inferences $\neg K[> 3]$ and $\neg K \neg[> 3]$ together amount to speaker ignorance about $[> 3]$. This is just an explicit instantiation of the generalization in (24): we derive ignorance about $[> 3]$ because it's presumably relevant and also not

---

entails *Ann owns more than two dogs* and is asymmetrically entailed by *Ann owns more than three dogs*. For Fox and Hackl (2006), these entailments are conceived of as consequences within a deductive system, but they can be made precise in various other ways, as well; see, for instance, Haida and Trinh 2016; Liebesman 2016.

11 For a more formal proof, see Gajewski 2009.

12 Fox and Hackl (2006), whose paper precedes the work on innocent exclusion by Fox (2007) and Chierchia, Fox, and Spector (2012), don't actually define their *exh* operator in terms of innocent exclusion, so their *exh* simply negates every alternative that doesn't (by itself) contradict the prejacent, S. We do assume that *exh* makes reference to innocent exclusion, but since every alternative in this case is IE, the two versions (mostly) coincide. ("Mostly" because on our account, *exh* even excludes alternatives that, by themselves, contradict the prejacent, e.g. $[> 1]$, since they, like every other alternative, are IE.)

13 Another possibility is that all the alternatives are 'pruned' (since not pruning them would derive a contradiction), thus rendering *exh* vacuous in this case. We discuss this possibility, and even provide evidence for a certain version of it (in the context of our amended proposal) in §8. For now, however, we follow the intuitions of Fox and Hackl (2006) that the rescue strategy is to omit *exh*.

settled by (3). The same holds for all propositions of the form $[>n]$ for $n>2$.

Moreover, we note that, although we could assume that every proposition $[>n]$, for $n>2$, is irrelevant, which would thereby block the derivation of ignorance inferences about them, such an assumption would be purely stipulative. Our ultimate theory will actually force these propositions, for principled reasons, to be irrelevant, thus deriving the right results in a non-stipulative way.

To recap, the UDM hypothesis successfully blocks the derivation of unattested scalar implicatures for *more than two*, but our assumptions about the maxims of quality and quantity and about relevance generate many ignorance inferences, provided there are relevant propositions. Given our claim that (3) cannot be used to convey/imply ignorance (though it's compatible with the speaker being ignorant), this is a bad result.

5.2 Superlative modified numerals: *At least two*

Now consider again (2), repeated below with the assumed LF containing *exh*.

(2) Ann owns at least two dogs.
    LF: exh [$_S$ Ann owns at least two dogs]

Fox and Hackl (2006) are doubtful that their account can be extended to handle cases like (2). Their worry is surely well founded if the alternatives of S are simply those obtained by replacing *two* with other numbers, i.e. if $\mathrm{alt}(S) = \{[\geq n] : n \in \mathbb{Q}^+\}$. In this case, the meaning of (2) would be the meaning of S, $[\geq 2]$, plus the conjunction of the negation of all IE alternatives of S. Every alternative 'higher' on the numerical scale than $[\geq 2]$ would be IE, for the following reason: negating all alternatives of the form $[\geq n]$, for $n>2$, in conjunction with the prejacent, $[\geq 2]$, entails $[=2]$, which is not itself an alternative. Put differently, there is only one maximal set of excludable alternatives, namely $\{[\geq n] : n \in \mathbb{Q}^+ \wedge n > 2\}$, so all of its members are IE. As a result, the meaning of (2) would boil down to just $[=2]$—a bad result, to be sure.

If, however, $[=2]$ were itself an alternative, then the set above would not be the only maximal set of excludable alternatives (because $[=2]$ itself would belong to another such set), and we might be able to block unwanted scalar implicatures. This idea, or rather a generalization of this idea, is what Schwarz (2016), following work by Büring (2008) and Mayr (2013), proposes.[14] Specifically, assume that $\mathrm{alt}^{FA}(S)$ (the

---

14 Schwarz's account is cast in neo-Gricean terms, in which primary inferences of the form $\neg K\phi$ are derived first, and then, if possible, are strengthened to secondary (scalar) inferences of the form $K\neg\phi$ (Sauerland 2004). Symmetry effectively entails ignorance at the primary level, thus preempting strengthening at the secondary level. We derive the same result in our framework, albeit in a reverse fashion: symmetry renders certain alternatives non-IE, thereby blocking the derivation of scalar implicatures by *exh*, which in turn allows for the derivation of ignorance inferences as quantity

set of formal alternatives of S) is the set of all alternatives obtained from S replacing *at least* with *exactly* (or *only*) and/or replacing *two* with any number, and assume that all such alternatives are relevant. Then $\text{alt}(S) = \{[\geq n] : n \in \mathbb{Q}^+\} \cup \{[= n] : n \in \mathbb{Q}^+\}$.[15]

The semantic meaning of (2) is therefore the meaning of S, $[\geq 2]$, plus the conjunction of the negation of all IE alternatives of S. Unlike before, the alternatives of the form $[\geq n]$, for $n > 2$, are not IE, because excluding any of them would entail some disjunction of $[= n]$ alternatives to be true. In fact, because of the introduction of $[= n]$ alternatives, *no* alternative is IE, as Schwarz (2016) shows. Take, for instance, the alternative $[\geq 4]$. Excluding $[\geq 4]$ would, together with the meaning of the prejacent, $[\geq 2]$, entail the following disjunction of alternatives: $\bigvee\{[= n] : n \in \mathbb{Q}^+ \land n \geq 2 \land n < 4\}$. Put differently, it's easy to construct a maximal set of excludable alternatives that doesn't contain $[\geq 4]$: take, for instance, $\{[= n] : n \in \mathbb{Q}^+ \land n \neq 4\} \cup \{[\geq n] : n \in \mathbb{Q}^+ \land n > 4\}$. As a result, $[\geq 4]$ is not IE. Similar reasoning applies to all other alternatives (including the $[= n]$ alternatives).

This result means that *exh* in (2) is vacuous, hence that (2) has the same semantic meaning regardless of whether it's parsed with *exh* or not. We therefore successfully replicate Schwarz's account of the lack of scalar implicatures for (2) within our own framework, even with the UDM hypothesis.

In addition, our account, like that of Büring (2008) and Schwarz (2016), derives speaker ignorance about whether Ann owns exactly two dogs and about whether Ann owns three or more dogs. To see this, assume that the proposition that Ann owns exactly two dogs, $[= 2]$, is relevant. Since $[= 2]$ is not entailed by (2), the maxim of quantity licenses the inference $\neg K[= 2]$. By the closure conditions of relevance, we also know that $\neg[= 2]$ is relevant, and since this proposition is likewise not entailed by (2), we derive $\neg K \neg[= 2]$. As a result, we derive $\neg K[= 2]$ and $\neg K \neg[= 2]$, which is to say, speaker ignorance about whether Ann owns exactly two dogs. By the same reasoning, we also derive $\neg K[\geq 3]$ and $\neg K \neg[\geq 3]$. These two pairs of inferences are precisely what Büring (2008) and Schwarz (2016) take to be the complete set of ignorance inferences licensed by (2).

Unfortunately, however, the reasoning outlined above generalizes from $[= 2]$ and $[\geq 3]$ to every conceivable relevant proposition that isn't settled by S. That is, for every proposition $[= n]$, for $n \geq 2$, and every proposition $[\geq n]$, for $n > 2$, we derive speaker ignorance about that proposition. This is just a consequence of the generalization stated in (24).

Thus, we derive what Schwarz (2016) calls "total", rather than "partial", igno-

---

implicatures.

15 Schwarz's account doesn't involve the UDM, but since we adopt it, our set of alternatives is strictly larger than his. This doesn't adversely affect the main result, for now. However, as we show in §6, adopting the UDM in the context of our amended proposal *will* have an effect on our analysis of *at least*, which will make us depart from Schwarz 2016 and instead put us more in line with Kennedy 2015. We discuss the consequences of this in §8.

rance. As Schwarz (2016) argues, however, there appears to be no evidence "for such an implication of total ignorance contributed by *at least*", and "more generally, it appears that the only ignorance implications consistently detectable for *at least n* are those about $[= n]$ and $[\geq (n + 1)]$, the propositions expressed by the alternative statements with *only n* and *at least n + 1*".[16] He supports this claim with the contrast in (25). As expected, (25a) is odd because the first sentence implies that the speaker is ignorant about whether the quintet has more than two German members, which contradicts the entailment of the second sentence that at most two members of the quintet are German. If the first sentence of (25b) implied total ignorance, then it, too, would imply that the speaker is ignorant about whether the quintet has more than two German members, which contradicts the entailment of the second sentence that at most two members of the quintet are German. Thus, (25b) should be equally odd. Since it's perfectly coherent, the first sentence of (25b) must only convey partial ignorance.

(25) a. #At least two members of the quintet were born in Germany. Exactly three were born in Canada.

   b. At least one member of the quintet was born in Germany. Exactly three were born in Canada.

We accept Schwarz's assessment and therefore view the predictions of our theory so far as incorrect. Moreover, we note that, although we could assume that every proposition except for $[= 2]$ and $[\geq 3]$ is irrelevant, such an assumption would be purely stipulative. Our ultimate theory will actually force all these other propositions, for principled reasons, to be irrelevant, thus deriving the right results in a non-stipulative way.[17]

5.3   Interim discussion

Our theory so far correctly predicts the lack of any scalar implicature associated with unembedded occurrences of *at least n* and *more than n*, as in (2) and (3). Density blocks the derivation of a scalar implicature for *more than* (e.g. parsing (3) with *exh* results in a contradiction, and so only the parse without *exh* is licit), as first discussed by Fox and Hackl (2006), while symmetry (even with the assumption of dense scales) blocks the derivation of scalar implicatures for *at least*, as discussed by Schwarz (2016). We take this to be a nice result: Fox and Hackl's account of *more*

---

16 We've taken the liberty of replacing Schwarz's notation by ours in the quotation above.

17 To be more precise, we will derive that $[= 2]$ and $[> 2]$ can be relevant, while all other $[= n]$ and $[\geq n]$ (and $[> n]$) alternatives, for $n > 2$, can't be relevant, for principled reasons. This result will then yield speaker ignorance inferences about $[= 2]$ and about $[> 2]$. We take it that post-semantic granularity considerations are responsible for the intuition that the speaker is ignorant about $[\geq 3]$ for a sentence like (2).

than does extend to *at least*, in the sense that it's compatible with a Schwarz-style analysis of *at least*, involving symmetric alternatives that block scalar implicatures.

Unfortunately, our predictions regarding ignorance are wrong. We incorrectly predict that *more than* licenses ignorance inferences, hence fail to predict the contrast between *at least* and *more than* regarding ignorance, and even in the case of *at least*, we derive total rather than partial ignorance. In short, we predict far too many ignorance inferences across the board. This shortcoming is a symptom of the generalization in (24): because relevance is closed under negation, it follows that if an utterance S doesn't settle a relevant proposition ϕ, then S gives rise to ignorance about ϕ.

Given this problematic result, perhaps the solution to our problem lies in relaxing the closure conditions of relevance.[18] For instance, assume that the set of relevant alternatives is not closed under negation. This assumption would prevent the introduction of unwanted symmetry into the equation for every proposition ϕ and would therefore block the inference of speaker ignorance about ϕ, *unless* ignorance about ϕ is already entailed on the basis of other relevant propositions.

For example, for (2), assume that the relevant propositions all have the form $[= n]$ or $[\geq n]$, for $n \in \mathbb{Q}^+$. Then we still derive $\neg K[= 2]$ by applying quantity reasoning to $[= 2]$. Together with the quality inference of (2), $K[\geq 2]$, this entails the possibility inference $\neg K \neg [> 2]$. Similarly, we still derive $\neg K[> 2]$ by applying quantity reasoning to all propositions of the form $[\geq n]$, for $n > 2$. Together with the quality inference $K[\geq 2]$, this entails the possibility inference $\neg K \neg [= 2]$. Thus, we still derive ignorance about $[= 2]$ and about $[> 2]$, because of the logical relations between the prejacent and the propositions under consideration. However, we don't derive ignorance (in the sense of (14)) for any other propositions, such as, say, $[= 7]$. For example, we do derive $\neg K[= 7]$, but we don't derive the possibility inference $\neg K \neg [= 7]$, because, even though we also derive $\neg K[> 7]$, the quality inference $K[\geq 2]$ and the quantity inference $\neg K[> 7]$ don't entail $\neg K \neg [= 7]$.[19]

As for (3), assume that the relevant propositions all have the form $[= n]$ and $[> n]$, for $n \in \mathbb{Q}^+$. In this case, we correctly don't derive ignorance about any such proposition, because now the logical relations between them and the prejacent are different. For example, let n be any arbitrary number greater than 2. We derive

---

18 The problem we observe here is essentially a corollary of the symmetry problem (mentioned in §1), according to which, if S' is a relevant alternative of S, and no restrictions are placed on alternativehood other than the closure conditions of relevance, then S *and not* S' is also a relevant alternative, and so no scalar inferences are ever expected to arise. As Fox and Katzir (2011) convincingly argue, relevance considerations don't appear to be able to break symmetry. Thus, the line we briefly pursue here is independently problematic. Nevertheless, we discuss it anyway to show that it still runs into another problem.

19 As Schwarz (2016) notes, it's not a problem that we (or he) derive(s) inferences like $\neg K[= 7]$. Given that they're already entailed by $\neg K[> 2]$, they aren't expected to be directly detectable.

$\neg K[= n]$, but we don't derive the possibility inference $\neg K\neg[= n]$, because there's no set of relevant propositions whose exclusion, together with the quality inference $K[> 2]$, entails $\neg K\neg[= n]$. For instance, we also derive $\neg K[> n]$ (by applying quantity reasoning to $[> n]$), but because of density, $\neg K[> n]$ and the quality inference $K[> 2]$ don't entail the possibility inference $\neg K\neg[= n]$. Thus, we don't derive $I[= n]$. In a parallel way, we derive $\neg K[> n]$, but we don't derive the possibility inference $\neg K\neg[> n]$, because there's no set of relevant propositions whose exclusion, together with the quality inference $K[> 2]$, entails $\neg K\neg[> n]$. Thus, we don't derive $I[> n]$, either. In sum, under these assumptions, we don't derive ignorance about $[= n]$ or $[> n]$ for any $n > 2$, which is good news. (Note that the same result would obtain if we assumed that the only relevant propositions were $[> n]$, and not also $[= n]$, for $n \in \mathbb{Q}^+$.)

If this were the end of the story, we'd be in a decent position because we'd manage to derive the basic contrast between *at least* and *more than*: *at least n* implies ignorance about $[= n]$ and $[> n]$, as characterized by the notion of ignorance defined in (14), while *more than n* doesn't imply ignorance about any relevant proposition. However, there are (at least) two serious problems with this line of analysis.

First, the main result (the contrast between *at least* and *more than* and the derivation of partial, rather than total, ignorance for *at least*) vanishes if we allow propositions of the form $[< n]$ or $[\leq n]$ to also be relevant. We'd then predict total ignorance for both *more than* and *at least*. For example, let n be any number greater than 2. Then we could derive both $\neg K[> n]$ and $\neg K[< n]$, which, together with the quality inference $K[\geq 2]$ for (2), or $K[> 2]$ for (3), would entail the possibility inference $\neg K\neg[= n]$, which would, together with $\neg K[= n]$, amount to ignorance about $[= n]$. This reasoning could be applied to any such proposition, thus bringing us back to total ignorance. Notably, it doesn't seem promising or insightful to just stipulate that the $[< n]$ and $[\leq n]$ propositions are irrelevant. (As we discuss in §8, our ultimate theory will actually force them to be irrelevant due in part to structural constraints, which we take to be much less stipulative.)

Second, even if we could somehow prevent those 'negative' propositions from being relevant, we would still derive many inferences of the form $\neg K\phi$, which we believe to be a problem for (3). For instance, (3) would license the inference $\neg K[> 3]$, the proposition that the speaker doesn't have the belief that Ann owns more than three dogs. We believe that such an inference is unavailable, and building on the example in (25) by Schwarz (2016), we illustrate our claim with the contrast in (26). As expected, (26a) is odd, because the first sentence of this sequence implies that the speaker is uncertain whether the quintet has exactly two German members, and/or because it implies that the speaker is uncertain whether the quintet has more than two German members—each of which is derived (in part) on the basis of the inference $\neg K[> 2]$. In other words, the inference $\neg K[> 2]$ is responsible for

the oddity of (26a). If (26b) licensed the inference $\neg K[> 2]$, then it should likewise be judged odd (albeit perhaps not quite as odd as (26a), since the latter entails real ignorance, in the sense of (14)). However, (26b) is perfectly coherent.

(26) a. #At least two members of the quintet were born in Germany. To be more precise, exactly three were born in Berlin.

b. More than one member of the quintet was born in Germany. To be more precise, exactly three were born in Berlin.

6 The missing ingredient: Relevance of belief

Our theory so far correctly predicts the lack of scalar implicatures for comparative modified numerals like *more than two* and superlative modified numerals like *at least two*. In the case of *more than two*, density blocks the computation of a scalar inference: if a sentence like (3) is parsed with *exh*, then all of its (densely populated) alternatives are innocently excludable (since, because of density, there simply is no maximal set of excludable alternatives), but excluding them all leads to a contradiction; thus, parsing (3) with *exh* is illicit, just as discussed by Fox and Hackl (2006). In the case of *at least two*, symmetry blocks the computation of a scalar inference: assuming the presence of both *at least n* and *exactly n* alternatives, *exh* is vacuous because no alternative is innocently excludable (even if density is assumed), and hence no scalar inference arises, just as described by Schwarz (2016).

However, our theory predicts far too many ignorance inferences. It predicts that *more than n* can license ignorance inferences, which we argued against in §2, and it predicts that *at least n* licenses total, rather than just partial ignorance, which Schwarz (2016) convincingly argues against. The reason we predict so many ignorance inferences is directly tied to the closure conditions of relevance: if ϕ is relevant and not settled by S, then the closure conditions of relevance ensure that $\neg\phi$ is likewise relevant, and basic logic ensures that $\neg\phi$ is likewise not settled by S; as a result, the maxim of quantity generates the inferences $\neg K\phi$ and $\neg K\neg\phi$, i.e. speaker ignorance about ϕ.

We now minimally extend our theory in the following way: we add to our closure conditions on relevance the condition that if ϕ is relevant, then so is $K\phi$. That is, the speaker's beliefs about the truth of what is relevant are themselves relevant. This small adjustment has a few important rippling effects. The first is that not only scalar implicatures, but also ignorance inferences must be derived in grammar. The second is that, for grammar to be capable of this, it needs to make available a grammatical belief operator (Chierchia 2006; Meyer 2013).[20] The third

---

20 Fox (2016) has observed these first two points.

is that, since scalar implicatures and ignorance inferences can only be derived in grammar, the maxim of quantity no longer derives additional inferences, i.e. no longer enriches the meaning of an utterance; rather, it acts as a filter on what can be relevant, and this crucial result is what solves our problem: certain alternative propositions are rendered obligatorily irrelevant for principled reasons, which in turn eliminates unattested ignorance inferences. We now discuss all this in more detail.

6.1 Closing relevance under belief

We now add to our closure conditions on relevance the condition that if ϕ is relevant, then it's also relevant whether the speaker believes ϕ.

(27) **Closure conditions of relevance** (revised)

    a. If ϕ and ψ are both relevant, then so is $\phi \wedge \psi$.

    b. If ϕ is relevant, then so is $\neg\phi$.

    c. If ϕ is relevant, then so is $K\phi$.

This move receives independent motivation from Fox (2016), who observes the following empirical fact: "Silence is uncooperative." For example, in the context of a murder trial, if the lawyer asks the witness, w, "Where was John at the time of the murder?", w can't just look the lawyer in the eye and remain silent. The intuition, as Fox puts it, is that "if w believes something that bears on John's whereabouts at the time of the murder, w is required to say so. If not, w is required to reveal this lack of opinion."

A consequence of closing relevance under belief, together with the maxim of quantity, is that ignorance can only be derived in grammar. We demonstrate this with the deductive procedure in (28).

(28) **Pragmatic reasoning procedure (as applied to K-prefixed alternative)**
For any context c, uttered sentence S, and proposition ϕ:

    a. $\text{relevant}_c(\phi)$, $[\![S]\!] \not\models \phi$, $[\![S]\!] \not\models \neg\phi$ (assumption)

    b. $\neg K\phi$ ((a) + maxim of quantity)

    c. $\text{relevant}_c(\neg\phi)$ ((a) + closure conditions of relevance)

    d. $\neg K\neg\phi$ ((a) + (c) + maxim of quantity)

    e. $\text{relevant}_c(\neg K\phi)$ ((a) + closure conditions of relevance)

    f. $\text{relevant}_c(\neg K\neg\phi)$ ((a) + closure conditions of relevance)

    g. $[\![S]\!] \models \neg K\phi \wedge \neg K\neg\phi$ or $[\![S]\!] \not\models \neg K\phi \wedge \neg K\neg\phi$ (tautology)

    h. Case 1: $[\![S]\!] \models \neg K\phi \wedge \neg K\neg\phi$
        $[\![S]\!] \models I\phi$ (above line + def. of ignorance)

i. Case 2: $[\![S]\!] \not\models \neg K\phi \wedge \neg K\neg\phi$

    Subcase 1: $[\![S]\!] \not\models \neg K\phi$
$\neg K\neg K\phi$	((e) + above line + maxim of quantity)
$K\neg K\phi$	((b) + negative introspection)
contradiction	(above two lines)

    Subcase 2: $[\![S]\!] \not\models \neg K\neg\phi$
$\neg K\neg K\neg\phi$	((f) + above line + maxim of quantity)
$K\neg K\neg\phi$	((d) + negative introspection)
contradiction	(above two lines)

This deductive procedure illustrates that, if the maxim of quantity is active in a given context c, and if ϕ is relevant in c and not settled by S, then one of two things follows: either S entails speaker ignorance about ϕ, or the maxim of quantity licenses the inference of a contradiction. We take it to be an empirical fact that the latter option is unavailable: we know of no cases where a listener actually draws a contradictory quantity inference.[21] Thus, ignorance must be able to be derived in grammar, and it can no longer be derived by quantity reasoning (since that would yield a contradiction); that is, ignorance must be derived in grammar alone.

Since ignorance must be derived in grammar (and since not every sentence begins with *I know* or *I believe*), language must make available a covert belief operator, such as the syntactic item K proposed for independent reasons by Meyer (2013). We adopt K and take its meaning to be the same as that of the belief operator we previously assumed; that is, overloading the symbol K, we have $[\![K]\!]\phi = K\phi$.

Crucially, in every context where the maxim of quantity is active, we take it that every sentence must be parsed with a matrix-level K, in order to be able to derive ignorance grammatically.[22, 23]

---

21 To be sure, formulating a diagnostic to test this claim is nontrivial, but at the very least, for sentences that the theory predicts to give rise to contradictory quantity inferences, we should detect some oddity. For example (to foreshadow a bit), for a sentence like (3) (*Ann has more than two dogs*), if any propositions of the form $[> n]$, for $n > 2$, are taken to be relevant, we'll end up deriving a contradictory quantity inference. But given that (3) isn't perceived to be odd in any way, we take this as an indication that no such propositions can in fact be relevant.

22 In contexts where the maxim of quantity is inactive, this condition doesn't hold. We briefly discuss such contexts in §8.

23 Meyer (2013, p. 42) comes to the same conclusion, for independent reasons, and calls this the "Matrix K hypothesis", because for her, it's an auxiliary assumption used to explain certain empirical facts. For us, the idea that relevance is closed under belief is the crucial ingredient (or hypothesis), and matrix K is essentially a consequence of this. More precisely, the consequence for us is that at least some sentences (the ones for which we observe ignorance inferences) can be parsed with matrix K. We could, therefore, adopt a weaker version of (29), replacing "must" with "can". However, this would raise an important question that we wish to ignore for now, namely why certain sentences (the ones for which we obligatorily observe ignorance inferences, like (2)) must be parsed with matrix K. We think that weakening (29) in this way, and answering this question, could lead to deeper insights;

(29) **Matrix *K***
    Every (assertive) sentence *S* must be c-commanded by an occurrence of *K*.

Given this, together with our assumption in (9) that *exh* may attach to any propositional constituent, it follows that *exh* can apply both above and/or below *K* (cf. Meyer 2013, p. 43).

6.2 Obligatory irrelevance

We now show how closing relevance under belief actually ends up conspiring with the maxim of quantity to act as a filter on relevance. First, note that the deductive procedure in (28) yields the generalization in (30), which is in fact just a stronger version of the generalization in (24) from §4, repeated below.

(24) **Old generalization about our pragmatic reasoning procedure**
    For any context *c*, uttered sentence *S*, and proposition ϕ, if the maxim of quantity is active in *c*, and if ϕ is relevant in *c* and *S* doesn't settle ϕ, then *S* gives rise to an inference of speaker ignorance about ϕ.

(30) **New generalization about our pragmatic reasoning procedure**
    For any context *c*, uttered sentence *S*, and proposition ϕ, if the maxim of quantity is active in *c*, and if ϕ is relevant in *c* and *S* doesn't settle ϕ, then $[\![S]\!] \vDash I\phi$.

The difference between the two—the effect of closing relevance under belief—has to do with the way in which *S* "gives rise to ignorance". Previously, ignorance was a purely pragmatic inference licensed directly by the maxim of quantity. Now, ignorance is an entailment of the semantic meaning of the uttered sentence.

As a consequence of the new generalization, we can now state a precise condition on relevance in contexts where the maxim of quantity is active.

(31) **Condition on relevance (a consequence of closing relevance under belief)**
    For any context *c*, uttered sentence *S*, and proposition ϕ, if the maxim of quantity is active in *c*, and if *S* doesn't settle or entail ignorance about ϕ, then ϕ isn't relevant in *c*.

For now, we assume that all the sentences we consider are interpreted in contexts where the maxim of quantity is active and obeyed (not flouted). (In §8, we

---

for instance, Spector (2015) argues that the status of *at least* as a global positive polarity item can be explained in terms of *at least* needing to be in the scope of *exh*, and if this is true, and if our proposal is correct, then it would actually need to be in the scope of *exh K*, since otherwise the maxim of quantity would yield a contradiction (there would be relevant alternatives, accessible to *exh*, which the sentence as a whole doesn't settle or entail ignorance about). However, we have to leave this inquiry for a future occasion, and so we instead adopt what is given in (29).

Draft: July 31, 2017

consider what our proposal predicts for contexts where the maxim of quantity is inactive.)

The condition in (31) restricts which propositions can be relevant in a specific context in which the maxim of quantity is active. In §7, we'll see that certain propositions can't be relevant (to certain sentences) in *any* context where the maxim of quantity is active. We call such propositions *obligatorily irrelevant*.

(32)   **Definition of *obligatory irrelevance***
ϕ is *obligatorily irrelevant* to S just in case for every context c, if the maxim of quantity is active in c, then S doesn't settle or entail ignorance about ϕ.

6.3   Disjunction revisited

Let's return to the textbook case of disjunction in (17), repeated below, with the new assumed LF. As before, we wish to show that this LF leads to just the right scalar and ignorance inferences (entailments now), but now we also need to show that, in doing so, the maxim of quantity can be active (i.e. doesn't yield a contradiction).

(17)   Ann owns a Boxer or a Collie.
LF: exh $[_{S_2}$ K exh $[_{S_1}$ Ann owns a Boxer or (Ann owns) a Collie]]

We take a bottom-up approach and start by considering the meaning of *exh* S_1. As before, we take alt(S_1) to be $\{b \lor c, b, c, b \land c\}$. *exh* S_1 is therefore precisely the same as what we derived in §3.2 as the semantic meaning of (17) itself on the old parse, with just one *exh* and no K: $(b \lor c) \land \neg(b \land c)$. Thus, the meaning of *K exh* S_1 is $K((b \lor c) \land \neg(b \land c))$.

The semantic meaning of (17) is therefore the meaning of *K exh* S_1, which we just computed, plus the conjunction of the negation of all IE alternatives of S_2. Following Meyer 2013, we take it that alt$_c(S_2) = \{[\![K \text{ exh } S']\!] : S' \in \text{alt}_c^{FA}(S_1) \land [\![S']\!] \in \text{relevant}_c\}$, the set of alternatives obtained by replacing in S_2 any relevant alternative of S_1.[24] (Below, we abbreviate *Ann owns a Boxer* by "B" and *Ann owns a Collie* by "C".)

(33)   alt$_c(S_2) = \{[\![K \text{ exh } S']\!] : S' \in \text{alt}_c^{FA}(S_1) \land [\![S']\!] \in \text{relevant}_c\}$
$= \{[\![K \text{ exh B or C}]\!], [\![K \text{ exh B}]\!], [\![K \text{ exh C}]\!], [\![K \text{ exh B and C}]\!]\}$
$= \{K((b \lor c) \land \neg(b \land c)), K(b \land \neg c), K(c \land \neg b), K(b \land c)\}$

The alternative corresponding to the prejacent itself (the first alternative in the set above) is of course not IE, but all of the other alternatives are IE, which means we exclude them all. Note that the negation of $K(b \land c)$ is already entailed by the meaning of $K S_1$, and so we omit it below.

---

24 See Meyer 2013 for arguments why deletion (in the sense of Katzir 2007) of K or *exh* is illicit in generating this set of formal alternatives.

(34) **Semantic meaning of (17)**

$$\underbrace{K((b \vee c) \wedge \neg(b \wedge c))}_{\text{Prejacent of matrix \textit{exh}}} \quad \wedge \quad \underbrace{\neg K(b \wedge \neg c) \wedge \neg K(c \wedge \neg b)}_{\text{Exclusions of matrix \textit{exh}}}$$

Crucially, $K((b \vee c) \wedge \neg(b \wedge c))$ and $\neg K(b \wedge \neg c)$ together entail both $\neg Kb$ and $\neg K\neg c$. By the same token, $K((b \vee c) \wedge \neg(b \wedge c))$ and $\neg K(c \wedge \neg b)$ together entail both $\neg Kc$ and $\neg K\neg b$. These four inferences together ($\neg Kb$, $\neg K\neg b$, $\neg Kc$, and $\neg K\neg c$) amount to a pair of speaker ignorance inferences concerning b and c.

Thus, what we end up with as the semantic meaning of (17) is precisely what we derived earlier as the complete semantic and pragmatic meaning of (17) parsed with just one *exh* and no K. That is, the semantic meaning of (17) that we arrive at can be equivalently stated as in (35).[25] The difference is that, what we derived earlier as a quantity inference given the closure conditions of relevance, we now derive grammatically, based purely on the semantics of the LF structure.

(35) **Semantic meaning of (17)** (equivalent version)

$$K((b \vee c) \wedge \neg(b \wedge c)) \wedge \neg Kb \wedge \neg K\neg b \wedge \neg Kc \wedge \neg K\neg c \qquad \text{(same as (21))}$$

The quality inference we derive is that the speaker knows the semantic meaning of (17), i.e. $K [\![(17)]\!]$. However, given that we assume the validity of both $K\phi \rightarrow KK\phi$ (positive introspection) and $\neg K\phi \rightarrow K\neg K\phi$ (negative introspection), this inference is already a valid inference regardless of the maxim of quality.[26]

Finally, as we noted earlier, the maxim of quantity no longer derives any additional inferences, but rather acts as a filter on relevance. In other words, the question now is not what inferences the maxim of quantity licenses, but rather whether the maxim of quantity can be active in a context where the structure in (17) is being interpreted, under the assumption that the formal alternatives we've assumed are relevant. We now show that the maxim of quantity can indeed be active, i.e. that no contradiction arises.

If the maxim of quantity is active, then it must hold that for every relevant proposition ϕ, either (17) settles ϕ or (17) entails ignorance about ϕ. This entails showing that for each alternative in $\{b \vee c, b, c, b \wedge c\}$, either (17) settles that alternative or entails ignorance about it. We also need to show the same thing for each proposition that can be derived from this set on the basis of the closure conditions of relevance. If we discover that at least one such proposition is such that (17) neither settles nor entails ignorance about that proposition, it means that the maxim of quantity would yield a contradiction, hence that (17), under the assumption that $\text{alt}(S_1) = \{b \vee c, b, c, b \wedge c\}$, can't be interpreted in a context where the maxim of

---

25 We've shown that (34) entails (35). For the other direction, note that $\neg Kb$ entails $\neg K(b \wedge \neg c)$, and that $\neg Kc$ entails $\neg K(c \wedge \neg b)$.

26 More precisely, $K\phi \wedge \neg K\psi$ (the basic form of the semantic meaning of (17)) already entails $KK\phi \wedge K\neg K\psi$ in S5 (by positive and negative introspection), which is equivalent to $K(K\phi \wedge \neg K\psi)$.

quantity is active (a bad result); otherwise, we're in good shape.

We begin with $b \vee c$. Since (17) entails $K(b \vee c)$, it also entails $b \vee c$ by the **T** axiom.[27] As a result, (17) settles $b \vee c$. Likewise, (17) entails $\neg(b \wedge c)$ (again by **T**), which means it also settles $b \wedge c$. Finally, as illustrated above, (17) entails speaker ignorance about the individual disjunct alternatives b and c. Thus, (17) settles or entails ignorance about every proposition in alt(S_1).

Second, we need to show that for every proposition ϕ in the closure of alt(S_1) under conjunction, negation, and belief, (17) settles ϕ or entails ignorance about ϕ. In the appendix (§9), we prove that if a sentence Σ of the form *exh K (exh) S* (like (17) here) settles or entails ignorance about every proposition in a set A, then Σ settles or entails ignorance about every proposition in the closure of A under conjunction, negation, and belief. (As a consequence of this, it follows that each alternative in alt(S_2), the domain of the matrix-level *exh*, can be relevant.)

To recap, we've shown that, if (17) is parsed as *exh K exh B or C*, and if we assume that the alternatives of a disjunctive sentence are the individual disjuncts and their conjunction, then we derive precisely the right meaning for (17) (Meyer 2013). We also showed that, in doing so, the maxim of quantity can be active, i.e. doesn't yield a contradiction. Put differently, the formal alternatives that yield the right inferences are all capable of being relevant without the maxim of quantity yielding a contradiction; hence, none of them are obligatorily irrelevant in the sense of (32).

As we'll now see, when we turn to *more than* and *at least*, the alternative propositions that previously overgenerated ignorance inferences are such that, if they're assumed to be relevant, then the maxim of quantity yields a contradiction. As a consequence, we obtain the result that, if the maxim of quantity is taken to be active, then those propositions can't be relevant to begin with—they're obligatorily irrelevant. Thus, the maxim of quantity acts as a filter on relevance, and it crucially renders irrelevant precisely those propositions that were previously problematic.

7 The maxim of quantity as a filter on relevance: Obligatory irrelevance

We now revisit our running examples (3), with *more than two*, and (2), with *at least two*, and show how our theory derives the right results.

---

27 Inferences from $K\phi$ to ϕ seem to be routinely drawn in cooperative conversational exchanges. Schwarz (2016) calls such inferences *bottom line* inferences.

7.1 Comparative modified numerals revisited: *More than two*

Let's discuss once more (3). In §5, we observed that, if we assume that the dense set of $[> n]$ propositions are relevant to (3), then parsing (3) with *exh* (and no *K*) leads to a contradiction, which means it isn't parsed with *exh*, hence why no scalar implicature arises. However, the closure conditions of relevance, specifically closure of relevance under negation, ensured that we incorrectly derived ignorance about each $[> n]$ alternative for $n > 2$; for example, the maxim of quantity derived both $\neg K[> 5]$, assuming $[> 5]$ is relevant, and $\neg K \neg [> 5]$, since relevance is closed under negation. We also argued that it'd be purely stipulative to say that $[> 5]$ is irrelevant. We now show that, on our current proposal, $[> 5]$ actually can't be relevant, for principled reasons, and the same for all other problematic propositions — they're all obligatorily irrelevant, in the sense of (32).

Consider the new assumed LF for (3) below, and assume, as before, that $\text{alt}^{FA}(S_1)$, the set of formal alternatives of S_1, are those obtained by replacing *two* with any number in $\mathbb{Q}^+$. The crucial point now is that, if we attempt to assume that any of these are relevant, for $n > 2$, then the maxim of quantity yields a contradiction. Let's see why. (For the purposes of the following discussion, assume that the lower *exh* is not present; otherwise, we'd immediately derive a contradiction, regardless of the maxim of quantity.)

(3) Ann owns more than two dogs.
   LF: exh [$_{S_2}$ K (exh) [$_{S_1}$ Ann owns more than two dogs]]

Let c be any arbitrary context such that $\{[> n] : n \in \mathbb{Q}^+\} \subseteq \text{relevant}_c$; hence, $\text{alt}_c(S_1) = \{[> n] : n \in \mathbb{Q}^+\}$. Then the semantic meaning of (3) is $K[> 2] \wedge \bigwedge \{\neg K[> n] : n > 2\}$. Now consider the alternative $[> 5]$. (3) neither entails $[> 5]$, nor its negation, hence doesn't settle $[> 5]$. Moreover, (3) doesn't entail ignorance about $[> 5]$ because, even though it entails $\neg K[> 5]$, it doesn't entail the possibility $\neg K \neg [> 5]$. The same holds for all alternative propositions of the form $[> n]$, for $n > 2$. As a result, if things are as we assumed, then the maxim of quantity can't be active in c for the interpretation of (3); otherwise, it yields a contradiction, illustrated by the procedure in (28).

If we additionally included $[= n]$ alternatives (those obtained by replacing *more than* with *exactly* and by replacing *two* with any number in $\mathbb{Q}^+$) as relevant, formal alternatives, nothing would change: (3) would still fail to settle or entail ignorance about any $[> n]$ proposition for $n > 2$ (and the same for any $[= n]$ proposition for $n > 2$). For instance, (3) would still fail to settle $[> 5]$, and while it would still entail $\neg K[> 5]$, it still wouldn't entail $\neg K \neg [> 5]$. Thus, the maxim of quantity, if it were active, would still yield a contradiction.

Importantly, this result — that the maxim of quantity can't be active in c without yielding a contradiction — also holds for any context c' in which any subset of

the $[> n]$ (or $[= n]$) alternatives are relevant; clearly, restricting the set of relevant alternatives doesn't settle or entail ignorance about more propositions than before.

If, however, we assumed that the set of formal alternatives of S_1 also included ones like *not more than n*, e.g. $[\leq 5]$ (*not more than five*), then the maxim of quantity could be active in c, because now (3) would settle or entail ignorance about every alternative, and so no contradiction would arise from the maxim of quantity. For instance, (3) would entail both $\neg K[> 5]$ (on the basis of the alternative $[> 5]$) and $\neg K[\leq 5]$ (on the basis of the alternative $[\leq 5]$), which is equivalent to $\neg K\neg[> 5]$. Thus, (3) would entail ignorance about $[> 5]$. More generally, (3) would settle all alternatives of the form $[> n]$ and $[\leq n]$ for $n \leq 2$, and entail ignorance about all other alternatives. This would be a bad result: as we observed in §2, (3) doesn't imply any ignorance. Thankfully, from the perspective of Katzir 2007, it's natural assume that *not more than n* isn't a formal alternative of *more than n*, since it's structurally more complex. But what about *fewer than n* or *at most n*, each of which would replicate the bad result just described? Here we have to assume that these are likewise structurally more complex than *more than n*. Such an assumption doesn't seem too outlandish to us, and it in fact follows from several proposals in the literature for independent phenomena like split-scope readings (cf. Heim 2006; Penka 2014). We take up this issue in §8.

Assuming, then, that the formal alternatives of S_1 are those obtained by replacing *two* with any number in $\mathbb{Q}^+$ (and perhaps also by replacing *more than* with *exactly*), we obtain the result that in *every* context where the maxim of quantity is active, *no* such alternative for $n > 2$ can be relevant — because if any such alternative were relevant, then as just illustrated, the maxim of quantity would yield a contradiction. In short, all such alternatives are obligatorily irrelevant to (3), in the sense of (32).

Since all propositions $[> n]$ (and $[= n]$) for $n > 2$ are obligatorily irrelevant to (3), it follows that alt(S_2), which feeds matrix exhaustification, doesn't include any alternatives of the form $K[> n]$ (or $K[= n]$) for $n > 2$, which means that the matrix-level *exh* doesn't exclude those propositions. Given all this, the semantic meaning of (3) is just $K[> 2]$. (Note that $[> 1]$, $[= 2]$, and so on are settled by (3), hence can be relevant to (3), and therefore alt(S_2) can include $K[> 1]$, $K[= 2]$, and so on. *exh* thus can exclude these, but their negations are already entailed by the prejacent, $K[> 2]$.)

(36) **Semantic meaning of (3)**
      $K[> 2]$

Lastly, since all alternatives $[> n]$ (and $[= n]$) for $n > 2$ are obligatorily irrelevant, parsing (3) with a second *exh*, below K, would no longer derive a contradiction (it would just be vacuous), hence why we put it in parentheses. In §8, we'll discuss an

example where we actually do need this lower *exh*.

Before concluding this subsection, we note that if *more than two* is embedded under a universal quantifier like *required*, as in (7c), repeated below, the alternatives $[> n]$ (and $[= n]$) for $n > 2$ can now all be relevant, and this is because they're all innocently excludable, hence settled when *exh* applies below K. The meaning we derive is thus $K(\square[> 2] \wedge \bigwedge\{\neg \square[> n] : n > 2\})$, which accounts for the inference discussed in §2. (The matrix *exh* doesn't exclude anything, hence could be omitted.)

(7c)   Bill is required to take more than two courses.
      LF: (exh) K exh [required [Bill takes more than two courses]]

7.2   Superlative modified numerals revisited: *At least two*

Now consider once again (2), repeated below, along with the new assumed LF. We first show that, if we assume that the formal alternatives of S_1 are all the *at least n* and *exactly n* alternatives, we derive the result that they're all obligatorily irrelevant, for $n > 2$, which in turn means that the matrix-level *exh* doesn't operate on any $K[\geq n]$ or $K[= n]$ alternatives, for $n > 2$, hence is vacuous and in particular fails to derive any ignorance whatsoever. We then show that if we additionally include *more than n* alternatives, we derive precisely what we want.

(2)   Ann owns at least two dogs.
      LF: exh [$_{S_2}$ K (exh) [$_{S_1}$ Ann owns at least two dogs]]

Assume, then, that *at least* can be replaced by *exactly* (or *only*) and that *two* can be replaced by any number in $\mathbb{Q}^+$. Now consider the alternative $[= 2]$, which we would like to be relevant, and specifically for (2) to entail ignorance about. If $[= 2]$ were relevant, then $K[= 2]$ would be relevant too, and hence (2) would entail $\neg K[= 2]$. To entail ignorance about $[= 2]$, (2) would also need to entail the possibility $\neg K \neg [= 2]$. This would be derivable if all the $[\geq n]$ alternatives for $n > 2$ were relevant (so that (2) would exclude all $K[\geq n]$ for $n > 2$), but not otherwise. Unfortunately, none of these alternatives can be relevant. Take an arbitrary alternative $[\geq n]$ for some $n > 2$. Clearly, (2) doesn't entail $[\geq n]$, nor its negation, given the structural alternatives we assume, so (2) doesn't settle $[\geq n]$. (2) does entail $\neg K[\geq n]$, since $K[\geq n]$ is IE, but (2) doesn't entail the possibility $\neg K \neg [\geq n]$. Thus, (2) neither settles nor entails ignorance about $[\geq n]$, which means that $[\geq n]$ is obligatorily irrelevant. Since n was arbitrary, every alternative $[\geq n]$ for $n > 2$ is obligatorily irrelevant, which in turn means that (2) neither settles nor entails ignorance about $[= 2]$, either, so $[= 2]$ is also obligatorily irrelevant. In short, every $[\geq n]$ and $[= n]$ alternative for $n > 2$ is obligatorily irrelevant. So, if we assume just these formal alternatives, (2) is predicted to just mean $K[\geq 2]$, i.e. not to entail any ignorance at all, which is wrong.

Our solution to this problem is to assume that the formal alternatives of *at*

least two include not just *at least n* and *exactly n* alternatives, but also *more than n* alternatives. Thus, for example, *Ann owns more than two dogs* is a formal alternative of S_1, hence *K (exh) Ann owns more than two dogs* is a formal alternative of the prejacent, S_2. Now consider again $[= 2]$, which, to repeat, we'd like to be relevant, and specifically for (2) to entail ignorance about. Now, $[= 2]$ and $[> 2]$, which are both denoted by formal alternatives, can *both* be relevant, because (2) can entail ignorance about them: if both are relevant, then (2) entails both $\neg K[= 2]$ and $\neg K[> 2]$, which together with the prejacent, $K[\geq 2]$, entails the two possibilities $\neg K\neg[= 2]$ and $\neg K\neg[> 2]$. Thus, (2) entails ignorance about both $[= 2]$ and $[> 2]$, as desired. This result obtains because $[= 2]$ and $[> 2]$ form a symmetric pair relative to $[\geq 2]$.

Crucially, every other $[\geq n]$, $[= n]$, and $[> n]$ alternative for $n > 2$ is obligatorily irrelevant because, given the formal alternatives we assume, (2) can't possibly settle or entail ignorance about them. So the semantic meaning we derive for (2) is the following.

(37) **Semantic meaning of (2)**
      $K[\geq 2] \wedge \neg K[= 2] \wedge \neg K[> 2]$
      $\equiv K[\geq 2] \wedge \neg K[= 2] \wedge \neg K\neg[= 2] \wedge \neg K[> 2] \wedge \neg K\neg[> 2]$

This is precisely the partial ignorance reading that we wished to derive, following Büring (2008), Kennedy (2013, 2015), and Schwarz (2016).

To recap, assuming that *at least two* has as its formal alternatives all alternatives of the form *at least n*, *exactly n* (or *only n*), and *more than n*, every such alternative for $n > 2$ ends up obligatorily irrelevant to (2), and the alternatives that we're left with, $[= 2]$ and $[> 2]$, yield precisely the attested inferences. Symmetry plays an important role for our account of the ignorance inferences associated with (2): symmetry is what allows (2) to entail ignorance about $[= 2]$ and about $[> 2]$, which in turn lets them be relevant. In the case of (3), because of density, no alternative has a symmetric partner, so (3) can't entail ignorance about (nor settle) any alternative for $n > 2$, which renders all such alternatives obligatorily irrelevant.

Finally, we note once more that, as in the case of *more than two*, if *at least two* is embedded under a universal quantifier like *required*, as in (7b), repeated below, the alternatives $[\geq n]$, $[> n]$, and $[= n]$ for $n > 2$ can now all be relevant, and this is again because they're all innocently excludable, hence settled when *exh* applies below K. The meaning we derive is thus $K(\Box[\geq 2] \wedge \neg \Box[= 2] \wedge \neg \Box[> 2])$, which again accounts for the inference discussed in §2.[28] (Again, the matrix *exh* doesn't

---

28 We actually derive something a little stronger than what we reported in (2). There, we reported only the inference $\neg \Box[> 2]$ (actually, $\neg \Box[\geq 3]$), which, together with the prejacent, entails that Bill is allowed to take exactly two courses. Here, we also derive $\neg \Box[= 2]$, which, together with the prejacent, entails that Bill is also allowed to take more than two courses. This result is completely in line with what Büring (2008) takes to be the "authoritative reading" of (7b).

Draft: July 31, 2017

exclude anything, hence could be omitted.)

(7c)   Bill is required to take at least two courses.
        LF: (exh) K exh [required [Bill takes at least two courses]]

8   Discussion

We conclude with a discussion of a couple issues our theory faces (and how to deal with them), a number of positive predictions it makes, and what it would take to recast our theory in neo-Gricean terms.

Obligatory admissibility of *more than* n alternatives.   Our proposal forces us to assume that sentences like (2) with *at least two* have *more than* n alternatives as formal alternatives. This result still keeps our analysis of the ignorance inferences of *at least* n in line with Büring 2008, since ignorance ultimately still arises on the basis of symmetry; in fact, our proposal coincides more or less with Kennedy 2015, who likewise derives ignorance for *at least two* by appealing to *exactly two* and *more than two* alternatives. (His theory is couched in neo-Gricean terms, so he stipulates that *at least*, *exactly*, and *more than* form a Horn scale.) However, Schwarz (2016) argues against the use of *more than* n alternatives, for the following reason. *At least* and *only* have wider syntactic distributions than *more than*, as illustrated in (38) and (39). If ignorance arose from *at least* on the basis of *only* and *more than*, then the fact that (38a) and (39a) both convey ignorance would be unexpected: since (38c) and (39c) are both ungrammatical (and assuming that ungrammatical structures can't serve as alternatives for scalar implicature computation), they can't be alternatives to (38a) and (39a), hence (38b) and (39b) wouldn't have symmetric partners.

(38)   a.   Ann is at least allowed to own two dogs.

        b.   Ann is only allowed to own two dogs.

        c.   *Ann is more than allowed to own two dogs.

(39)   a.   Ann owns two dogs, at least.

        b.   Ann owns two dogs, only.

        c.   *Ann owns two dogs, more than.

Our response is that we don't think our account is necessarily insufficiently general. Although replacing *at least* with *more than* in the surface structures above yields ungrammaticality, the question is whether the LF of (38a) and the LF of (39a) have structural alternatives involving *more than* which are grammatical LFs, and answering this question relies on explicit assumptions about the syntax and (focus) semantics of *at least*. For instance, (38a) may well have the LF in (40a), and (39a) the

LF in (40b), despite surface appearances, and clearly, replacing *at least* with *more than* in such structures preserves grammaticality.[29]

(40) a. exh K allowed [Ann owns at least two dogs]

  b. exh K [Ann owns at least two dogs]

Moreover, note that, to the extent that the examples in (38a) and (39a) are a problem for an analysis of English *at least* that relies on *more than* alternatives, a similar problem could be constructed even for a Schwarz-style analysis of *at least* in other languages in which (the equivalent of) *at least* has a wider syntactic distribution than (the equivalent of) *only*. German *mindestens* 'at least' is a case in point.

(41) *Ann hat zwei Hunde besessen, {mindestens / *nur / *mehr als}.*
    Ann has two dogs owned {at least   / only / more than}

  'Ann owned {at least / only / more than} two dogs.'

These data show that, if the problem Schwarz (2016) raises for an analysis of English *at least* relying on *more than* alternatives is valid, then an analysis of German *mindestens* relying on *nur* alternatives faces a nearly identical problem. We leave a more detailed exploration of these issues for future work.

Inadmissibility of *fewer than* n and *at most* n alternatives. For sentences with *more than n* and *at least n*, we crucially need to assume that the sentences one would obtain by replacing *more than* and *at least* with *fewer than* or *at most* are not viable alternatives; otherwise, we'd derive total ignorance both for *more than n* and for *at least n*. From a structural perspective, we think this isn't too outlandish: it's plausible that *fewer than* and *at most*, which possibly contain some kind of negation at some level, are structurally more complex than either *more than* or *at least*. For example, Penka (2014) proposes that *at most* is decomposed into a negative antonymizer *ANT* plus *at least*, which explains why a sentence like *Ann is allowed to own at most two dogs* can have an upper-bounded authoritative reading: it has a parse where *ANT* scopes above *allowed*, while the *at least* remnant stays below *allowed*, yielding the reading 'it's not the case that Ann is allowed to own at least *n* dogs, for any *n* > 2'.

---

29 Alternatively, we might simply object to the premise of the problem, namely that unassertible LFs can't serve as alternatives for the purposes of implicature calculation. As Buccola, Križ, and Chemla (2016) have argued, and attempted to demonstrate experimentally, alternative-based reasoning may occur at a more conceptual, less linguistic level, and in particular the alternatives that enter into such reasoning may not be totally constrained by one's linguistic lexicon. If so, then perhaps the (conceptual) alternative that one would obtain by replacing *at least* with *more than* in (the conceptual counterparts of) (38a) or (39a) play a role, even though they can't be vocalized.

By the same token, however, for sentences like *Ann owns fewer than/at most two dogs*, we now need the *more than n* and *at least n* alternatives to be inadmissible, lest we again derive total ignorance in both cases. Following the above discussion, if *fewer than* and *at most* are strictly more complex than *more than* and *at least*, then we expect to be able to derive such alternatives, simply by deleting the relevant negative expression, e.g. *ANT*. To solve this issue, we might simply argue that negation cannot be deleted for the purposes of structurally deriving alternatives. Again, we think this isn't too outlandish an assumption, and in fact we observe evidence for it, as in (42).

(42)   Bill didn't eat all of the cookies.

An utterance of (42) implicates that Bill ate some of the cookies. This implicature is expected if the only alternative of (42) is *Bill didn't eat some (any) of the cookies*, whose negation is equivalent to 'Bill ate some or all of the cookies'. Together with the basic the meaning of (42), the overall meaning derived is that Bill ate some but not all of the cookies, which is exactly right.

If, however, *Bill ate some of the cookies* were also an alternative of (42), as it would be on a simple, structural theory of formal alternatives, then it would form a symmetric pair with *Bill didn't eat some (any) of the cookies*, hence incorrectly lead to speaker ignorance about whether Bill ate some of the cookies, rather than the implicature that he did in fact eat some. Thus, this example shows that it's not too implausible to think that negation can't be deleted for the purposes of deriving formal alternatives. See Trinh and Haida 2015 for further details.

Role of the UDM hypothesis.   Our proposal is compatible *more than n* having *exactly n* as a structural alternative, which we take to be a positive feature since, presumably, they are structural alternatives of one another. This result relies crucially on the UDM hypothesis. If the alternatives weren't densely populated, then $[= 3]$ and $[> 3]$ would form a symmetric pair relative to $[> 2]$, and so (3) would entail ignorance about $[= 3]$ and ignorance about $[> 3]$. Specifically, $K[> 2]$, $\neg K[= 3]$, and $\neg K[> 3]$ (all entailments of (3)) would entail the possibility inferences $\neg K\neg[> 3]$ and $\neg K\neg[= 3]$. As a result, the alternatives $[= 3]$ and $[> 3]$ would no longer be obligatorily irrelevant, and so instead of deriving as the basic meaning of (3) $K[> 2]$, we'd incorrectly derive the stronger meaning $K[> 2] \wedge I[= 3] \wedge I[> 3]$. Adopting the UDM hypothesis forces these alternatives to be irrelevant, because density in this case removes symmetry.

By contrast, when it comes to *at least*, as in (2), density fails to remove the symmetry between the $[= 2]$ and $[> 2]$ alternatives, relative to $[\geq 2]$, and so, even with the UDM hypothesis, we still correctly derive as the meaning of (2) $K[\geq 2] \wedge I[= 2] \wedge I[> 2]$.

Class A/B distinction. More generally, we correctly predict the class A/B distinction discussed by Nouwen (2010) purely as a consequence of the ordering relations introduced by class A vs. class B modifiers. Class B modifiers, like *at least*, *at most*, *maximally*, and so on, introduce non-strict relations like ≥ and ≤, which, given reasonable assumptions about the formal alternatives they trigger, give rise to symmetry, hence to partial ignorance. Class A modifiers, by contrast, like *more than*, *fewer than*, and so on, introduce strict relations like > and <, which, together with density, fail to give rise to symmetry, hence fail to yield ignorance.

Double quantifier sentences. In section §5.1, we assumed, following Fox and Hackl (2006), that a sentence with unembedded *more than* like (3) can't be parsed with *exh*, since doing so leads to a contradiction. In the context of our amended proposal, however, we pointed out that (3) could be in principle parsed with *exh*, simply because it has no non-weaker alternatives to operate on (it's vacuous), due to obligatory irrelevance. We now show that this latter possibility is a reality.

Consider the following two sentences involving two quantifiers each (*at least/ more than* and *some*). As Schwarz (2016, fn. 12) observes, such sentences appear to license scalar inferences about *some*. For instance, (43a) implicates that it's not the case that more than two students were interested in all of the topics, while (43b) implicates that it's not the case that at least two students were interested in all of the topics.

(43) a. More than two students were interested in some of the topics.

   b. At least two students were interested in some of the topics.

To derive these implicatures, we need to assume that the sentences are parsed with an occurrence of *exh* below K, and we need it to be the case that the alternatives *more than two ... all ...* ($[> 2, \forall]$) and *at least two ... all ...* ($[\geq 2, \forall]$) can be relevant (aren't obligatorily irrelevant), so that this *exh* can exclude them.

(44) a. exh $[_{S_1}$ K exh $[_{S_2}$ more than two ... some ... $]$ $]$

   b. exh $[_{S_1}$ K exh $[_{S_2}$ at least two ... some ... $]$ $]$

Crucially, if we maintain our assumptions about the formal alternatives induced by *more than two* and *at least two*, and in addition assume that *some* has *all* as a formal alternative, then it turns out that $[> 2, \forall]$ can be relevant to (43a) and that $[\geq 2, \forall]$ can be relevant to (43b) (and the same for greater-numbered alternatives of those forms), and the reason for this is simple: they're both innocently excludable, hence settled when relevant. Note that the existential quantifier alternatives, $[> n, \exists]$ and $[\geq n, \exists]$, for $n > 2$, are all obligatorily irrelevant to (43a) and (43b), respectively, for the same reasons as before.

Importantly, we note that the original account of Fox and Hackl (2006) can't

derive a scalar inference for (43a) (they don't discuss *at least*): if *exh* applies, it yields a contradiction, and so it simply can't apply. On our proposal, contradiction is avoided because obligatory irrelevance eliminates from consideration the alternatives whose exclusion would yield a contradiction, and for those that remain, some are IE, and they yield precisely the scalar inferences that are attested.

Comparison with a neo-Gricean account.   We now briefly argue that our account has several advantages over a neo-Gricean account. First, we note that it's possible within a neo-Gricean framework to derive the partial ignorance inferences associated with *at least*, either by taking the formal alternatives of *at least two* to be just *exactly two* and *more than two* (Büring 2008; Kennedy 2013, 2015), or by taking them to be all *at least n* and *exactly n* alternatives (Schwarz 2016). (The latter case requires that the maxim of quantity make reference to innocent exclusion.) In both cases, primary inferences of the form $\neg K\phi$ are derived, such as $\neg K[= 2]$ and $\neg K[> 2]$ or $\neg K[\geq 3]$ in the case of (2) (*Ann owns at least two dogs*), which, together with the quality inference of (2), $K[\geq 2]$, entail speaker ignorance about $[= 2]$ and about $[> 2]$. Moreover, none of these primary $\neg K\phi$ inferences can be strengthened into secondary inferences of the form $K\neg\phi$ without yielding a contradiction, and hence no scalar implicatures arise.

However, these accounts run into trouble when it comes to sentences with *more than*, like (3) (*Ann owns more than two dogs*). If density isn't assumed, then either *more than two* triggers only other *more than n* alternatives, or it triggers both *more than n* and *exactly n* alternatives, where *n* is an integer. (Recall that it must trigger at least the *more than n* alternatives in order to account for the scalar inferences we observe when *more than two* is embedded under a universal modal.) In the former case, the scalar implicature $K\neg[> 3]$ is incorrectly predicted: the primary inference $\neg K[> 3]$ can be strengthened without a problem. In the latter case, $\neg K[> 3]$ can't be strengthened, but together with the inference $\neg K[= 3]$, ignorance about $[= 3]$ and $[> 3]$ is incorrectly predicted.

Perhaps, then, we can adopt the UDM hypothesis and try to recast the account of Fox and Hackl (2006) in neo-Gricean terms. For instance, assume that *more than two* triggers the dense set of *more than n* alternatives. Already at this point, it's unclear that a post-semantic, pragmatic algorithm would be sensitive to such alternatives: at this level of reasoning, it's plausible that granularity considerations (which render alternatives like *Ann owns exactly 2.5 dogs* contextually irrelevant, since they yield a contextual contradiction) are already in play. Nevertheless, let's grant that the dense set of alternatives are visible to the neo-Gricean algorithm. Then it derives the primary implicatures $\neg K[> n]$ for all $n > 2$, but they can't be strengthened into secondary implicatures without yielding a contradiction (even if the neo-Gricean maxim of quantity refers to innocent exclusion, since, as we

already observed, all these alternatives are IE). At this point, one may argue that, because secondary strengthening yields a contradiction, it simply doesn't happen. There are two problems this approach faces.

First, we saw in the case of the double quantifier sentences that *more than two* (and *at least two*) don't necessarily suppress all scalar implicatures, if another quantifier, like *some*, is present. If the step of strengthening from primary to secondary implicatures doesn't occur, then this fact remains unexplained. To be sure, many subsets of primary implicatures can be strengthened consistently, but there's no principled way, on this approach, to determine precisely the subset that is strengthened to yield the attested inferences.

Second, even for basic, single-quantifier sentences like (3), we would still be left with many primary inferences of the form $\neg K[> n]$, for $n > 2$. As we argued in §5.3, however, even these inferences appear to be unattested.

Finally, if our arguments in this article are correct, then some notion of obligatory irrelevance, i.e. some non-stipulative restriction on which alternatives can be relevant to an uttered sentence, is necessary for explaining the range of ignorance inferences observed in natural language. One wonders, then, whether such a notion can be devised within a neo-Gricean framework. We're doubtful that this is possible. For instance, the move to close relevance under belief is not an option for the neo-Gricean theory: doing so entails moving away from a neo-Gricean view to a grammatical view of both scalar implicatures and ignorance inferences. The reason for this is that closing relevance under belief entails adopting matrix K, and this in turn entails adopting *exh*, in order to avoid deriving a contradiction in contexts where the maxim of quantity is active.

Contexts where the maxim of quantity is inactive. In this article, we've only considered contexts where the maxim of quantity is active. In such contexts, our proposal states that every sentence must either settle its relevant alternatives or entail ignorance about them. One wonders, then, what predictions our account makes for contexts where the maxim of quantity is inactive, such as the game show contexts discussed in Fox 2014. Specifically, consider a game show, where utterances by the host don't obey the maxim of quantity. There are 100 boxes of different colors, and five of the 100 boxes contain a million dollars each (the rest are empty). The host gives hints to the players without revealing exactly which boxes contain money.

Fox (2014) observes that, in such a context, (45) fails to imply that the host is ignorant about whether box 20 contains money and about whether box 25 contains money, yet it does imply that boxes 20 and 25 don't both contain money. His point is that we can dissociate the computation of scalar inferences from that of ignorance inferences.

(45)   There is money in box 20 or 25.

On a proposal like ours, which assumes matrix *K*, this indicates that (45) has an LF like (46), with *exh* but without any matrix *K*. In other words, matrix *K* serves to make an utterance obey the maxim of quantity, so if the maxim of quantity is inactive, then it has no role to play. Note, by the way, that we can still derive the inference that the speaker believes the content of her utterance because we still assume that the maxim of quality is active.

(46)   exh [there is money in box 20 or 25]

Consider now (47). In a game show context, this sentence isn't parsed with matrix *K*, and so we instead have one of the two parses below. These two LFs are truth-conditionally equivalent because the *exh* in LF_1 is vacuous: no alternative is IE, for the same reasons discussed in §5. As a result, (47), uttered by a host in a game show context, is predicted just to mean that there is money in two or more green boxes, without any speaker ignorance, which is correct.[30]

(47)   There is money in at least two green boxes.
       LF_1: exh [there is money in at least two green boxes]
       LF_2: [there is money in at least two green boxes]

Similarly, (48) would have the parse in LF_2, again a simple quantificational statement with no speaker ignorance. Note that in this case, LF_1 is ruled out because it's a contradiction: for the same reasons discussed in §5, every alternative is IE, but excluding them all contradicts the complement of *exh*. (Note that the maxim of quantity, being inactive, doesn't render any of these alternatives obligatorily irrelevant, hence doesn't rescue the LF from contradiction.)

(48)   There is money in more than two green boxes.
       LF_1: exh [there is money in more than two green boxes]
       LF_2: [there is money in more than two green boxes]

9   Appendix: Proof

Let Σ be a sentence of the form *exh K (exh) S*, and let *A* be a set of propositions. We wish to prove that if Σ either settles or entails ignorance about every proposition in *A*, then Σ either settles or entails ignorance about every proposition in the

---

30 This explanation seems, at least on the surface, to be incompatible with Spector (2015)'s account of the global PPI status of *at least*. He argues that *at least* must occur in the scope of *exh*, which rules out LF_2, and that the application of *exh* be non-vacuous, which rules out LF_1. His prediction, then (even apart from our own theory), is that *at least* can't be used in contexts where the maxim of quantity is inactive. According to our judgments, this prediction is incorrect.

Draft: July 31, 2017

closure of A under conjunction, negation, and belief, as described by the closure conditions on relevance in (27). Assume, then, that for every $\phi \in A$, Σ settles or entails ignorance about ϕ.

Case 1 (closure under conjunction). Let $\phi, \psi \in A$. Then we need to show that Σ either settles or entails ignorance about $\phi \wedge \psi$. If Σ settles both ϕ and ψ, then Σ either entails $\phi \wedge \psi$ (if Σ entails both) or entails the negation of $\phi \wedge \psi$ (if Σ entails the negation of at least one of them), hence settles $\phi \wedge \psi$. If Σ entails ignorance about both ϕ and ψ, then Σ also entails ignorance about $\phi \wedge \psi$. Finally, without loss of generality, assume that Σ settles ϕ and entails ignorance about ψ. Then, if Σ settles ϕ by entailing it, then Σ entails ignorance about $\phi \wedge \psi$, and if Σ settles ϕ by entailing its negation, then Σ settles $\phi \wedge \psi$ by entailing its negation.

Case 2 (closure under negation). Let $\phi \in A$. If Σ settles ϕ, then Σ also settles $\neg\phi$, and if Σ entails ignorance about ϕ, then Σ also entails ignorance about $\neg\phi$.

Case 3 (closure under belief). Let $\phi \in A$. We must show that Σ either settles or entails ignorance about $K\phi$. We proceed by induction on the form of ϕ and show that Σ settles $K\phi$.

- **Subcase 0 (base case)** If ϕ doesn't contain an occurrence of K, then, if Σ settles ϕ, either ϕ or $\neg\phi$ is an entailment of Σ (by **T**) and so Σ either entails $K\phi$ or entails $K\neg\phi$ (which entails $\neg K\phi$), hence settles $K\phi$; and if Σ entails ignorance about ϕ, then Σ entails $\neg K\phi$, hence settles $K\phi$.

- **Subcase 1 (closure under conjunction)** If ϕ is of the form $\psi \wedge \chi$, and if Σ settles $K\psi$ and settles $K\chi$, then Σ settles $K\phi$. (If Σ entails both $K\psi$ and $K\chi$, then it entails, hence settles, $K\phi$; and if Σ entails the negation of at least one of $K\psi$ or $K\chi$, then it entails the negation of $K\phi$, hence settles it.)

- **Subcase 2 (closure under negation)** If ϕ is of the form $\neg\psi$ and if Σ settles $K\psi$, then Σ either entails $K\psi$, which entails $\neg K\phi$, hence settles $K\phi$, or Σ entails $\neg K\psi$, hence $\neg K\neg\phi$, and since by hypothesis Σ either settles ϕ or entails ignorance about ϕ ($\phi \in A$), we have that either Σ entails, hence settles, $K\phi$, or Σ entails speaker ignorance about ϕ, hence entails $\neg K\phi$, hence settles $K\phi$.

- **Subcase 3 (closure under belief)** If ϕ is of the form $K\psi$ and Σ settles $K\psi$, then Σ either entails $K\psi$, which (by positive introspection) entails, hence settles, $K\phi$, or it entails $\neg K\psi$, which (by negative introspection) entails $K\neg\phi$, which in turn entails $\neg K\phi$, hence settles $K\phi$.

QED

References

Buccola, Brian, Manuel Križ, and Emmanuel Chemla (2016). Conceptual Alternatives. URL: http://ling.auf.net/lingbuzz/003208.

Büring, Daniel (2008). The Least *at least* Can Do. In: *West Coast Conference on Formal Linguistics (WCCFL)*. Ed. by Charles B. Chang and Hannah J. Haynie. Vol. 26. Somerville, MA: Cascadilla Proceedings Project, pp. 114–120. URL: http://www.lingref.com/cpp/wccfl/26/paper1662.pdf.

Chierchia, Gennaro (2004). Scalar Implicatures, Polarity Phenomena, and the Syntax/Pragmatics Interface. In: *The Cartography of Syntactic Structures*. Ed. by Adriana Belletti. Vol. 3. Structures and Beyond. Oxford, UK: Oxford University Press, pp. 39–103.

Chierchia, Gennaro (2006). Broaden Your Views: Implicatures of Domain Widening and the "Logicality" of Language. In: *Linguistic Inquiry* 37.4, pp. 535–590. DOI: 10.1162/ling.2006.37.4.535.

Chierchia, Gennaro, Danny Fox, and Benjamin Spector (2012). Scalar Implicature as a Grammatical Phenomenon. In: *Semantics: An International Handbook of Natural Language Meaning*. Ed. by Claudia Maienborn, Klaus von Heusinger, and Paul Portner. Vol. 3. Berlin, Germany: Mouton de Gruyter, pp. 2297–2331.

Cohen, Ariel and Manfred Krifka (2014). Superlative Quantifiers as Meta-Speech Acts. In: *Linguistics and Philosophy* 37.1, pp. 41–90. DOI: 10.1007/s10988-014-9144-x.

Coppock, Elizabeth and Thomas Brochhagen (2013). Raising and Resolving Issues with Scalar Modifiers. In: *Semantics and Pragmatics* 6.3, pp. 1–57. DOI: 10.3765/sp.6.3.

Cummins, Chris, Uli Sauerland, and Stephanie Solt (2012). Granularity and Scalar Implicature in Numerical Expressions. In: *Linguistics and Philosophy* 35.2, pp. 135–169. DOI: 10.1007/s10988-012-9114-0.

Fox, Danny (2007). Free Choice and the Theory of Scalar Implicatures. In: *Presupposition and Implicature in Compositional Semantics*. Ed. by Uli Sauerland and Penka Stateva. Palgrave Studies in Pragmatics, Language and Cognition Series. New York, NY: Palgrave Macmillan. Chap. 4, pp. 71–120. DOI: 10.1057/9780230210752_4.

Fox, Danny (2014). Cancelling the Maxim of Quantity: Another Challenge for a Gricean Theory of Scalar Implicatures. In: *Semantics and Pragmatics* 7.5, pp. 1–20. DOI: 10.3765/sp.7.5.

Fox, Danny (2016). Topics in Exhaustification. Mini-course taught at the Hebrew University of Jerusalem. Jerusalem, Israel.

Fox, Danny and Martin Hackl (2006). The Universal Density of Measurement. In: *Linguistics and Philosophy* 29.5, pp. 537–586. DOI: 10.1007/s10988-006-9004-4.

Fox, Danny and Roni Katzir (2011). On the Characterization of Alternatives. In: *Natural Language Semantics* 19.1, pp. 87–107. DOI: 10.1007/s11050-010-9065-3.

Gajewski, Jon (2009). Innocent Exclusion is not Contradiction Free. Manuscript. URL: http://jon-gajewski.uconn.edu/wp-content/uploads/sites/1784/2016/08/IEnotCF.pdf.

Geurts, Bart and Rick Nouwen (2007). *At least* et al. The Semantics of Scalar Modifiers. In: *Language* 83.3, pp. 533–559. URL: http://www.jstor.org/stable/40070901.

Haida, Andreas and Tue Trinh (2016). Beyond Density. Talk presented at the Exh workshop, MIT. Cambridge, MA. URL: https://people.uwm.edu/tuetrinh/files/2016/11/HO-trinh2016beyond-2kdl7aa.pdf.

Heim, Irene (2006). *Little*. In: *Semantics and Linguistic Theory (SALT)*. Vol. 16, pp. 35–58. DOI: 10.3765/salt.v16i0.2941.

Horn, Laurence R. (1972). On the Semantics of Logical Operators in English. PhD thesis. New Haven, CT: Yale University.

Katzir, Roni (2007). Structurally-Defined Alternatives. In: *Linguistics and Philosophy* 30.6, pp. 669–690. DOI: 10.1007/s10988-008-9029-y.

Kennedy, Chris (2013). A Scalar Semantics for Scalar Readings of Number Words. In: *From Grammar to Meaning: The Spontaneous Logicality of Language*. Ed. by Ivano Caponigro and Carlo Cecchetto. Cambridge, England: Cambridge University Press, pp. 172–200. DOI: 10.1017/cbo9781139519328.010.

Kennedy, Chris (2015). A "de-Fregean" Semantics (and neo-Gricean Pragmatics) for Modified and Unmodified Numerals. In: *Semantics and Pragmatics* 8.10, pp. 1–44. DOI: 10.3765/sp.8.10.

Krifka, Manfred (1999). At Least Some Determiners aren't Determiners. In: *The Semantics/Pragmatics Interface from Different Points of View*. Ed. by Ken Turner. Vol. 1. New York, NY: Elsevier, pp. 257–291.

Liebesman, David (2016). Counting as a Type of Measuring. In: *Philosophers' Imprint* 16.12, pp. 1–25. URL: http://hdl.handle.net/2027/spo.3521354.0016.012.

Mayr, Clemens (2013). Implicatures of Modified Numerals. In: *From Grammar to Meaning: The Spontaneous Logicality of Language*. Ed. by Ivano Caponigro and Carlo Cecchetto. Cambridge, England: Cambridge University Press, pp. 139–171.

Meyer, Marie-Christine (2013). Ignorance and Grammar. PhD thesis. Cambridge, MA: Massachusetts Institute of Technology. URL: http://hdl.handle.net/1721.1/84420.

Nouwen, Rick (2010). Two Kinds of Modified Numerals. In: *Semantics and Pragmatics* 3.3, pp. 1–41. DOI: 10.3765/sp.3.3.

Nouwen, Rick (2015). Modified Numerals: The Epistemic Effect. In: *Epistemic Indefinites: Exploring Modality Beyond the Verbal Domain*. Ed. by Luis Alonso-Ovalle and Paula Menéndez-Benito. Oxford Linguistics. Oxford, England: Oxford University Press. Chap. 11, pp. 244–266.

Penka, Doris (2014). *At most* at Last. In: *Sinn und Bedeutung*. Ed. by Eva Csipak and Hedde Zeijlstra. Vol. 19, pp. 463–480. URL: http://nbn-resolving.de/urn:nbn:de:bsz:352-0-304168.

Sauerland, Uli (2004). Scalar Implicatures in Complex Sentences. In: *Linguistics and Philosophy* 27.3, pp. 367–391. DOI: 10.1023/B:LING.0000023378.71748.db.

Schwarz, Bernhard (2016). Consistency Preservation in Quantity Implicature: The Case of *at least*. In: *Semantics and Pragmatics* 9.1, pp. 1–47. DOI: 10.3765/sp.9.1.

Spector, Benjamin (2015). Why are Class B Modifiers Global PPIs?: Disjunctions as a Model for Class B Modifiers. Talk presented at the Workshop *New Directions in Negation and Polarity*, The Hebrew University of Jerusalem. Jerusalem, Israel.

Trinh, Tue and Andreas Haida (2015). Constraining the Derivation of Alternatives. In: *Natural Language Semantics* 23.4, pp. 249–270. DOI: 10.1007/s11050-015-9115-y.

```
Word count: 13343 (texcount, without bibliography)
           17634 (Monterey PDF word count service)
```