

Obligatory irrelevance and the computation of ignorance inferences

Brian Buccola¹ and Andreas Haida²

¹Laboratoire de Sciences Cognitives et Psycholinguistique (ENS, EHESS, CNRS),
Département d'Études Cognitives, École Normale Supérieure, Université PSL

²Edmond and Lily Safra Center for Brain Sciences, The Hebrew University of Jerusalem

Abstract

In recent work, Fox (2016) has argued, on the basis of both empirical and conceptual considerations, that relevance (the set of propositions relevant in an utterance context) is closed under speaker belief: if ϕ is relevant, then it's also relevant whether the speaker believes ϕ . We provide a formally explicit implementation of this idea and explore its theoretical consequences and empirical predictions. As Fox (2016) already observes, one consequence is that ignorance inferences (and scalar implicatures) can only be derived in grammar, via a covert belief operator of the sort proposed by Meyer (2013). We show, further, that the maxim of quantity 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 show that the resulting system predicts a quite restricted range of interpretations for sentences with the scalar item *some*, as compared to both neo-Gricean (Horn 1972; Sauerland 2004; Geurts 2010) and grammatical (Chierchia, Fox, and Spector 2012; Fox 2007; Meyer 2013) theories of scalar implicature, and we argue that these predictions seem largely on the right track.

Keywords: ignorance inferences, scalar implicatures, exhaustivity, relevance, Gricean maxims

1 Introduction

According to the standard grammatical approach to scalar implicatures, as developed by Chierchia (2004), Fox (2007), Chierchia, Fox, and Spector (2012), and Fox (2014), scalar implicatures are only derived in grammar, while 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 potentially be quite pervasive, for the following reason: given an utterance S , if ϕ is a relevant alternative of S such that S doesn't entail ϕ nor $\neg\phi$ (that is, S doesn't settle ϕ), then the maxim of quantity licenses the inference that the speaker doesn't have the belief that ϕ is true; but on standard assumptions about relevance (viz. that it's closed under negation), $\neg\phi$ is also relevant, and since S doesn't settle ϕ , S doesn't settle $\neg\phi$ either, so S also licenses the inference that the speaker doesn't have the belief that $\neg\phi$ is true (i.e. that ϕ is false). In other words, the speaker doesn't have a belief one way or the other about the truth of ϕ : she's ignorant about ϕ . The result is that, for any relevant alternative ϕ of an utterance S which S doesn't settle, the maxim of quantity licenses an inference of speaker ignorance about ϕ . Therefore, without any principled restrictions on relevance, ignorance inferences are predicted to potentially be quite pervasive.

In recent work, Fox (2016) has argued that relevance is closed not just under negation (and conjunction), but also under speaker belief: if ϕ is relevant, then it's also relevant whether the speaker believes ϕ . This closure condition is motivated by the observation that "silence is uncooperative": if ϕ is relevant, but the speaker doesn't have a belief one way or the other about the truth of ϕ , then she is obligated to say so, which suggests that her beliefs about ϕ are relevant. As Fox points out, a consequence of closing relevance under belief is that ignorance inferences, like scalar implicatures, can only be derived in grammar, by way of a covert belief operator (Chierchia 2006; Meyer 2013), not pragmatically. Put differently, the maxim of quantity no longer derives any additional inferences, i.e. doesn't enrich the meaning of an uttered sentence.

In this article, we take seriously Fox's proposal and explore its theoretical and empirical consequences in detail. The article is divided roughly into two halves. In the first half, we begin by developing a formally explicit semantic-pragmatic framework (§2) 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 (§3). (As far as we know, ours is the first formally explicit theory of this kind.) We then minimally extend this system by closing relevance under belief (§4). We show that the maxim of quantity then 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*.

In the second half of the article, we investigate whether we can find any empirical traces of obligatory irrelevance by examining textbook cases of sentences with the scalar expression *some*. We show that the new theory predicts a rather restricted range of interpretations for such sentences: they cannot generally be used to convey speaker ignorance about their *all* alternative; instead, they convey either that their *all* alternative is false, or that it's irrelevant. This prediction contrasts sharply with those of the standard neo-Gricean view of implicature (Sauerland 2004; Geurts 2010), the standard grammatical view of implicature (Chierchia, Fox, and Spector 2012), and the Matrix *K* theory of Meyer 2013, which we discuss and compare with in turn.

The system we explore here makes quite clear predictions, which, however, are at times tough to test introspectively. Nevertheless, our goal is to present them as is, with the hope that more sophisticated empirical tests can be conducted in the future. With that caveat in mind, we do find that the facts arguably point in favor of closing relevance under belief. If it turns out that this is correct, then these facts would also constitute another argument against a neo-Gricean approach to scalar implicature, and in favor of a grammatical approach—one that is even more radically grammatical than the standard one.

2 The standard grammatical theory of scalar implicature

We begin by laying out the ingredients that form the basis of the standard grammatical theory of scalar implicature, and we illustrate how it works for a textbook example involving disjunction.

2.1 The ingredients

On the grammatical view of scalar implicature, 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 $\text{alt}(S)$, the set of relevant formal alternatives of S , is false. Intuitively, a proposition $q \in \text{alt}(S)$ is IE (given $\llbracket S \rrbracket$) just in case the negation of q doesn't contradict $\llbracket S \rrbracket$ and doesn't force any disjunction of members of $\text{alt}(S)$ to be true, unless that disjunction is already entailed by $\llbracket S \rrbracket$. The formal semantics of *exh* is provided in (1).

- (1) **Definition of *exh*** (Chierchia, Fox, and Spector 2012)
- a. $\llbracket \text{exh } S \rrbracket = \llbracket \text{exh} \rrbracket (\text{alt}(S))(\llbracket S \rrbracket)$.
 - b. $\llbracket \text{exh} \rrbracket (A)(p) = p \wedge \bigwedge \{-q : q \in \text{IE}(p, A)\}$.

- c. $IE(p, A) = \bigcap \{M : M \text{ is a maximal set of excludable alternatives of } A \text{ given } p\}$.
- d. $M \subseteq A$ is a set of excludable alternatives of A given p iff $p \wedge \bigwedge \{-q : q \in M\}$ 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 .

(2) **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 $alt_c^{FA}(S)$ 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:

(3) **Alternatives of a sentence (as used by *exh*)**

For any sentence S , $alt_c(S) = \{\llbracket S' \rrbracket : S' \in alt_c^{FA}(S)\} \cap relevant_c$.

We almost always omit the subscript c , like in the definition of *exh* in (1).

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).^{1,2} Fox (2016) refers to the two maxims together as “virtual truisms”, because we seem to be aware of them without explicitly being taught them.³

¹ 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” (see fn. 3). 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”.

² 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 ...”

³ 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.

(4) **Maxim of quality**

The speaker should only utter sentences that the speaker believes to be true. (“Tell nothing but the truth.”)

(5) **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 (6) 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.

(6) **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 the theory) be tied to the derivation of ignorance inferences. We characterize ignorance about a proposition ϕ as in (7).

(7) **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.)

For our semantic representation language, we adopt an epistemic logic containing the modal operator K , where $K\phi$ means that the speaker believes the proposition ϕ (cf. Hintikka 1962). 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 $\neg K\phi \wedge \neg K\neg\phi$.

2.2 Example: Disjunction

To illustrate, let’s take the classic case of disjunction, as in (8). Intuitively, (8) 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.

(8) Ann owns a Boxer or a Collie.

LF: $\text{exh } [_S \text{ Ann owns a Boxer or (Ann owns) a Collie }]$

What we wish to show is that parsing (8) with exh leads to just the right scalar and ignorance inferences. For the sake of clarity, when we write “(8)”, we refer to the LF provided below (8), containing exh , and we write “ S ” to refer to the complement of exh , also known as the *prejacent*. The meaning of S is thus $b \vee c$,

the proposition that Ann owns a Boxer or a Collie (or both). We take $\text{alt}(S)$ to be $\{b \vee c, b, c, b \wedge c\}$, since each of these is (denoted by) a formal alternative of S (Sauerland 2004) and is presumably relevant. Then the meaning of (8) 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 (8) is $(b \vee c) \wedge \neg(b \wedge c)$.

(9) **Semantic meaning of (8):**

$$(b \vee c) \wedge \neg(b \wedge c)$$

Next, based on the maxim of quality, a hearer is licensed to infer that the speaker of (8) believes the semantic content of (8). Thus, (8) licenses the quality inference $K((b \vee c) \wedge \neg(b \wedge c))$.

(10) **Quality inference of (8):**

$$K((b \vee c) \wedge \neg(b \wedge c))$$

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 (8), 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 (8), the maxim of quantity also licenses the inferences $\neg K\neg b$ and $\neg K\neg c$.

(11) **Quantity inferences of (8):**

$$\neg Kb \wedge \neg K\neg b \wedge \neg Kc \wedge \neg K\neg c$$

Importantly, the pair of inferences $\neg Kb$ and $\neg K\neg b$ together amount to speaker ignorance about b , as characterized by (7). 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 (8) given below.

(12) **Semantic and pragmatic meaning of (8):**

$$\underbrace{K((b \vee c) \wedge \neg(b \wedge c))}_{\text{Quality inference}} \wedge \underbrace{\underbrace{\neg Kb \wedge \neg K\neg b}_{\text{Ignorance about } b} \wedge \underbrace{\neg Kc \wedge \neg K\neg c}_{\text{Ignorance about } c}}_{\text{Quantity inferences}}$$

3 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 (8), 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, (8) 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. The deductive procedure in (13) 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 .

(13) **Pragmatic reasoning procedure**

For any context c , uttered sentence S , and proposition ϕ :

- a. $\text{relevant}_c(\phi), \llbracket S \rrbracket \neq \phi, \llbracket S \rrbracket \neq \neg\phi$ (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. $I\phi$ ((b) + (d) + def. of ignorance)

If, following Fox 2016, we introduce the notion of *settlement*, given in (14), then we can state our generalization as in (15).

(14) **Definition of settlement**

S settles ϕ just in case $\llbracket S \rrbracket \models \phi$ or $\llbracket S \rrbracket \models \neg\phi$.

(15) **Generalization about the pragmatic reasoning procedure in (13)**

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, (8) gives rise to inferences of speaker ignorance about the individual disjunct alternatives b and c because (8) 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.

4 Closing relevance under belief

We now minimally extend the standard grammatical 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).⁴ The third 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, leading to a notion we call *obligatory irrelevance*, whose (potential) empirical footprint we explore in §5.

4.1 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 .

(16) **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 $K\phi$.

Fox (2016) motivates this move by observing 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.”

We note that this observation is not limited to highly formal or rigid legal contexts, but rather holds across everyday conversational contexts, too. For example, if Ann asks Bob, “Where is John?”, and Bob doesn’t know where John is, Bob cannot simply remain silent; he must express his ignorance to Ann in some way.

⁴ Fox (2016) has observed these first two points.

As Fox (2016) informally demonstrates, if relevance is closed under belief, then ignorance inferences are no longer pragmatic, but rather must be derived in grammar. To see this, suppose that Ann utters a sentence S , and that ϕ is any relevant alternative not settled by S . Then, as a consequence of the earlier generalization in (15), Ann must be ignorant about ϕ . But in that case, Ann must also convey her ignorance about ϕ , since that information is both relevant and available to Ann. Thus, S must entail ignorance about ϕ .

We now formalize this result by making explicit the underlying logic necessary for it to obtain.

4.2 A generalization about ignorance as semantic entailments

Formally, we assume an S5 modal logic, which contains the axiom schemata in (17).

(17) **S5 epistemic modal logic**

- | | | |
|----|---|--|
| a. | $(K\phi \wedge K(\phi \rightarrow \psi)) \rightarrow K\psi$ | <i>distribution axiom</i> (K) |
| b. | $K\phi \rightarrow \phi$ | <i>truth axiom</i> (T) |
| c. | $K\phi \rightarrow KK\phi$ | <i>positive introspection axiom</i> (4) |
| d. | $\neg K\phi \rightarrow K\neg K\phi$ | <i>negative introspection axiom</i> (5) |

We take it that this logic appropriately models that component of the mind that governs speaker belief in grammar and semantics. (We don't necessarily take this logic to be an appropriate model of human belief or knowledge more generally.)

The **K** axiom essentially represents our implicit knowledge of the logical rule of modus ponens.⁵ We assume that speakers know certain basic logical validities like $(\phi \wedge \psi) \rightarrow \phi$. This means that from $K(\phi \wedge \psi)$ we may conclude $K\phi$ on the basis of **K**.⁶

The **T** axiom corresponds to what Schwarz (2016) calls "bottom line inferences" and effectively represents a hearer's update of the common ground (or of their own knowledge state) on the basis of what the speaker uttered.

Finally, axioms **4** and **5** formalize the notion that we know what our knowledge state is (about a given proposition ϕ). These axioms are important for a speaker to be able to obey the maxim of quality. Without **4**, belief in ϕ would not suffice for being able to utter ϕ (because the speaker wouldn't necessarily know that she believes ϕ); and without **5**, a lack of belief in ϕ wouldn't prevent the speaker from

⁵ Logical reasoning with modus ponens has been shown to be relatively easy for humans, as compared to, say, modus tollens, which is much harder (Evans, Newstead, and Byrne 1993).

⁶ As far as we can tell, the **K** axiom is not actually required for any of the deductions that follow in this article, but we list it here because **K** is valid in all normal modal logics, and because S5 is a logic commonly employed in epistemic modal logic.

uttering ϕ (because the speaker wouldn't necessarily know that she doesn't believe ϕ). Together they also entail that a speaker knows what she's ignorant about.

We now demonstrate formally that, as a consequence of closing relevance under belief, together with the maxim of quantity, ignorance can only be derived in grammar.

(18) **Pragmatic reasoning procedure** (after closure of relevance under K)

For any context c , uttered sentence S , and proposition ϕ :

- a. $\text{relevant}_c(\phi), \llbracket S \rrbracket \neq \phi, \llbracket S \rrbracket \neq \neg\phi$ (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. $\llbracket S \rrbracket \models \neg K\phi \wedge \neg K\neg\phi$ or $\llbracket S \rrbracket \neq \neg K\phi \wedge \neg K\neg\phi$ (tautology)
- h. Case 1: $\llbracket S \rrbracket \models \neg K\phi \wedge \neg K\neg\phi$
 $\llbracket S \rrbracket \models I\phi$ (above line + def. of ignorance)
- i. Case 2: $\llbracket S \rrbracket \neq \neg K\phi \wedge \neg K\neg\phi$
 - Subcase 1: $\llbracket S \rrbracket \neq \neg K\phi$
 $\neg K\neg K\phi$ ((e) + above line + maxim of quantity)
 $K\neg K\phi$ ((b) + axiom 4)
 contradiction (above two lines)
 - Subcase 2: $\llbracket S \rrbracket \neq \neg K\neg\phi$
 $\neg K\neg K\neg\phi$ ((f) + above line + maxim of quantity)
 $K\neg K\neg\phi$ ((d) + axiom 5)
 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.⁷ 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

⁷ A reviewer suggests the following intuitive way to think about the consequence of closing relevance under belief. Given an utterance S , if ϕ is relevant, then by the closure conditions of relevance, so is $I\phi$. Thus, either S settles $I\phi$, or else the maxim of quantity licenses the inference $II\phi$. The latter case cannot hold under plausible assumptions about epistemic states, viz. that speakers know what they know and what they are ignorant about (see Case 2 of (18), as well as the preceding discussion of axioms 4 and 5). So, S must settle $I\phi$. If, furthermore, S does not settle ϕ itself, then S cannot entail $\neg I\phi$, and so S must entail $I\phi$.

contradictory quantity inference.⁸ Consequently (and given the obvious fact that things can be relevant that speakers are ignorant about), 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. Thus, the pertinent generalization can be stated as in (19).

(19) **Generalization about the pragmatic reasoning procedure in (18)**

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 $\llbracket S \rrbracket \models I\phi$.

Since speakers need to be able to convey their ignorance grammatically (and since not every sentence begins with *I (don't) know, I'm (un)certain*, etc.), language evidently 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 $\llbracket K \rrbracket \phi = K\phi$.

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 convey ignorance grammatically.^{9, 10}

⁸ 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 either perceive the contradiction directly (viz. that the speaker is expressing a contradictory belief state) and/or detect some oddness similar perhaps to the cases of oddness arising from contradiction discussed elsewhere in the literature. For example (to foreshadow a bit), if Bob is interested in knowing whether or not Ann is Canadian (and explicitly asks Carl this), and Carl says to Bob, *It's raining in Paris*, Carl's reply is certainly odd, but that's simply because he isn't addressing the question at hand; he's explicitly flouting or disobeying the maxims of quantity and relevance (for whatever reason). In no way is the oddity of the sentence on a par with that of sentences like *#Somebody but Kim called* (Fintel 1993) or *#There are any cookies left* (Chierchia 2013), whose oddity has been argued to arise due to their expressing a contradiction.

⁹ 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 (20), 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 observe ignorance inferences obligatorily/by default) must be parsed with matrix K . We think that weakening (20) in this way, and answering this question, could lead to deeper insights. However, we have to leave this inquiry for a future occasion, and so we instead adopt what is given in (20).

¹⁰ In contexts where the maxim of quantity is inactive, e.g. game show contexts, we take it that sentences are not parsed with K , but could still be parsed with *exh*. Thus, scalar implicatures are predicted to still arise, but ignorance implicatures are not — predictions that Fox (2014) argues are borne out.

(20) **Matrix K**

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

Given this, together with our assumption in (2) 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).

4.3 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 generalization in (19) is in fact just a stronger version of the generalization in (15) from §3, repeated below.

(15) **Generalization about the pragmatic reasoning procedure in (13)**

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 .

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.

(21) **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 .

We assume that all the sentences we consider are interpreted in contexts where the maxim of quantity is active and obeyed (not flouted) (but see fn. 10).

The condition in (21) restricts which propositions can be relevant in a specific utterance context in which the maxim of quantity is active. Crucially, 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*.

(22) **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 .

To give an extreme example, the utterance *Ann is Canadian* ($= S$) neither settles nor entails ignorance about the proposition ‘it is raining in Paris’ ($= \phi$). So, ϕ is

obligatorily irrelevant to S — which seems intuitive enough. (S is predicted to be assertable only in contexts where ϕ is not taken to be relevant.) Note that this result, and the definition in (22) more generally, relies on the assumption that settlement is defined in terms of logical, not contextual, entailment — for otherwise, ϕ would not (necessarily) be obligatorily irrelevant to S , since there may well be conceivable contexts in which S contextually settles or contextually entails ignorance about ϕ . We discuss the issue of logical vs. contextual entailment in further detail in §6, but we note right away that switching to contextual entailment does not change the predictions that the system makes concerning the textbook cases of *or* (§4.4) and *some* (§5).

In the second half of this article, we’ll explore whether less trivial examples can be found for empirical traces of obligatory irrelevance, in order to help us decide whether closure of relevance under belief yields an empirically adequate semantic-pragmatic system. But first, we show how the new system handles the classic case of disjunction.

4.4 Disjunction revisited

Let’s return to the classic case of disjunction in (8), 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). This involves checking that the LF settles or entails speaker ignorance about every relevant proposition ϕ .

(8) Ann owns a Boxer or a Collie.

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

We take a bottom-up approach and start by considering the meaning of $\text{exh } S_1$. As before, we take $\text{alt}(S_1)$ to be $\{b \vee c, b, c, b \wedge c\}$. $\text{exh } S_1$ is therefore precisely the same as what we derived in §2.2 as the semantic meaning of (8) itself on the old parse, with just one exh and no K : $(b \vee c) \wedge \neg(b \wedge c)$. Thus, the meaning of $K \text{exh } S_1$ is $K((b \vee c) \wedge \neg(b \wedge c))$.

The semantic meaning of (8) is therefore the meaning of $K \text{exh } S_1$, which we just computed, plus the conjunction of the negation of all IE alternatives of S_2 . Following Meyer 2013, we take it that $\text{alt}_c(S_2) = \{[[K \text{exh } S'] : S' \in \text{alt}_c^{FA}(S_1) \wedge [S'] \in \text{relevant}_c\}$, the set of alternatives obtained by replacing in S_2 any relevant alternative of S_1 .¹¹ (We abbreviate *Ann owns a Boxer* by “B” and *Ann owns a Collie* by “C”.)

¹¹ 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.

$$\begin{aligned}
 (23) \quad \text{alt}_c(S_2) &= \{ \llbracket K \text{ exh } S' \rrbracket : S' \in \text{alt}_c^{FA}(S_1) \wedge \llbracket S' \rrbracket \in \text{relevant}_c \} \\
 &= \{ \llbracket K \text{ exh } [B \text{ or } C] \rrbracket, \llbracket K \text{ exh } B \rrbracket, \llbracket K \text{ exh } C \rrbracket, \llbracket K \text{ exh } [B \text{ and } C] \rrbracket \} \\
 &= \{ K((b \vee c) \wedge \neg(b \wedge c)), K(b \wedge \neg c), K(c \wedge \neg b), K(b \wedge c) \}
 \end{aligned}$$

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 $K(b \wedge c)$ is already entailed by the meaning of $K S_1$, and so we omit it below.

$$(24) \quad \text{Semantic meaning of (8)} \\
 \underbrace{K((b \vee c) \wedge \neg(b \wedge c))}_{\text{Prejacent of matrix } \textit{exh}} \quad \wedge \quad \underbrace{\neg K(b \wedge \neg c) \wedge \neg K(c \wedge \neg b)}_{\text{Exclusions of matrix } \textit{exh}}$$

Crucially, $K((b \vee c) \wedge \neg(b \wedge c))$ and $\neg K(b \wedge \neg c)$ together entail both $\neg K b$ 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 $\neg K c$ and $\neg K \neg b$. These four inferences together ($\neg K b$, $\neg K \neg b$, $\neg K c$, 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 (8) is precisely what we derived earlier as the complete semantic and pragmatic meaning of (8) parsed with just one *exh* and no K . That is, the semantic meaning of (8) that we arrive at can be equivalently stated as in (25).¹² 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.

$$(25) \quad \text{Semantic meaning of (8)} \qquad \text{(equivalent version)} \\
 K((b \vee c) \wedge \neg(b \wedge c)) \wedge \neg K b \wedge \neg K \neg b \wedge \neg K c \wedge \neg K \neg c \qquad \text{(same as (12))}$$

The quality inference we derive is that the speaker knows the semantic meaning of (8), i.e. $K \llbracket (8) \rrbracket$. 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.¹³

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 additional inferences the maxim of quantity licenses, but rather whether the maxim of quantity can be active in a context where the structure in (8) 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

¹² We've shown that (24) entails (25). For the other direction, note that $\neg K b$ entails $\neg K(b \wedge \neg c)$, and that $\neg K c$ entails $\neg K(c \wedge \neg b)$.

¹³ More precisely, $K\phi \wedge \neg K\psi$ (the basic form of the semantic meaning of (8)) already entails $KK\phi \wedge K\neg K\psi$ in S_5 (by positive and negative introspection), which is equivalent to $K(K\phi \wedge \neg K\psi)$.

proposition ϕ , either (8) settles ϕ or entails ignorance about ϕ . This entails showing that for each alternative in $\{b \vee c, b, c, b \wedge c\}$, either (8) 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 (8) neither settles nor entails ignorance about that proposition, it means that the maxim of quantity would yield a contradiction, hence that (8), 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 quantity is active (a bad result); otherwise, we're in good shape.

We begin with $b \vee c$. Since (8) entails $K(b \vee c)$, it also entails $b \vee c$ by the **T** axiom. As a result, (8) settles $b \vee c$. Likewise, (8) entails $\neg(b \wedge c)$ (again by **T**), which means it also settles $b \wedge c$. Finally, as illustrated above, (8) entails speaker ignorance about the individual disjunct alternatives b and c . Thus, (8) settles or entails ignorance about every proposition in $\text{alt}(S_1)$.

Second, we need to show that for every proposition ϕ in the closure of $\text{alt}(S_1)$ under conjunction, negation, and belief, (8) settles ϕ or entails ignorance about ϕ . In the appendix (§A), we prove that if a sentence Σ of the form $\text{exh } K(\text{exh } S)$ (like (8) 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 $\text{alt}(S_2)$, the domain of the matrix-level exh , can be relevant.)

To recap, we've shown that, if (8) is parsed as $\text{exh } K \text{ exh } B \text{ or } C$, 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 (8) (cf. 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 (22).

5 The textbook case of *some*

Having established the formal result of closing relevance under belief — namely, the generalization in (19), which gives rise to the notion of obligatory irrelevance, defined in (22) — we now turn to the empirical question of whether we can find traces of obligatory irrelevance in natural language. To do so, we look at the textbook case of the scalar item *some*.

5.1 Predicted denotations of *some* sentences

We focus on a textbook example like (26), borrowed from Geurts 2010. Which readings of (26) does the current system predict to be available?

(26) Bonnie stole some of the pears.

First, we need to consider all the possible parses of (26). Since we assume that every sentence must be c-commanded by *K*, (20), and that *exh* may be adjoined to any propositional constituent, (2), this leaves us with four main options, depending on whether *exh* attaches above *K*, below *K*, neither, or both.¹⁴ For clarity, we write “some” to refer to the denotation of S_1 .

- (27) a. $K [_{S_1}$ Bonnie stole some of the pears]
 b. $exh [_{S_2}$ $K [_{S_1}$ Bonnie stole some of the pears]]
 c. $K exh [_{S_1}$ Bonnie stole some of the pears]
 d. $exh [_{S_3}$ $K exh [_{S_1}$ Bonnie stole some of the pears]]

We start with (27c), and we assume, uncontroversially, that (28) is a formal alternative of S_1 . (We write “all” to refer to its denotation.) As such, we assume that $\text{alt}(S_1) = \{\text{some}, \text{all}\}$, i.e. that all is also relevant, and then ask what interpretation we predict, and whether the maxim of quantity yields a contradiction.

(28) $[_{S_4}$ Bonnie stole all of the pears]

First, we note that all is IE (it has no symmetric partner), and so the overall denotation we get for (27c) is $K(\text{some} \wedge \neg \text{all})$. Second, we must ask whether (27c), given this denotation, settles or entails ignorance about every relevant proposition. Clearly, it settles both some and all (by the T axiom). (It also settles everything in the closure of $\{\text{some}, \text{all}\}$ under negation, conjunction, and belief, as per the result in §A.) Therefore, no contradiction arises: the maxim of quantity can be active, and the alternative all can be relevant. This is a good result, to be sure: this parse corresponds to the ‘strong’ implicature reading of (26), which is universally agreed to be available.

Next, we note that (27d) delivers the exact same reading as (27c), because the outer *exh* is semantically vacuous: the alternatives of S_3 are $K exh \dots \text{some} \dots$ (S_3 itself) and $K exh \dots \text{all} \dots$; the former is not IE (the prejacent of *exh* never is), and although the latter is IE, its negation is already entailed by S_3 . So, (27d) doesn’t represent a different reading from (27c).

Turning next to (27a), since there is no *exh* present, the only possible denotation is K some. This neither settles nor entails ignorance about all, no matter in which

¹⁴ Adding more occurrences of *K* or *exh* wouldn’t generate any further readings, since they would be semantically vacuous. For *K*, this is due to introspection (axioms 4 and 5). For *exh*, the lower *exh* already fully strengthens the meaning of the sentence. (More precisely, for (27c), *exh* already excludes the *all* alternative, so adding a second *exh* above or just below *K* would not exclude anything further; and for (27b), the *all* alternative is obligatorily irrelevant, as we’ll soon see, so adding a second *exh* would again not exclude anything further.)

	relevant(all)	¬relevant(all)
<i>(exh)</i> $K \dots \text{some} \dots$	\perp	$K \text{some}$
<i>(exh)</i> $K \text{exh} \dots \text{some} \dots$	$K(\text{some} \wedge \neg \text{all})$	$K \text{some}$

Table 1: Predicted denotations of (26), depending on parse and (ir)relevance of all. The symbol “ \perp ” means that the maxim of quantity yields a contradiction.

context the sentence is used.¹⁵ Therefore, all is obligatorily irrelevant to (27a), in the sense of (22). In other words, if all is taken to be relevant (and the maxim of quantity is active), then (27a) is not an available parse of (26), since it would lead to a contradiction; otherwise (if all is irrelevant), then (27a) is available and means just $K \text{some}$.

Finally, we turn to (27b). We first assume, again, that $\text{alt}(S_1) = \{\text{some}, \text{all}\}$, and then ask whether the maxim of quantity yields a contradiction. In this case, the denotation we end up with is $K \text{some} \wedge \neg K \text{all}$. Clearly, (27b), given this denotation, settles some; however, it neither settles nor entails ignorance (as characterized in (7)) about all. (Although it entails $\neg K \text{all}$, it doesn’t entail the possibility inference $\neg K \neg \text{all}$.) Therefore, if all were relevant, then the maxim of quantity would yield a contradiction. Thus, all is obligatorily irrelevant to (27b). In other words, if all is taken to be relevant, then (27b), like (27a), is not an available parse of (26); otherwise (if all is irrelevant), then (27b) is available and, like (27a), means just $K \text{some}$.

To recap, all can be relevant both to (27c) and to (27d), in which case the sentence denotes $K(\text{some} \wedge \neg \text{all})$. By contrast, all is obligatorily irrelevant both to (27a) and to (27b); hence, they each denote simply $K \text{some}$. This all means that if all is *not* taken to be relevant in a given context, then (26) can simply be parsed as (27a) (or (27b)), with a vacuous *exh*, yielding the simple meaning $K \text{some}$, which seems like a fine result. However, if all *is* taken to be relevant in a given context, then (26) must be parsed as (27c) (or (27d)), with a vacuous outer *exh*, yielding the strong implicature meaning $K(\text{some} \wedge \neg \text{all})$. These predictions are summarized in Table 1.

5.2 Three paths to ignorance

We just saw that if all is relevant, then (26) is predicted to only be usable to convey the strong implicature meaning $K(\text{some} \wedge \neg \text{all})$. Thus, if the speaker is ignorant

¹⁵ This claim rests on the assumption that the underlying notion of entailment (in terms of which settlement is defined) is logical, not contextual, entailment: if the underlying notion were contextual entailment, then in contexts where some and all are contextually equivalent, (27a) would indeed settle all. We discuss such contexts in §6.2, where we provide arguments for maintaining that the underlying notion of entailment is logical.

about all, then she shouldn't be able to use (26) (assuming still that all is relevant, and of course that the speaker is obeying the maxim of quality). This is because the speaker is compelled (by the maxim of quantity) to grammatically express her ignorance about all, which (26), no matter how it's parsed, cannot do. Before we explore whether this prediction (and the others) is correct, we first ask the following question: in such a scenario, what options are predicted to be available to the speaker, if not (26)? We mention here three options.

5.2.1 Disjunction: *Some or all*

The first option is to use the overt disjunction *some or all*.

(29) Bonnie stole some or all of the pears.

This is an example of a Hurford disjunction, which has been argued to involve embedded exhaustification of *some*. (Otherwise, there would be an entailment relation between the disjuncts, which Hurford's constraint disallows, as evidenced by the oddness of #*Bonnie went to France or Paris*.) Thus, (29) is parsed as (30).¹⁶

(30) $\text{exh } [S_4 \text{ K } [S_3 [\text{exh } [S_1 \text{ Bonnie stole some of the pears}]] \text{ or } [S_2 \dots \text{all} \dots]]]$

The denotation of the first disjunct, $\text{exh } S_1$, is thus $\text{some} \wedge \neg \text{all}$, while the denotation of the second disjunct, S_2 , is all. These two disjuncts exclude one another, which means that they form a symmetric pair relative to the whole disjunction. This in turn means that the outer *exh* yields ignorance about each disjunct, precisely like in run-of-the-mill cases of disjunction like (8). The denotation of (30) is given in (31); we write "some*" to abbreviate $\text{some} \wedge \neg \text{all}$.

(31)
$$\begin{aligned} \llbracket (30) \rrbracket &= \text{K}(\text{some}^* \vee \text{all}) \wedge \neg \text{K} \text{some}^* \wedge \neg \text{K} \text{all} \\ &\equiv \text{K}(\text{some}^* \vee \text{all}) \wedge \text{I} \text{some}^* \wedge \text{I} \text{all} \end{aligned}$$

Since (30) entails ignorance about all, all can be relevant to (30) (unlike what we saw for (27b)). Thus, if all is taken to be relevant, then (29) (specifically, parsed as (30)) is predicted to be one path for the speaker to convey her ignorance about all.

5.2.2 Dedicated expression for conveying ignorance: *At least some*

A second option is to use a dedicated expression for conveying ignorance, such as *at least*.

(32) Bonnie stole at least some of the pears.

¹⁶ We omit any occurrence *exh* just below *K*, with the entire disjunction as its complement, because that occurrence of *exh* would be vacuous.

The effect of *at least* is essentially to undo, or preempt, the scalar implicature that would arise in its absence (Krifka 1999), viz. \neg all in the case of *some*. How exactly this works is a topic of much debate, and for our purposes it doesn't really matter, as long as (32) can indeed convey (grammatically) ignorance about all. One straightforward way to do this is to assume that *at least* has *only* as an alternative (Schwarz 2016). Then (32), or more precisely S_1 in (33a), has (33b) and (33c) as alternatives. These alternatives form a symmetric pair, just like in the *some or all* case discussed above, and so we derive precisely the same meaning as there.

- (33) a. $\text{exh } [_{S_2} K [_{S_1} \text{ Bonnie stole at least some of the pears}]]$
 b. $[_{S_3} \text{ Bonnie stole only some of the pears}]$
 c. $[_{S_4} \text{ Bonnie stole (at least) all of the pears}]$
- (34) $[[\text{(33a)}]] = K(\text{some}^* \vee \text{all}) \wedge \neg K \text{some}^* \wedge \neg K \text{all}$
 $\equiv K(\text{some}^* \vee \text{all}) \wedge I \text{some}^* \wedge I \text{all}$

5.2.3 Special intonation (rise-fall-rise)

The third and final option we mention here is to use a special intonation contour for conveying ignorance, such as the rise-fall-rise (RFR) contour in English (Ward and Hirschberg 1985). We won't provide an explicit semantics for RFR (for a specific proposal, see Constant 2012); it suffices that RFR on a sentence with a scalar item like *some* conveys speaker uncertainty about the relevant alternative(s), viz. all. For simplicity, we model RFR as a kind of sentential operator, but the specifics don't concern us here.

- (35) $[[\text{RFR } [_{S_2} K [_{S_1} \text{ Bonnie stole some of the pears}]]]] = K \text{some} \wedge I \text{all}$

When we discussed the four possible parses for (26) given in (27), we therefore implicitly assumed that (26) has falling (or, at least, not RFR) intonation.

As far as we're aware, little attention has been paid to the role of specific intonation contours (like RFR) in computing scalar implicatures (but see Westera 2017, most recently). For example, textbook discussions of sentences with *some* generally don't mention the prosody with which they should be uttered, nor do experimental studies on scalar implicatures generally control for prosody (participants often just read sentences on a screen). In the discussions that follow, therefore, we assume falling (non-RFR) intonation on *some* when it comes to predictions.

	relevant(all)	¬relevant(all)
$K \text{ some} \wedge K \text{ all}$	✗	✓
$K \text{ some} \wedge K \neg\text{all}$	✓	✓
$K \text{ some} \wedge I \text{ all}$	✗	✓

Table 2: Predictions about when (26) is usable, depending on whether or not all is relevant and on the speaker’s epistemic state about all.

5.3 Predictions in more detail

We now lay out the predictions of the current system in more detail. Our discussion is divided into two parts. First, we discuss the predicted usability conditions of *some* sentences: in which contexts is (26) predicted to be usable (and under which parses)? Second, we discuss the predicted inferential patterns: which inferences does (26) license in which contexts? After laying out the predictions, we discuss in §5.4 a number of tests that help us decide whether the predictions are borne out.

5.3.1 Predicted usability conditions

If we fix the context—specifically, whether or not all is relevant, and what the epistemic state of the speaker is regarding all—then we can ask whether or not (26) is predicted to be usable (under this or that parse). We assume in each case that the speaker believes *some*, since otherwise, (26) is not usable, because using it would violate the maxim of quality.

If all is *irrelevant*, then the speaker should be able to use (26)—on any parse, in principle—regardless of her epistemic state about all.

If, however, all is *relevant*, then the speaker can only use (26) (specifically, parse (27c) or (27d)) if she believes ¬all. That is, if she believes all or is ignorant about all, then she should not be able to use (26) (with falling intonation). These predictions are summarized in Table 2.

5.3.2 Predicted inferential patterns

Rather than fixing both relevance and the speaker’s knowledge state, we can instead fix just one of them and then ask which inferences about the other one a hearer should be licensed to draw from a speaker’s use of (26). We first fix relevance about all and ask which inferences a hearer should be able to draw about the speaker’s epistemic state about all. Afterwards, we do the opposite.

If all is *irrelevant*, then a hearer can’t draw any inference at all about the speaker’s epistemic state about all, and this is because the speaker is able to use (26) regardless of which epistemic state she’s in, as Table 2 illustrates.

relevant(all)	¬relevant(all)
$\rightsquigarrow K\neg\text{all}$	$\rightsquigarrow \emptyset$

Table 3: Predictions about which inferences a hearer is licensed to draw about the speaker’s epistemic state about all when she uses (26), depending on whether or not all is relevant.

Hearer knows: $K\text{all} \vee I\text{all}$	Otherwise
$\rightsquigarrow \neg\text{relevant}(\text{all})$	$\rightsquigarrow \emptyset$

Table 4: Predictions about which inferences a hearer is licensed to draw about the relevance of all when a speaker uses (26), depending on the hearer’s knowledge of the epistemic state of the speaker about all.

If, however, all is *relevant*, then a hearer can draw the inference that the speaker believes $\neg\text{all}$, since that is the only epistemic state that she can be in to be able to use (26). These predictions are summarized in Table 3.

We now fix the speaker’s epistemic state about all and then ask which inferences a hearer should be able to draw about the relevance of all. In other words, we assume that the hearer knows the speaker’s epistemic state about all, but isn’t completely sure about whether or not all is relevant (or is ready to revise her assumptions about it).

If the hearer knows that the speaker either believes all or is ignorant about all, then from the speaker’s use of (26) she can deduce that all is *irrelevant*, for otherwise, the speaker would be conveying $K(\text{some} \wedge \neg\text{all})$, hence would be violating the maxim of quality.

In all other cases, the speaker can’t conclude one way or the other (purely from the speaker’s use of (26)) whether or not all is relevant. These predictions are summarized in Table 4.

5.4 Testing the predictions

We now turn to the question of whether these predictions are indeed borne out. We note right away that the judgments here are extremely subtle: it’s very difficult to introspect about whether (26) licenses this or that inference in this or that context, and about whether (26) is usable in this or that context. To help sharpen the judgments, we will rely on the *Are you saying that...?* test of Meyer 2013 to test inferences, and on a modified version of it, *What I’m saying is that...*, to test usability. The idea is that *Are you saying that...?* and *What I’m saying is that...* can only target meaning components that are somehow explicitly (grammatically) conveyed by a

speaker's utterance.

- (36) *Are you saying that (might) ϕ ...* can only target grammatically derived implicatures. (Meyer 2013, p. 71)

To see this, witness the following contrasts. In (37), B can target either the 'not both' implicature of A's utterance, or the ignorance implicature of A's utterance, since these are grammatically expressed (by a suitable LF structure like *exh K exh...or...*, as explained in §4.4).

- (37) A. Table 3 ordered a sandwich or a hot chocolate.
 B₁. Are you saying they didn't order both?
 B₂. Are you saying you don't know whether they ordered the sandwich or the hot chocolate?

By contrast, in (38) and (39), B attempts to target a proposition that isn't at all expressed by A's utterance (but is merely compatible with it).

- (38) A. I will buy a Gamay or a Valdiguié.
 B. #Are you saying you will buy two bottles?
 (39) A. Al went to South America for vacation.
 B. #Are you saying he went to Chile?

The question now is what this diagnostic can tell us about the interpretation of *some* sentences like (26), specifically whether the predictions laid out above are correct. We proceed in the reverse order of how we presented them, starting first with the question of which inferences the use of (26) licenses, followed by the question of which contexts support the use of (26).

5.4.1 Testing the predicted inferential patterns

If all is taken to be relevant, then (26) is usable only under a strongly exhaustive parse, hence should license the inference *K*-all, but not, for instance, *I* all. To ensure that all is relevant, we consider a dialogue in which Alice explicitly asks whether Bonnie stole all of the pears. As a consequence of Clyde's use of (26) as a response, Alice ought to be able to target the strong implicature of Clyde's response, but not any ignorance implicature.¹⁷ (As a reminder, Clyde's sentence must be read with falling (not RFR) intonation.)

¹⁷ We find that the addition of *so* before *are you saying that...* sharpens the judgments even more, perhaps because it makes it clearer that the speaker is drawing an inference from what the other person said (or at least is making reference to such a potential inference).

- (40) A. Did Bonnie steal all of the pears?
C. She stole some of them.
A₁. So, are you saying she didn't steal all of them?
A₂. ??So, are you saying you don't know whether she stole all of them?

Compare with *at least some*, where now the judgments should be reversed, since Clyde's response now (grammatically) conveys ignorance about all, as discussed in §5.2.2.

- (41) A. Did Bonnie steal all of the pears?
C. She stole at least some of them.
A₁. ??So, are you saying she didn't steal all of them?
A₂. So, are you saying you don't know whether she stole all of them?

It seems to us that there is a real contrast between the A₁ and A₂ responses in each example (as indicated by our judgment mark of "??"), as well as between the A₂ responses across both examples: A₂ is a worse response to *some* than to *at least some*. These judgments therefore support the prediction that the use of (26) in a context where all is relevant licenses the inference *K*-all, but not *I* all.¹⁸

Suppose now that Alice (the hearer) knows that Clyde (the speaker) believes that Bonnie stole all of the pears, e.g. because this is common knowledge between the two of them. Suppose further that Alice is curious about whether Clyde thinks Bonnie will go to jail and assumes that going to jail depends (at least in part) on whether Bonnie stole all of the pears, rather than just some. In this case, Clyde's use of (26) is predicted to be usable, but only if all is irrelevant. Thus, Clyde should be able to use (26) to convey to Alice that going to jail doesn't depend on Bonnie's having stole *all* of the pears; it only matters that she stole any at all. Our judgment is that this prediction is borne out.¹⁹

- (42) Context: A and C both know that Bonnie stole all of the pears.
A. Do you think Bonnie will go to jail?
C. Absolutely — she stole some of the pears.
A₁. ??So, are you saying she didn't steal all of them?
A₂. So, are you saying it's irrelevant whether (that) she stole all of them?

¹⁸ This argument is taken straight from Meyer 2013.

¹⁹ Strictly speaking, it's not clear to us that this irrelevance inference (–relevant(all)) is a grammatical inference on a par with, e.g., scalar inferences (like –all). If not, then the acceptability of the A₂ response in (42), as compared to the A₁ response, suggests that the generalization about the *Are you saying...* diagnostic in (36) is not fully general. We set these details aside for now. The important point here is that irrelevance inferences are as targetable by *Are you saying...* as scalar inferences are, given the right contextual assumptions, whereas ignorance inferences are not.

The same contrast is predicted in a context where Alice knows that Clyde is ignorant about whether or not Bonnie stole all of the pears, e.g. because this is common knowledge between the two of them. This prediction also seems correct to us.

- (43) Context: A and C are both ignorant about whether or not Bonnie stole all of the pears.
- A. Do you think Bonnie will go to jail?
 - C. Absolutely — she stole some of the pears.
 - A₁. ??So, are you saying she didn't steal all of them?
 - A₂. So, are you saying it's irrelevant whether she stole all of them?

5.4.2 Testing the predicted usability conditions

We now fix both the relevance (or irrelevance) of all and the epistemic state of the speaker regarding all, and ask whether the predictions about the usability of (26) are borne out. As before, we assume that the speaker believes some.

If all is *irrelevant*, then it appears that the speaker can use (26), regardless of his epistemic state about all: all of Clyde's responses below sound fine to us.

- (44) Context: Anyone who steals any pears goes to jail.
- A. Do you think Bonnie will go to jail?
 - C₁. Absolutely — she stole some of the pears. In fact, she stole all of them.
 - C₂. Absolutely — she stole some of the pears, though not all of them.
 - C₃. Absolutely — she stole some of the pears. She may even have stolen all of them — I don't know.

If all is *relevant*, then (26) should only be usable if $K\neg\text{all}$. In that case, the meaning conveyed is $K(\text{some} \wedge \neg\text{all})$. Consequently, the speaker should be able to follow his utterance by targeting $K\neg\text{all}$, but not, for instance, $I\text{all}$, and this seems correct.

- (45) A. Did Bonnie steal all of the pears?
- C₁. Bonnie stole some of them. So, what I'm saying is she didn't steal all of them.
 - C₂. ??Bonnie stole some of them. So, what I'm saying is I don't know whether or not she stole all of them.

Compare with *at least some*, where now the judgments are once again reversed.

- (46) A. Did Bonnie steal all of the pears?
C₁. ??Bonnie stole at least some of them. So, what I'm saying is she didn't steal all of them.
C₂. Bonnie stole at least some of them. So, what I'm saying is I don't know whether or not she stole all of them.

5.5 The status of weak implicatures

So far, the predictions of the current system seem to be on the right track. We now turn to a potential problematic prediction.

On the standard neo-Gricean view of implicatures, e.g. the Standard Recipe laid out in Geurts 2010, building on Sauerland 2004, the derivation of a scalar implicature proceeds in two steps. First, a weak (or primary) implicature of the form $\neg K\phi$ is computed. Second, if the speaker is assumed to be competent (or opinionated) about ϕ (i.e. $K\phi \vee K\neg\phi$), then the strong (or secondary) implicature $K\neg\phi$ follows. If this is correct, then a possible reading of (26), in a context where all is relevant but the Competence Assumption is neither endorsed nor denied, is $K\text{some} \wedge \neg K\text{all}$. This weak implicature reading is unexpected on the present theory, because the denotation neither settles nor entails ignorance about all.

Geurts (2010, p. 30) writes that “one of the main virtues of the Standard Recipe is that it distinguishes between weak and strong implicatures, and connects them via the Competence Assumption”. However, as far as we can tell, there has been little evidence put forward to suggest that a sentence with a scalar item like *some*, as in (26), can be interpreted as actually *conveying* a weak implicature. First, if the context is such that the Competence Assumption is explicitly denied, then speaker ignorance, and a fortiori the weak implicature, follows as a contextual entailment: the negation of $K\phi \vee K\neg\phi$ is $\neg K\phi \wedge \neg K\neg\phi (= I\phi)$. So, in such a case, nothing about the sentence meaning conveys speaker ignorance.

What if the Competence Assumption is neither endorsed nor denied, i.e. a context where the hearer doesn't know whether the speaker is competent about ϕ ? Geurts (2010, p. 30) writes about this possibility:

In a context that doesn't support the Competence Assumption, the strong implicature cannot be derived, and the Standard Recipe gives us only the weak one. To illustrate how this might happen, consider [(26)] in two slightly different scenarios. In both cases Clyde has witnessed Bonnie stealing some pears, but while in scenario S_1 he saw the platter from which she took the pears, in S_2 his view of the platter was obstructed by (say) a fat cat. Consequently, in S_1 but not in S_2 Clyde knows whether any pears are left; or in other words, while the

Competence Assumption holds in S_1 , it doesn't hold in S_2 , and the Standard Recipe forecasts a strong implicature in the former case but not in the latter. This seems to be correct: while in S_1 we would infer from Clyde's uttering [(26)] that [K -all], S_2 only allows for the weaker inference that [$\neg K$ all].

Unfortunately, in scenario S_2 it's explicitly stated that Clyde's view of the platter is obstructed, hence that Clyde is ignorant about whether Bonnie stole all of the pears. So, again, ignorance (hence also the weak implicature) is a contextual entailment.

What this all means is that it's not so clear that the present system makes any wrong predictions (yet), despite the fact that there is no parse of (26) that actually conveys $\neg K$ all (either as an entailment or as a quantity implicature): in scenario S_1 , the parses (27c) and (27d) are available and mean $K(\text{some} \wedge \neg\text{all})$, while in scenario S_2 , the parses (27a) and (27b) are available and mean just K some, with $\neg K$ all (or I all) following as a contextual entailment. (In the latter case, the obligatory irrelevance of all is compatible with the contextual entailment $\neg K$ all; indeed, the fact that the context is already transparent about the speaker's epistemic state concerning all plausibly renders all irrelevant to begin with.)

But what about a scenario S_3 where the hearer is genuinely unsure whether or not Clyde saw the platter from which Bonnie took pears (i.e. a scenario where it really is the case that the Competence Assumption is neither endorsed nor denied)? In this scenario, the current system predicts that Clyde's utterance of (26) either conveys that all is irrelevant, or conveys K -all. Thus, if all is explicitly made relevant, e.g. by the question *Did Bonnie steal all of the pears?*, then Clyde's utterance of (26) should imply K -all. What this means is that if Clyde did in fact see the platter (and saw that Bonnie stole just some of the pears), then Clyde can felicitously and truthfully use (26). However, if Clyde did not see the platter (but only saw Bonnie running away with pears presumably taken from the platter), then Clyde *cannot* use (26), because doing so would either convey a belief he doesn't actually have (K -all), or convey that all is irrelevant. Instead, Clyde would have to use a different utterance that *does* convey his ignorance about all, e.g. use *at least some* rather than plain *some*, or use RFR rather than falling intonation. We believe that this is a correct prediction: it feels inappropriate for Clyde to utter (26) (with falling intonation) when he's in fact ignorant about all. (Naturally, intuitions here are subtle, so more careful empirical testing is necessary to fully assess whether the predictions of the current system are indeed correct.)²⁰

²⁰ Westera (2017) provides a similar example below. The point here is that A neither endorses nor denies the assumption that B is competent about whether, say, Alice was at the picnic. And yet B's response can be interpreted exhaustively to convey that B believes that Alice wasn't there. Indeed,

The only piece of evidence that we're aware of for the existence of strictly weak (i.e. not ignorance) inferences associated with *some* comes from Meyer (2013), who uses the *Are you saying...?* test to distinguish weak implicatures from ignorance implicatures. First, consider (47), where the felicity of B's response question indicates that the strong implicature K -all can be targeted.

- (47) A. Al returned some of your books yesterday.
B. Are you saying he didn't return all of them?

Contrast (47) with (48), where the infelicity of B's response question indicates that the (potential) ignorance inference I all cannot be targeted.

- (48) A. Al returned some of your books yesterday.
B. #Are you saying you don't know whether he returned all of them?

As we just saw in §5.4.1, the present theory predicts this contrast between strong scalar implicatures, which can arise (grammatically) from plain *some* sentences, and ignorance inferences, which cannot.

Now consider (49). The felicity of B's response suggests that the weak implicature $\neg K$ all can be targeted. (Note that $\neg K\phi$ means that $\neg\phi$ is epistemically possible, i.e. 'it might be that $\neg\phi$ ').²¹

- (49) A. Al returned some of your books yesterday.
B. Are you saying he might not have returned all of them?

Contrast (49) with (50), where the infelicity of B's response question indicates that the (potential) inference $\neg K$ -all cannot be targeted. (Note that $\neg K\neg\phi$ means that ϕ is epistemically possible, i.e. 'it might be the case that ϕ ').²²

- (50) A. Al returned some of your books yesterday.
B. #Are you saying he might have returned all of them?

it seems to us that B's response *has* to be interpreted that way: it would be odd for A to follow up B's response with *Do you know if Alice was there?*, since B's response already settles that question. If, however, B were to use RFR intonation, or explicitly say *... but I don't know who else*, then of course B's response *Wow, only four* would be infelicitous.

- (i) A. I *may* be asking the wrong person, but who were [sic.] at the picnic?
B. John, Bob, Mary, and Sue. (with falling intonation)
A. Wow, only four huh. Must have been boring. Well, thanks!

²¹ Meyer (2013) credits Irene Heim (p.c.) with the idea of using *might not ϕ* to diagnose the presence or absence of weak implicatures.

²² The judgment in (50) is expected given the infelicity of B's response in (48) and the felicity of B's response in (49).

The contrast between (49) and (50) seems to suggest that A’s utterance could grammatically entail $\neg K$ all without entailing $\neg K$ –all, hence that A’s utterance has a reading corresponding strictly to the weak implicature: K some \wedge $\neg K$ all. If so, then the current theory is in trouble, because it predicts only the plain meaning K some or the strong meaning K (some \wedge \neg all).

We take the contrast between (49) and (50) to be real, but we wish to offer a potential alternative explanation. Consider first (51). B’s response would normally be considered quite odd, since A’s sentence is automatically taken to entail that Al went to France; thus, it’s evidently odd, at least sometimes, to target a direct entailment of a previous sentence by using *Are you saying...?*

- (51) A. Al went to Paris last week.
 B. (#)Are you saying he might have gone to France?

However, suppose that A and B both live in Texas, which has a city named Paris. In that case, B’s response all of a sudden becomes much more acceptable. Compare, for example, with #*Are you saying that someone went to Paris last week?* (which is clearly entailed by A’s utterance), or with #*Are you saying that he might have gone to Germany?* (which is clearly not entailed by A’s utterance). If so, then it would appear that *Are you saying that ϕ* can target entailments of sentences, especially if there is some kind of disambiguation (syntactic, lexical, or otherwise) taking place. In the case of (51), if A and B live in Texas, then B’s utterance is a way of asking whether A meant Paris, France, as opposed to Paris, Texas.²³

Returning now to (49), this sentence has two parses: one without *exh*, which conveys just K some, and one with *exh* below K , which conveys K (some \wedge \neg all). Importantly, the latter parse entails K –all, which in turn entails $\neg K$ all (‘might not all’), but does not entail $\neg K$ –all (‘might all’); the former parse does not entail either one. As such, it’s plausible that in (49), B’s response is a way of asking whether A meant K (some \wedge \neg all) (which entails the content of B’s response question) as opposed to just K some. By contrast, since neither parse entails $\neg K$ –all, B’s response in (50) is odd.²⁴

23 B could also reply with *Are you saying he went to France?*, without the modal *might*. Our impression is that the use of *might* is a kind of hedge: it feels most natural when B is mostly sure that A meant Paris, Texas, but has some reason to believe that A might have meant Paris, France.

24 A reviewer points out that this explanation seems to only work for contexts in which it’s not known whether all is relevant: if all is (known to be) irrelevant, then the only possible interpretation of A’s utterance is K some, and if all is (known to be) relevant, then the only possible interpretation is K (some \wedge \neg all); in neither case is there any ambiguity about the parse of the sentence. However, as the reviewer observes, if we ensure that all is relevant, e.g. if B explicitly asks A whether Al returned all of the books, the contrast persists. We agree with this intuition; however, we don’t need to argue that disambiguation is a necessary condition on the use of *Are you saying that...?* (even if that may be its most common function). In particular, we’re arguing that the contrast boils down to the difference between saying “Are you saying that ϕ ?” when ϕ is entailed vs. not entailed by the

As in the case of the intonation data, the intuitions here are quite delicate (though they seem to point in the right direction, in our opinion). Needless to say, therefore, these preliminary remarks deserve further scrutiny.

6 Discussion

6.1 Comparison with Meyer 2013

The theory we've developed here has a lot of similarity, but also several crucial points of difference, with the Matrix K Theory of Meyer 2013. One major difference (noted already in fn. 9) is that Meyer's theory rests on the *hypothesis* that K is attached to every assertively used sentence (Meyer 2013, p. 42), whereas for the present theory, matrix K is effectively a *consequence* of the hypothesis that relevance is closed under belief. To appreciate the difference between the two systems, take the case of disjunction, (8). It's uncontroversial that the use of (8) can convey speaker ignorance about the individual disjuncts, inferences traditionally analyzed in pragmatic terms. Meyer argues that these inferences are best captured in grammar, and to do so hypothesizes that (8) is parsed with *exh* K . By contrast, on the present theory, the *only* way for ignorance to arise with a disjunction like (8) is for (8) to be parsed with *exh* K ; otherwise, the individual disjunct alternatives are obligatorily irrelevant.

Another major point of difference has to do with which parses, and which readings, each theory predicts to be available. As already mentioned, on the basis of the *Are you saying that...?* test, Meyer concludes that when it comes to *some* sentences and their *all* alternatives, strong implicatures (K -all) and weak implicatures ($\neg K$ all) should be derivable in grammar, but not ignorance implicatures (I all). To ensure this result, Meyer needs to restrict which parses/readings her theory generates, and to do this she proposes a regulatory principle she calls *epistemic transparency*.

(52) **Epistemic transparency** (Meyer 2013, p. 58)

An LF of the form [$\dots K S$] is licensed iff it entails the speaker's state of mind about every $\phi \in \text{alt}(S)$. States of mind are:

previous utterance. Indeed, one can construct an example where the same sort of contrast appears, even though A 's response is totally unambiguous and completely answers A 's question:

- (i) B. Was Al born in France?
 A. He was born in Canada.
 B₁. Are you saying he might not have been born in France??
 B₂. #Are you saying he might have been born in France??

We thank the reviewer for this important comment, and we hope to study the felicity conditions of *Are you saying that...?* in more detail in future work.

- a. $K\psi$
- b. $\neg K\psi$

So, for Meyer, the parse in (27b) for (26), repeated in (53), is licensed even in a context in which all is relevant, and hence denotes K some $\wedge \neg K$ all (the weak implicature reading), because this parse is epistemically transparent about all, according to the definition of epistemic transparency given in (52).

(53) $\text{exh}_{[s_2]} K [s_1 \text{ Bonnie stole some of the pears}]$

Interestingly, closing relevance under belief automatically induces its own notion of epistemic transparency, but one which differs in a small but crucial respect from that in (52). Specifically, a consequence of closing relevance under belief is that a speaker’s utterance S must settle or entail ignorance about every relevant proposition ϕ (cf. the generalization in (19)). An equivalent, and perhaps more intuitive, way of understanding this consequence is to say that a speaker’s utterance S must convey her epistemic state about every relevant proposition ϕ , where epistemic states come in *three*, not two, logically exhaustive and mutually exclusive varieties: certainty that ϕ , certainty that $\neg\phi$, or uncertainty about ϕ .²⁵

(54) **Epistemic transparency 2** (equivalent to (19))

An utterance S must entail the speaker’s epistemic state about every relevant proposition ϕ . Epistemic states about ϕ are:

- a. $K\phi$
- b. $K\neg\phi$
- c. $\neg K\phi \wedge \neg K\neg\phi$ (= $I\phi$)

Given this equivalence, we stress once more that the requirement to be epistemically transparent arises automatically in our system as a consequence of closure under belief, together with the maxim of quantity.

The difference between Meyer’s stipulated principle and the one that falls out of our system boils down to the notion of epistemic state. For Meyer, there are two, $K\phi$ and $\neg K\phi$, while our system breaks the latter down into $K\neg\phi$ and $I\phi$. That is, Meyer’s principle of epistemic transparency doesn’t require speakers to be transparent about whether they are in a $K\neg\phi$ vs. an $I\phi$ state, whereas ours does. Intuitively, it seems to us that hearers would care about this difference; that is, we find our more restrictive notion of epistemic state more intuitively satisfying.

²⁵ The equivalence between (19) and (54) holds for the following reasons. (19) states that a speaker’s utterance S must entail ϕ , $\neg\phi$, or $I\phi$, for every relevant ϕ . Since every sentence must be c-commanded by K , (20), the only way that S could entail ϕ or $\neg\phi$ is by entailing $K\phi$ or $K\neg\phi$, hence entailing ϕ or $\neg\phi$ by the T axiom. In other words, obeying the generalization in (19) entails being epistemically transparent. Conversely, being epistemically transparent entails obeying (19), as well. Hence, the two notions are equivalent.

However, whether or not language really does care about this distinction is an empirical question.

6.2 Cases of oddness (Magri 2009)

In §4.1, we laid out the far-reaching consequences of adding the condition of closure of relevance under belief to the standard grammatical theory. Therefore, we may expect that this condition will interact in nontrivial ways with yet other closure conditions on relevance that we might have reason to adopt. In this section, we consider one such condition, viz. closure of relevance under contextual equivalence, proposed in Magri 2009.

Magri wishes to explain the oddness of sentences like (55).

(55) #Some Italians come from a warm country.

To do so, he makes three assumptions: (i) that relevance is closed under contextual equivalence (the aforementioned closure condition), (ii) that *exh* is mandatory in matrix clauses, and (iii) that a sentence sounds odd if it expresses a proposition that contradicts common knowledge. By (i), the *all* alternative of (55) is relevant in all contexts in which all Italians come from the same country (i.e. in all natural contexts, since this is common knowledge).²⁶ By (ii), (55) is parsed with *exh* and hence denotes the proposition that some but not all Italians come from a warm country.²⁷ Therefore, by (iii) and the fact that this proposition contradicts common knowledge, (55) sounds odd.

Within the current system, we can explain this judgment by adopting (i) and (iii), while at the same time *not* stipulating (ii). More precisely, in the current system we automatically derive that (55) must be parsed with *exh* (below *K*) since otherwise the maxim of quantity would yield a contradiction. Here's why. Consider the four parses in (56), which exhaust what can possibly be expressed by (56).²⁸ As before, we write "some" to refer to the proposition denoted by S_1 and "all" to refer to its universal counterpart.

- (56) a. $K [S_1 \text{ some Italians come from a warm country}]$
 b. $exh [S_2 K [S_1 \text{ some Italians come from a warm country}]]$
 c. $K exh [S_1 \text{ some Italians come from a warm country}]$
 d. $exh [S_3 K exh [S_1 \text{ some Italians come from a warm country}]]$

²⁶ More precisely, the relevance of the *all* alternative follows given that, furthermore, the proposition denoted by (55) is itself relevant. See Magri 2009 for a stipulation to this effect.

²⁷ As discussed in Magri 2009, exhaustification must be assumed to be contextually blind for the strengthened meaning to arise. That is, innocent excludability must be defined by means of the relation of logical entailment and not contextual entailment.

²⁸ That is, additional occurrences of *K* and/or *exh* would not lead to different meanings (see fn. 14).

Note first that regardless of which of the structures in (56) is uttered, some is relevant in the utterance context c (i.e. some \in relevant $_c$) because the utterance contains S_1 . By the common knowledge that all Italians come from the same country, some and all are contextually equivalent. Therefore, all \in relevant $_c$ because some \in relevant $_c$.

Now consider the parse in (56a). As discussed in §5.1, this parse doesn't settle or entail ignorance about all; thus, all is obligatorily irrelevant to (56a).²⁹ However, we also have all \in relevant $_c$, given that some \in relevant $_c$. As such, there is a conflict between the obligatory irrelevance of all, in the sense of (22), on the one hand, and the (obligatory) relevance of all due to contextual equivalence with some, on the other hand. The end result is that the maxim of quantity yields a contradiction, and so we take it that the parse in (56a) is therefore unavailable. Likewise, the parse in (56b) neither settles nor entails ignorance about all, so the maxim of quantity again leads to a contradiction since all \in relevant $_c$.

This leaves us with the parses in (56c) and (56d). Again as discussed in §5.1, the *exh* operator below K yields the inference \neg all so that both (56c) and (56d) entail the contextual contradiction that some but not all Italians come from a warm country.

Thus, overall we derive that (55) can only convey a contextual contradiction or bring about a contradiction by the maxim of quantity. This explains why (55) sounds odd. Importantly, unlike in Magri 2009, we derive this result without having to stipulate the syntactic distribution of *exh*. This means that in certain circumstances (here provided by common knowledge that leads to contextual equivalence of some and all), obligatory irrelevance can explain the distribution of *exh*.

6.3 Contextual expansion of formal alternatives

In §1, we adopted the view that the domain of the exhaustivity operator *exh*, applied to a sentence S in an utterance context c , is the intersection of two sets, viz. the set of propositions that are relevant in c and the set of all propositions that are denoted by formally derivable alternatives of S in c (see the definition in (3)). As before, we refer to the set of formally derivable alternatives of S in c with $\text{alt}_c^{FA}(S)$, and we cited Katzir 2007 as an example of how this set can be defined. Importantly, Katzir assumes that the extension of $\text{alt}_c^{FA}(S)$ depends in part on the utterance context (as indicated by the subscript c). This assumption is motivated by the observation that the context can extend what implicatures are available. For an illustration, consider the sequence of sentences in (57), which we borrow from

²⁹ This claim relies crucially on the assumption that the underlying notion of entailment (in terms of which settlement is defined) is logical, not contextual, entailment (see fn. 15). Thus, if our explanation of odd sentences like (55) is on the right track, then it provides an argument for making this assumption.

Trinh and Haida 2015 (modeled after an example presented in Matsumoto 1995).

(57) Yesterday, it was warm. Today, it is warm and sunny with gusts of wind.

The crucial observation with respect to (57) is that the two-sentence sequence licenses the implicature that it wasn't sunny with gusts of wind yesterday. Importantly, this implicature cannot be drawn from the first sentence outside of the particular context provided by the second sentence (or a similar context). Such implicatures hence are called *particularized implicatures* as opposed to *generalized implicatures*, which arise without contextual support (Grice 1975). Katzir (2007) accounts for particularized implicatures by assuming that phrases that occur in the utterance context of a sentence *S* can be used to derive formal alternatives of *S*. In the case of (57), the second sentence provides the context for the first sentence, which we assume to be parsed as given in (58a). Hence, the phrase *warm and sunny with gusts of wind* (= α), contained in the second sentence, (58b), can be used to derive a formal alternative of *yesterday it was warm*, namely by substituting α for the adjective *warm*, (58c).

- (58) a. K [exh [_S yesterday it was warm]]
 b. today it is [_{α} warm and sunny with gusts of wind]
 c. yesterday it was warm $\xrightarrow{[\alpha/\text{warm}]}$ yesterday it was [_{α} warm and sunny with gusts of wind]

Consequently, $\text{alt}_c(S)$ contains the proposition denoted by *yesterday it was warm and sunny with gusts of wind*,³⁰ and since this proposition is IE given $\llbracket S \rrbracket$,³¹ the denotation of *exh S* (and hence the denotation of *K exh S*) entails that it wasn't sunny with gusts of wind yesterday. Thus, Katzir's analysis accounts for the particularized implicature observed for (57).³²

Now, the existence of particularized implicatures and Katzir's analysis of this phenomenon raise the question for what structures we predict that obligatory irrelevance can even arise. That is, if context can extend the set of formally derivable alternatives of a sentence *S* and hence extend the domain of *exh*, then it may turn out that *any* proposition can, in a particular context, be settled or entailed ignorance about by, e.g., the parse *exh K exh S*. That is, it may turn out

³⁰ That is, it is plausible to assume that, in addition to being the denotation of a formal alternative, this proposition is also relevant in *c*, given that the speaker adds to *c* that it is warm and sunny with gusts of wind on the day of utterance.

³¹ The innocent excludability of this proposition rests on the assumption that the lexicon does not contain a lexical item with the meaning of 'not warm and not sunny with gusts of wind'. See Trinh 2018 for discussion.

³² This means that there may even be a context in which an utterance of the sentence *Ann is Canadian* (parsed as K [exh [Ann is Canadian]]) can implicate that it is(n't) raining in Paris and hence settle this proposition, counter to what we claimed at the end of section §4.3.

that obligatory irrelevance is only predicted for parses without *exh*.³³ The example in (57) illustrates that there are particularized scalar implicatures and thus also contexts that extend the set of formal alternatives to the effect that sentences of the form *K exh S* settle propositions that wouldn't be settled outside of these contexts. What we will discuss next is that there seem to be no particularized ignorance inferences. If indeed such inferences cannot be contextually triggered, this means that there are no contexts that extend the set of formal alternatives to the effect that sentences of the form *exh K (exh) S* can entail ignorance about a proposition that wouldn't be entailed ignorance about outside of these contexts. To see that there are no particularized ignorance inferences, consider the sentence in (59), taken from Katzir 2007.

- (59) John talked to some of the girls yesterday, and he talked to some but not all of the girls today.

Katzir notes that according to his theory (59) has the two formal alternatives in (60a) and (60b), where the former is derived from (59) by replacing *some* in the first conjunct with *all*, which is taken from the lexicon, and the latter by replacing *some of the girls* in the first conjunct with *some but not all of the girls*, which is taken from the second conjunct (i.e. from the context of the first conjunct).

- (60) a. John talked to all of the girls yesterday, and he talked to some but not all of the girls today.
 b. John talked to some but not all of the girls yesterday, and he talked to some but not all of the girls today.

Moreover, Katzir observes that (59) is “distinctly odd”, a judgment that we very much agree with. We take the oddness of (59) to show that it cannot entail ignorance about the propositions denoted by the alternatives in (60).^{34, 35} Importantly, the structure *exh K (exh) S*, where *S* is the sentence in (59), would entail ignorance about these propositions if the sentences in (60) were formally derivable alternatives of

³³ Note that this would be an interesting result since in this case obligatory irrelevance would regulate the occurrence of *exh* in particular contexts. See the concluding remarks in §7.

³⁴ To see that ignorance inferences about the alternatives in (60) can salvage a conjunctive assertion like that in (59), note that the conjunction in (i) below is perfectly acceptable.

- (i) John talked to at least some of the girls yesterday, and he talked to some but not all of the girls today.

We can explain the contrast between (59) and (i) by the fact that the occurrence of *at least* in (i) conveys ignorance about whether John talked to just some of the girls today or to all of them.

³⁵ Katzir states that given the alternatives in (60) “all we get for [(59)] is an ignorance inference”, but concedes in a footnote, “[t]hough I do not know why the result is odd”.

S.^{36,37} Thus, we are led to conclude that Katzir’s theory overgenerates alternatives (see also Trinh and Haida 2015; Breheny et al. 2017; Trinh 2018 for observations that justify the same conclusion). Specifically, it falsely predicts that there are contexts in which (60b) is a formal alternative of (59).

It is not our goal to criticize the analysis in Katzir 2007. Rather, our goal is to find out what exact predictions regarding obligatory irrelevance of alternatives arise from our theory on the background of the correct theory of formal alternatives. In particular, we concluded in §5.1 that the proposition that Bonnie stole all of the pears is obligatorily irrelevant to the structure (*exh*) *K Bonnie stole some of the pears*. What the foregoing discussion illustrates is that, despite the ability of the set of formal alternatives to be expanded in cases like (57), our claim about structures like *exh K...some...* is still correct, given the apparent inability of the set of formal alternatives to be expanded (in the relevant way) in cases like (59).

Before closing this section, we want to reinforce the empirical claim that there are no particularized ignorance inferences. Support for this claim comes from the observation that there are no particularized free-choice inferences either. Consider, for instance, the sequence of sentences in (61). The second sentence of this sequence cannot be understood to convey that today John is allowed to talk to just some of the girls and is also allowed to talk to all of them.

(61) Yesterday, John was allowed to talk to some but not all of the girls. Today, he is allowed to talk to some of the girls.

Assuming the theory of free-choice readings proposed in Fox 2007, the second sentence of (61) (= *S*), parsed as *exh exh S*, entails that both (62a) and (62b) are true if these two sentences are formally derivable alternatives of *S*.

- (62) a. Today, John is allowed to talk to all of the girls.
 b. Today, John is allowed to talk to some but not all of the girls.

Since the sequence in (61) doesn’t license this inference, we can conclude that

³⁶ Neither alternative is IE given *S* since the conjunction of (59) (= *S*₁) with the negation of (60a) (= *S*₂) entails (60b) and vice versa. Both *K* [*S*₁] and *K* [*S*₂] are IE given *K* [*S*], and *K* [*S*] ∧ ¬*K* [*S*₁] ∧ ¬*K* [*S*₂] entails *I* [*S*₁] and *I* [*S*₂], which in turn entail ignorance about whether John talked to just some of the girls today or to all of them.

³⁷ In a recent paper, Trinh (2018) considers the discourse in (i) and observes that “disjunctions give rise to ignorance inferences much more naturally than such cases as [(i)].”

(i) A. Bill ran. He did not smoke.

B. What about John?

A. John ran.

Like in the case of (59), (i) can reasonably (but incorrectly) be expected to give rise to an ignorance inference since by Katzir’s theory the last sentence of (i) has both *John smoked* and *John didn’t smoke* as formal alternatives.

the sentences in (62) are not both formally derivable alternatives of the second sentence of this sequence. That is, we can conclude that (62b) is not a formally derivable alternative of the second sentence of (61). For reasons of space, we cannot discuss why this restriction holds, i.e. why there are no particularized ignorance or free-choice inferences, but see Trinh 2018 for an amendment of Katzir’s theory, which explains this fact.

6.4 Logical vs. contextual entailment

Throughout this article, we’ve assumed that the underlying notion of entailment in terms of which settlement is defined is logical entailment, i.e. that the maxim of quantity is articulated in terms of logical entailment. However, it’s often assumed that the maxim of quantity compels a speaker to make an utterance that merely contextually settles every relevant proposition. We now briefly discuss the consequences for the current theory of adopting this view.

First, by switching to contextual entailment, we lose the strong notion of obligatory irrelevance as defined in (22), in the sense that no proposition could ever count as obligatorily irrelevant. The reason for this is that for any sentence S and any proposition ϕ , there is, presumably, always some context c in which S contextually settles or contextually entails ignorance about ϕ . However, even still, all the predictions we discussed concerning *or* (§4.4) and *some* (§5) would still hold, because there we only considered contexts where there was no contextual entailment relation between a disjunction and its disjuncts, or between some and all. Our adoption of logical entailment could thus be seen as a methodological decision: it allowed us to easily define obligatory irrelevance, which in turn allowed us to easily explain the textbook predictions. Using contextual entailment would’ve required a more complicated definition of obligatory irrelevance (or an abandonment of the notion, which we find useful).

Second, we lose the advantage, discussed in §6.2, that the present system has over Magri 2009. In contexts where some and all are contextually equivalent, some contextually settles all. Thus, in such contexts, all can be relevant to a sentence like (55) (#*Some Italians come from a warm country*), even when it isn’t parsed with any *exh*. But in that case, no contradiction arises, and so the oddity of the sentence remains unexplained. In other words, by moving from logical to contextual entailment, we lose the strong notion of obligatory irrelevance (see above), which we had previously shown forced the application of *exh* in sentences like (55). Of course, we don’t fare any worse than Magri 2009 (or any other theory): we could, like Magri, stipulate that *exh* is mandatory, which would re-deliver the contextual contradiction, hence explain the oddity.

Finally, by adopting contextual entailment, we run the risk of undermining our conclusion that ignorance inferences are derived in grammar, and that a syntactic

item like K is available in language. However, showing that there is no K (with relevance closed under belief) would involve showing that for every utterance S , every context c , and every relevant alternative ϕ not (contextually) settled by S , S contextually (in c) entails ignorance about ϕ . This seems like a tall order: it's hard to see how a disjunction like (8), uttered out of the blue, could contextually entail ignorance about its individual disjuncts. In the absence of any proposal about how this might work, we think that we can maintain our conclusion that matrix K follows from closure of relevance under belief.

Before concluding, we would like to offer one small piece of evidence for the idea that settlement is sensitive to logical entailment. Consider the scenario below and the sentences that follow it.

- (63) Context: Ann played a card game in which, given the rules, the final score is always an even number of points. Bob knows this, and reports to Carl:
- a. #Ann scored at least 3 points.
 - b. Ann scored at least 4 points.

According to our judgments, the sentence with *at least 3* is decidedly odd, while the sentence with *at least 4* is not. A straightforward way to explain the contrast is to observe, first, that an utterance of the form *Ann scored at least n points* implies speaker ignorance about whether Ann scored exactly n points or more than n points (Geurts and Nouwen 2007; Büring 2008; Mayr 2013; Kennedy 2015; Schwarz 2016). As such, (63a) implies that the speaker is uncertain whether Ann scored exactly 3 points, (64a), which contradicts the common knowledge that final scores in this game always consist of an even number of points, and that the speaker is uncertain whether Ann scored more than 3 points, (64b). By contrast, the ignorance inferences associated with (63b) don't contradict contextual knowledge, and so the sentence is fine.

- (64) a. Ann scored exactly 3 points.
b. Ann scored more than 3 points.

The ignorance inferences associated with *at least* have been argued to arise as quantity implicatures (Mayr 2013; Kennedy 2015; Schwarz 2016). If so, then the explanation of the contrast between (63a) and (63b) sketched above provides evidence that the maxim of quantity cares specifically about logical, not contextual, entailment. Here's why: (63a) contextually settles both (64a) and (64b)—specifically, (64a) is a contextual contradiction, and (63a) contextually entails (64b)—and so no ignorance about either one is predicted to arise, which means that the oddity of (63a) remains a mystery.

Needless to say, these remarks on the nature of the underlying notion of entailment remain preliminary. We hope to examine this issue in more detail in

future research.

7 Conclusion

Fox (2007) argues that exhaustification, i.e. the use of the operator *exh*, serves to reduce ignorance inferences: if S is a parse of a sentence and the maxim of quantity produces a speaker ignorance inference about an alternative ϕ because S doesn't settle ϕ , then the sentence can be parsed as *exh* S to potentially assign it a stronger meaning that settles ϕ and hence eliminates the speaker ignorance inference about it. In this article, we have shown that, surprisingly, the hypothesis that relevance is closed under speaker belief leads to the opposite conclusion: *exh* serves to *produce* speaker ignorance inferences, when it occurs above the epistemic operator K , and to prevent obligatory irrelevance, when it occurs below K , so that the maxim of quantity can remain active and the utterance can conform to what is contextually relevant. We have discussed empirical data, mostly centered around textbook examples, that suggest that the consequences of closing relevance under belief are indeed empirically attested: (i) a sentence gives rise to ignorance inferences if and only if it has symmetric *formal* alternatives (§4.4 and §5.4.1), (ii) a sentence can give rise to inferences about the relevance of certain alternatives (again §5.4.1), and (iii) the (obligatory) relevance of certain alternatives can force *exh* to be present in a sentence (§6.2). We also discussed the potentially problematic prediction regarding the absence of weak implicatures and found that the available empirical facts don't adjudicate between the theory at hand and theories that can derive weak implicatures (§5.5). Overall, we therefore conclude that the theory discussed in this article deserves further attention, not least because it puts the spotlight on new empirical questions and theoretical challenges that need yet to be fully addressed (§5.5, §6.3, and §6.4).

A 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 closure of A under conjunction, negation, and belief, as described by the closure conditions on relevance in (16). 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

Breheny, Richard, Nathan Klinedinst, Jacopo Romoli, and Yasutada Sudo (2017). The Symmetry Problem: Current Theories and Prospects. In: *Natural Language Semantics*, pp. 1–26. DOI: [10.1007/s11050-017-9141-z](https://doi.org/10.1007/s11050-017-9141-z).

- 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, England: Oxford University Press, pp. 39–103.
- Chierchia, Gennaro (2006). Broaden Your Views: Implicatures of Domain Widening and the “Logicity” of Language. In: *Linguistic Inquiry* 37.4, pp. 535–590. DOI: [10.1162/ling.2006.37.4.535](https://doi.org/10.1162/ling.2006.37.4.535).
- Chierchia, Gennaro (2013). *Logic in Grammar: Polarity, Free Choice, and Intervention*. Oxford, England: Oxford University Press.
- 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 Stechow, and Paul Portner. Vol. 3. Berlin, Germany: Mouton de Gruyter, pp. 2297–2331.
- Constant, Noah (2012). English rise-fall-rise: A Study in the Semantics and Pragmatics of Intonation. In: *Linguistics and Philosophy* 35.5, pp. 407–442. DOI: [10.1007/s10988-012-9121-1](https://doi.org/10.1007/s10988-012-9121-1).
- Evans, Jonathan St. B. T., Stephen E. Newstead, and Ruth M. J. Byrne (1993). *Human Reasoning: The Psychology of Deduction*. Hove, England: Lawrence Erlbaum Associates Ltd.
- Fintel, Kai von (1993). Exeptive Constructions. In: *Natural Language Semantics* 1.2, pp. 123–148. DOI: [10.1007/BF00372560](https://doi.org/10.1007/BF00372560).
- Fox, Danny (2007). Free Choice and the Theory of Scalar Implicatures. In: *Pre-supposition 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](https://doi.org/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](https://doi.org/10.3765/sp.7.5).
- Fox, Danny (Sept. 10, 2016). On Why Ignorance Might Not Be Part of Literal Meaning. Commentary on Marie-Christine Meyer, MIT Workshop on Exhaustivity. Cambridge, MA. URL: <https://drive.google.com/file/d/0B0oeGQ78K8a1Q1BncUVpa1A0eW8/view>.
- Geurts, Bart (2010). *Quantity Implicatures*. Cambridge, England: Cambridge University Press.

- 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>.
- Grice, H. P. (1975). Logic and Conversation. In: *Syntax and Semantics, Volume 3: Speech Acts*. Ed. by Peter Cole and Jerry L. Morgan. New York, NY: Academic Press, pp. 41–58.
- Hintikka, Jaakko (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ithaca, NY: Cornell University Press.
- Horn, Laurence R. (1972). On the Semantic Properties of Logical Operators in English. PhD thesis. Los Angeles, CA: University of California, Los Angeles.
- Katzir, Roni (2007). Structurally-Defined Alternatives. In: *Linguistics and Philosophy* 30.6, pp. 669–690. DOI: [10.1007/s10988-008-9029-y](https://doi.org/10.1007/s10988-008-9029-y).
- 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](https://doi.org/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.
- Magri, Giorgio (2009). A Theory of Individual-Level Predicates Based on Blind Mandatory Scalar Implicatures. In: *Natural Language Semantics* 17.3, pp. 245–297. DOI: [10.1007/s11050-009-9042-x](https://doi.org/10.1007/s11050-009-9042-x).
- Matsumoto, Yo (1995). The Conversational Condition on Horn Scales. In: *Linguistics and Philosophy* 18.1, pp. 21–60. DOI: [10.1007/BF00984960](https://doi.org/10.1007/BF00984960).
- 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>.
- 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](https://doi.org/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](https://doi.org/10.3765/sp.9.1).
- Trinh, Tue (2018). Keeping It Simple. In: *Natural Language Semantics*, pp. 1–14. DOI: [10.1007/s11050-018-9143-5](https://doi.org/10.1007/s11050-018-9143-5).
- 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](https://doi.org/10.1007/s11050-015-9115-y).

- Ward, Gregory and Julia Hirschberg (1985). Implicating Uncertainty: The Pragmatics of Fall-Rise Intonation. In: *Language* 61.4, pp. 747–776. DOI: [10.2307/414489](https://doi.org/10.2307/414489). URL: <http://www.jstor.org/stable/414489>.
- Westera, Matthijs (2017). Exhaustivity and Intonation: A Unified Theory. University of Amsterdam.