

# WHEREOF ONE CANNOT SPEAK\*

AMIR ANVARI

*Massachusetts Institute of Technology*

## 1 Introduction

Normally one thinks of relevance as being closed under negation and conjunction (Lewis, 1988). Fox (2016) suggests, however, that relevance is also closed under speakers' beliefs.<sup>1</sup>

- (1) If a question  $Q$  is directed at a participant  $S$  and the propositions  $\mathbf{p}$  and  $\mathbf{q}$  are among those made relevant by  $Q$  then the propositions  $\neg\mathbf{p}$ ,  $\mathbf{p} \wedge \mathbf{q}$  and  $\mathbf{K}_s(\mathbf{p})$  are made relevant by  $Q$  as well, where  $\mathbf{K}_s(\mathbf{p})$  means  $S$  believes that  $\mathbf{p}$ .

One might ask whether (1) requires argument. After all, if a certain question  $Q$  is asked of a participant  $S$  then  $S$ 's beliefs with respect to  $Q$  are surely pertinent. However, relevance plays a particular role in cooperative information-exchange and the idea that relevance is closed under speakers' beliefs is in fact far from innocuous. Specifically, in conjunction with a standard formulation of norms of cooperative information-exchange, (1) has the drastic consequence that the *semantic* content of one's utterance is required to settle one's *epistemic attitude* towards all relevant propositions. A bit more specifically, from (1) together with the principle of Quantity, i.e. that one ought to say everything that one believes to be true so long as it is relevant, the statement in (2) follows (see Fox 2016, and especially Buccola and Haida 2019 for detailed discussion).<sup>2</sup>

---

**\*Thanks to Gabor Brody for jump starting this project by raising a problem for Fox's treatment of silence.**

<sup>1</sup>Fox's claim is in fact stronger. Specifically, he proposes that if  $\mathbf{p}$  is relevant then so is  $\mathbf{K}_x(\mathbf{p})$ , for each discourse participant  $x$ . As far as I can see, it is the weaker statement in the text that's justified by the considerations that Fox brings to the table and I will elaborate on immediately below.

<sup>2</sup>*Proof Sketch:* Assumption #1: If  $\mathbf{p}$  is relevant and  $S$ 's utterance  $\phi$  does not settle  $\mathbf{p}$  then it follows from Quantity and the assumption that relevance is closed under negation that  $S$  is ignorant about  $\mathbf{p}$ . Assumption #2: if  $S$  believes  $\mathbf{p}$  then  $S$  believes that he believes  $\mathbf{p}$  ("positive introspection") and if  $S$  does not believe  $\mathbf{p}$  then  $S$  believes that he does not believe  $\mathbf{p}$  ("negative introspection"). Now suppose  $\mathbf{p}$  is relevant. From (1), it follows that  $\mathbf{K}_s(\mathbf{p})$  is relevant. If  $\phi$  does not settle  $\mathbf{K}_s(\mathbf{p})$  then it follows from Assumption #1 that  $S$  is ignorant about  $\mathbf{K}_s(\mathbf{p})$ . But from Assumption #2 it follows that  $S$  cannot possibly be ignorant about whether he believes  $\mathbf{p}$  or not. Therefore it must be the case that  $\phi$  settles  $\mathbf{K}_s(\mathbf{p})$ . Now if  $\mathbf{p}$  is relevant then  $\neg\mathbf{p}$  is relevant too which means that  $\mathbf{K}_s(\neg\mathbf{p})$  is relevant as well. Same reasoning as above leads to the conclusion that  $\phi$  must settle  $\mathbf{K}_s(\neg\mathbf{p})$ . Since  $\mathbf{K}_s(\mathbf{p})$  and  $\mathbf{K}_s(\neg\mathbf{p})$  are mutually incompatible, for  $\phi$  to settle both  $\mathbf{K}_s(\mathbf{p})$  and  $\mathbf{K}_s(\neg\mathbf{p})$  consistently it must be the case that either  $\phi$  entails  $\mathbf{K}_s(\mathbf{p})$ , or it entails  $\mathbf{K}_s(\neg\mathbf{p})$ , or it entails  $\neg\mathbf{K}_s(\mathbf{p}) \wedge \neg\mathbf{K}_s(\neg\mathbf{p})$ .

- (2) If a question  $Q$  is directed at a participant  $S$  then, for each proposition  $\mathbf{p}$  made relevant by  $Q$ ,  $S$ 's utterance ought to entail either that  $S$  believes  $\mathbf{p}$  to be true, or that  $S$  believes  $\mathbf{p}$  to be false, or that  $S$  is ignorant as to whether  $\mathbf{p}$  to true or false.

Indeed (2) is the consequence of (1) that Fox is interested in. Fox's primary concern is to understand the "conceptual underpinnings" of a specific proposal made by Meyer (2013), namely that ignorance inferences are computed within the grammar, and his proposal is made as a contribution to that project in light of (2).<sup>3</sup> In this note my primary concern is not Meyer's proposal or whether (1) actually motivates it. Instead I'd like to focus on a particular motivation for (1) that Fox provides based on an independent observation, namely that "silence is uncooperative".

Suppose the Lawyer asks the Witness "*Where was John at the time of the murder?*". Fox's observation is that "if [the Witness] believes something that bears on John's whereabouts at the time of the murder, [he] is required to say so. If not, he is required reveal this lack of opinion". The Witness cannot "look the Lawyer in the eye but remain silent". More generally, the point is that silence is intuitively not a legitimate way of conveying ignorance. If relevance is not closed under speakers' beliefs, this observation is mysterious. After all, if the speaker is unopinionated about all relevant propositions in the context then there is no relevant proposition that he can truthfully assert and silence appears as the advisable option. But if relevance is closed under the speakers' beliefs, the situation changes. In that case the speaker is never truly unopinionated; he always has *some* relevant information to truthfully contribute, and therefore there is always *something* the speaker can (and by Quantity should) say. For an unopinionated speaker, this relevant information is simply that he is ignorant.

Now, it is debatable whether this is an adequate way of thinking about silence. After all, it appears that turn-taking interaction is quite generally premised on the expectation that the participants ought to do *something* when their turn comes. Silence is reasonably viewed as a kind of doing nothing and as such its unacceptability may be viewed as rooted in the very conceptualization of what it means to engage in turn-taking interaction as such. If so then an explanation based on the principles of cooperative information-exchange, a very specific kind of turn-taking interaction, seems too limited to do justice to the generality of phenomenon.<sup>4</sup>

I therefore believe that it is reasonable to ask whether a motivation for (1) can be provided that does not rely on the reasons behind the uncooperativeness, or more generally, the unacceptability of silence. I'd like to argue that the answer is positive, that even if silence is presumed to be unacceptable for independent reasons Fox's argument still goes through. Specifically, I'd

<sup>3</sup>The statement in (2) is nearly identical with what Meyer dubs the Epistemic Transparency principle. The only difference between the two is that ET is formulated with respect to relevant propositions that are expressed by the utterance's formal alternatives while (2) is formulated with respect to relevant propositions as such. In other words, (2) is stronger than ET. In Meyer's system, ET acts as a constraint that filters out certain grammatically available readings. Fox's idea is that if (2) can be motivated via (1) then the core grammatical commitment of Meyer's proposal, namely the Matrix K hypothesis, which together with covert exhaustification allows the grammar to derive not only scalar implicatures but also ignorance inferences, can be motivated as a grammatical response to a pragmatic pressure, namely to somehow meet the demand in (2).

<sup>4</sup>As Gabor Brody points out (p.c.), the acute stress that infants experience in Tronick's "still face" experiment (see Adamson and Frick 2003) is perhaps an indication of how deeply rooted the aversion to silence is. The urge to avoid silence is so strong that it is difficult to repress even at the risk of self-incrimination. Given freedom to do so, a police officer can extract a lot of information by leveraging this observation without engaging in overt interrogation. It is telling that the Miranda warning serves not only as a declaration that one is no longer in a cooperative context ("everything you say will be used against you") but also as a reminder of the Fifth Amendment Right to Silence.

like to show that Fox's argument for (1) can be reconstructed using the uncooperativeness of *uninformative utterances* as the starting point.

Above I said "if the speaker is unopinionated about all relevant propositions in the context then there is no relevant proposition that he can truthfully assert". This statement is misleading, however, because there are *always* relevant propositions that any speaker is opinionated about, namely those that amount to contextual trivialities.

- (3) Let  $\mathbf{p}$  be a proposition and  $c$  the context of utterance.  $\mathbf{p}$  is uninformative in  $c$  if it is common ground in  $c$  that  $\mathbf{p}$  is true.  $\mathbf{p}$  is trivial in  $c$  if either  $\mathbf{p}$  is uninformative in  $c$  or  $\neg\mathbf{p}$  is uninformative in  $c$ .

As Lewis (1988) stressed, trivial statements are always relevant<sup>5</sup> and those that express uninformative propositions can of course always be asserted truthfully by discourse participants.<sup>6</sup> Consequently, even if silence is left aside, an unopinionated speaker is expected to always have the option of asserting an uninformative statement to express ignorance. Yet what we observe is that uninformative utterances are perceived as uncooperative across the board, even as vehicles to convey ignorance. I will argue that this observation, quite independently of the question of why silence is perceived as unacceptable, can be used to motivate (1).

## 2 Uninformative Utterances

From now on, I will assume that silence is not a legitimate conversational move. My goal is to show that regardless of why this might be the case, an argument for the claim that relevance is closed under speakers' beliefs can be built using uninformative utterances.

- (4) **Stipulation.** Saliency is not a legitimate conversational move (for whatever reason).

Now consider the following dialogues.

- (5) a. Who is John's wife?  
b. #John is married.
- (6) a. Is John in Paris?  
b. #Either he's in Paris or he's not.

It is strongly felt that the speaker  $S$  is not addressing the question asked and is therefore being uncooperative. One can imagine various explanations for why this might be so. Perhaps  $S$  has reasons to be evasive or perhaps he is being facetious. Regardless, it is very intuitive that  $S$ 's response falls beyond the cooperative norm because it is uninformative. In (5), for example, the context that results from the target question is a context in which it is taken for granted that John is married. Consequently  $S$ 's response fails to communicate any new information. The same observation applies to (6), with the only difference being that in the latter case  $S$ 's utterance

<sup>5</sup>This is a consequence of the assumption that relevant is closed under negation and conjunction. If  $\mathbf{p}$  is a relevant proposition then so are  $\neg\mathbf{p}$ ,  $\mathbf{p} \wedge \neg\mathbf{p}$  and  $\neg(\mathbf{p} \wedge \neg\mathbf{p})$ . The latter two of course express trivialities,  $\perp$  and  $\top$ . It follows that any proposition that is contextually equivalent with  $\perp$  or  $\top$  has to be relevant, since it is a truism that relevance does not distinguish between contextually equivalent propositions (Magri, 2009).

<sup>6</sup>If it is common ground in  $c$  that  $\mathbf{p}$  is true then every discourse participant believes that  $\mathbf{p}$  is true. Therefore, every discourse participant can assert  $\mathbf{p}$  truthfully.

is tautological to begin with and is therefore uninformative in any context. Stalnaker's (1974) Principle of Informativity (PI), then, seems to supply the relevant characterization of what's wrong with *S*'s response.

- (7) **Principle of Informativity.** Don't assert something that's already known to be true.

Stalnaker formulated PI in the context of introducing the common ground as a key concept in pragmatic theory. From that perspective, it is perhaps tempting to reason as follows. The goal of assertion is to change the common ground. Uninformative utterances fail to change the common ground and therefore are naturally viewed as defective on the ground of vacuity. However, postulating PI as a primitive principle of pragmatic theory along the lines just sketched seems unsatisfactory. For note that all utterances, including uninformative ones, change the common ground. At the bare minimum, once an utterance is made the common ground is changed to reflect the fact that that utterance was made. Consequently, what's special about uninformative utterances is not that they fail to change the common ground. Rather, what's special about them is that they fail to change the common ground *by introducing new information on the basis of their semantic content*. If so then it seems plausible that PI should be thought of as a consequence of norms of cooperative information-exchange.

- (8) **Desideratum.** A good theory of the norms of cooperative information-exchange ought to have PI as a derived consequence.

There are two approaches to this desideratum, as far as I can see. One option is to think of PI as a consequence of the principle of Brevity. This is what Stalnaker (1974) seems to suggest: "I will not say things that are already taken for granted, since that would be redundant". But exactly what makes uninformative utterances redundant? Redundancy is a matter of comparison. An utterance is redundant if all else equal there is a better, i.e. in some sense simpler, conversational move that could've been made to achieve the same outcome. But in the case of uninformative utterances, it is unclear if such a simpler alternative exists in general. In any context, there are many uninformative sentences that cannot be further simplified in any obvious, formal sense. Consider e.g. the sentence "*it is raining*" in a context in which it is taken for granted that it is raining. What makes such maximally simple yet uninformative utterances redundant? As far as I can see, the inevitable response is to say that utterances generally compete with silence. Suppose silence is the simplest alternative to any speech act. If so then it is perhaps reasonable to think that, in the context of cooperative information-exchange, silence is preferred to an uninformative contribution. But we are operating on the assumption that silence is not a legitimate conversational move to begin with. Quite regardless of why this might be so, if silence is not a legitimate conversational move then it can never be the preferred conversation move. I conclude that, given (4), the attempt to derive PI from Brevity is doomed.

The remaining option is to think of PI as a consequence of Quantity. Let us focus on the following statement which is modeled after what Fox (2016) refers to as Virtual Truism (VT).

- (9) As speakers we are required to say
- a. only things that we believe to be true (Quality) that are relevant (Relevance),
  - b. and everything that we believe to be true that is relevant (Quantity).

The idea, then, is that VT has PI as a consequence. Now for opinionated speakers this is certainly true. If *S*'s beliefs settle some proposition that's made relevant by *Q* then *S* can't assert an uninformative sentence without violating Quantity because if he does so, there is something that is relevant that he believes to be true that he is not saying. It follows that as far as opinionated speakers are concerned, PI can be viewed as a direct consequence of VT. But what about unopinionated speakers?

On standard assumptions, nothing prevents an unopinionated speaker to reveal his ignorance using an uninformative utterance. Specifically, the unopinionated speaker does not violate VT if he were to utter an uninformative sentence. More specifically, he is not violating Quality because it is impossible for a participant not to believe an uninformative proposition (see fn. (6)), he is not violating Relevance because it is impossible for an uninformative statement to be irrelevant (see fn. (5)), and he is not violating Quantity because by assumption he is unopinionated about all relevant propositions that are non-trivial and therefore the uninformative utterance *is* the most informative and relevant proposition that he can truthfully assert. The standard assumptions, then, lead us to expect that using uninformative utterances is a legitimate way of communicating ignorance. This is an incorrect expectation, of course. Even if I have no idea who John's wife is, I cannot respond to "*Who is John's wife?*" with "*John is married*". If I do so, my response will be understood as an attempt to evade the question instead of revealing ignorance. So what prevents using uninformative utterances as vehicles to communicate ignorance?

If we assume that relevance is closed under speakers' beliefs, we get a handle on this problem in a manner that parallels Fox's (2016) treatment of silence. If relevance is closed under speakers' beliefs (as well as conjunction and negation) then a question like "*Who is John's wife?*" makes relevant, among other things, whether the speaker knows who John's wife is. Therefore even the unopinionated speaker is in a position to assert a non-trivial yet relevant statement, namely "*I don't know*". Stated differently, closing relevance under speakers' beliefs makes genuine unopinionatedness impossible: There are always relevant and informative propositions that any given speaker is opinionated about, namely those pertain to his own beliefs. Consequently, so long as Quantity is operative, there is never any justification for making an uninformative statement.<sup>7</sup> I conclude that PI can be viewed as a consequence of VT, specifically Quantity, on the assumption that relevance is closed under belief.

Of course, this is not a conclusive argument. It is conceivable that there are other modifications of standard assumptions that would allow us to derive PI. Let us therefore look at a further prediction that this approach to PI makes. If PI is fundamentally a consequence of Quantity then we predict that in contexts in which Quantity is inoperative, PI can be violated. As Fox (2014) argues, hints are a good case-study of this sort. By definition, the purpose of a hint is to guide the audience's attention in a manner helps them solve some problem, say, to answer a certain question. A hint is not expected to be maximally informative because that would defeat the purpose. The

<sup>7</sup>Of course the speaker's beliefs with regard to a particular question may be common ground. If so then propositions that pertain to the speaker's relevant beliefs will end up being uninformative. For example, if it is already common ground that the speaker *S* doesn't know who John's wife is (or if it is common ground that *S* believes John's wife is, say, Mary) then there really is no non-trivial and relevant proposition that *S* can (truthfully) contribute in response to the question "*Who is John's wife?*". However, from the stand point of cooperative information-exchange, the problem with such cases is not *S*'s response but the fact that the question was asked of him to begin with. After all, what's the point of asking a question of someone if we already know what that person's answer would be?

point of a hint is to help others to arrive at the answer, not to give them the answer. If so then we predict that a good hint may very well fail to deliver *any* new information. This seems correct.

- (10) [Context: *There is a prize in one of the three boxes. The other two boxes are empty. Each box is subject to the following rule: If it is empty then the box that it points to (indicated with the arrow) contains the prize. Which box contains the prize?*]

Hint: Either Box 1 is empty or Box 3 is empty.

[Box 1 →] [Box 2 →] [Box 3 ←]

The hint provided in the example above is of course uninformative. Since it is common ground that two of the boxes are empty, Box 1 and Box 3 cannot both be non-empty. Nevertheless, the hint is judged as fully acceptable and quite appropriate, i.e. helpful. This observation corroborates the claim that PI and Quantity go hand in hand.

### 3 Strengthening the argument

Above I have argued that there is motivation for the claim that PI is rooted in VT, specifically Quantity. I also argued that PI can be derived from Quantity if it is assumed that relevance is closed under speakers' beliefs. In this section I'd like to strengthen the main conclusion by addressing a way in which the argument it behind may be attacked.

The question I posed in the previous section was this.

- (11) Why can't "*John is married*" be asserted in order to communicate that the speaker doesn't know who John's wife is in a context in which it is taken for granted that John is married?

Now, one may wonder if this question is the right question to ask to begin with. For note that there is no evidence whatsoever that a sentence like "*John is married*" can *ever* be used to communicate ignorance about who John is married to *regardless of whether the proposition that John is married is common ground or not*. In other words, even in contexts in which the proposition that John is married is informative, the sentence "*John is married*" cannot be used to communicate ignorance. To see this, consider the following dialogue (inspired by Feinmann 2023).

- (12) [Context: *It is common ground that either John is single or he is married to Ann or Sue.*]
- a. **A:** Is John single or is he married to Ann or Sue?
  - b. **B:** He is married.
  - c. **A:** OK, but is he married to Ann or Sue?!

In the context of (12), **B**'s response in (b) is of course informative. Yet it still fails to communicate ignorance and is most likely perceived as an attempt by **B** to withhold information, whence the felicity of the reaction in (c). Compare with (13).

- (13) [Context: *It is common ground that either John is single or he is married to Ann or Sue.*]
- a. **A:** Is John single or is he married to Ann or Sue?
  - b. **B:** He is married, that's all I know!
  - c. **A:** #OK, but is he married to Ann or Sue?!

In the absence of evidence that a sentence like "*John is married*" can ever be used to communicate ignorance, one may argue that the question to ask is (14) rather than (11).

- (14) Why can't "*John is married*" be asserted in order to communicate that the speaker doesn't know who John's wife is *in any context*?

The argument I discussed in the previous section made crucial reference to un informativity, of course. That is, I have an answer to (11) but not (14). Consequently, one might argue that the conclusion I argue for, i.e. closing relevance under speakers' beliefs, is not justified because (i) the bigger problem, namely (14), remains open, (ii) the answer to (14) will resolve (11) as well, and (iii) as things stand there is no reason to think that the answer to (14), whatever it may be, will also come with the commitment that relevance must be closed under speakers' beliefs. Let me be clear that I do not have an answer to (14). Instead, what I'd like to do is to show is that,

- (15) It is possible to construct a sentence  $\phi$  such that,
- a. It is uncontroversial that  $\phi$  can be used to communicate ignorance but
  - b.  $\phi$  cannot be used to communicate ignorance if it is otherwise uninformative.

If this claim is true then the argument made in the previous section goes through regardless of the proper treatment of questions like the one in (14). To see that this claim *is* true, consider the following dialogue first.

- (16) [*Context: It is common ground that either John is single or he is married to Sue.*]
- a. **A:** Is John single or is he married to Sue?
  - b. **B:** #Either he is single or he is married to Sue.

The assertion in (b) is contextually uninformative and therefore unacceptable much like the example I discussed before. Note, however, that unlike the previous example it is uncontroversial that a disjunctive sentence like "*John is single or he is married to Sue*" can be used to communicate ignorance about the individual disjuncts.

- (17) [*Context: It is common ground that either John is single or he is married to Ann or Sue.*]
- a. **A:** Is John single or is he married to Ann or Sue?
  - b. **B:** Either he is single or he is married to Sue.
  - c. **A:** #OK, but is he single or is he married to Sue?!

I conclude that while there is an open problem pertaining to why certain sentences fail to trigger certain kinds of ignorance inferences even in highly favorable contexts (see Feinmann 2023), it is also true that sentences that do very easily trigger ignorance inferences in general fail to serve as acceptable means of communicating *mere* ignorance, that is, if they are otherwise uninformative.

## 4 A loose end

If the argument in the previous sections is successful then we need to deal with the consequences of closing relevance under speakers' belief. As Fox (2016) observes, the most immediate consequence is that the grammar now needs to have a way to settle the speaker's epistemic attitude towards various relevant alternatives in order to meet the demand in (2). Meyer's (2013) Matrix K hypothesis is the most obvious way of satisfying this desideratum.

(18) **Matrix K Hypothesis.** An assertion of  $\phi$  by  $S$  may be parsed as  $\mathbf{K}_s\phi$  or  $exh\mathbf{K}_s\phi$  at LF.

For a sentence like “*John is single or he is married to Sue*” the following parse, among others, is available (since the disjuncts are mutually incompatible, I ignore exhaustification below  $\mathbf{K}$ ).

(19)  $exh\mathbf{K}_s$  [SINGLE or SUE]

Assuming that SINGLE and SUE are both relevant,  $\mathbf{K}_s$ SINGLE and  $\mathbf{K}_s$ SUE are relevant too. Since the latter also happen to be formal alternatives of  $exh$ ’s prejacent, the LF above is predicted to be true iff the speaker believes that either John is single or he is married to Sue but is otherwise ignorant (see Meyer 2013 for details). Therefore, this LF is well-equipped to meet the demand set up by (2) in many contexts. The loose end I’d like to make explicit is that (19) is, unfortunately, predicted to satisfy (2) even in contexts in which it is common ground that John is either single or married to Sue. That is, we predict the target sentence to be acceptable both in example (16) and example (17) discussed in the previous section. Our job is not done.

Since I’m running out of space, let me merely point at a way to think about this problem. Brian Buccola has observed (p.c. circa 2018) that the following contrast is surprising. To appreciate the puzzle in its purest form, the reader may ignore the Matrix K hypothesis.

- (20) a. #Every grandparent of mine is either dead or alive.  
 b. Every grandparent of mine is either dead or sick.

Intuitively, one might attribute this contrast to the fact that (a) is uninformative while (b) is not. After all all humans, including the speaker’s grandparents, are either dead or alive but it is of course quite conceivable that the speaker has a grandparent who is neither dead nor sick. But note that (b) is most naturally understood with the inference that at least one of the speaker’s grandparents is dead and at least one of them is sick. Very plausibly, these inferences are due to exhaustification (see Bar-Lev and Fox 2023 and references therein). Regardless of the details, by parity of reasoning we should expect that (a) can also be parsed to entail that at least one of  $S$ ’s grandparents is dead and at least one of  $S$ ’s grandparents is alive. But of course this would make (a) informative, and the explanation of its degradedness based on un informativity is undermined. What is needed, evidently, is a theory of exhaustification that can derive the so-called “distributive inferences” in the case of (b) but not (a). But the only difference between the two cases seems to be that the prejacent of  $exh$  in the case of (a) but not (b) is contextually uninformative. Stated differently, there is motivation for a generalization along the following lines.

(21) **Generalization.** An expression of the form  $[exh \alpha]$  is unacceptable in context  $c$  if  $\alpha$  is uninformative  $c$ .

From this generalization it follows that the parse in (19) is unavailable in any context in which it is common ground that if John is not single then he is married to Sue, as desired. In such contexts, the prejacent of  $exh$  denotes a trivial proposition (recall if  $\mathbf{p}$  is common ground and  $S$  is a discourse participant then by definition  $\mathbf{K}_s\mathbf{p}$  is common ground as well). For some ways to think about the generalization in (21) see section 7.2 of Bar-Lev and Fox (2023).

## References

- Adamson, Lauren B., and Janet E. Frick. 2003. The Still Face: A History of a Shared Experimental Paradigm. *Infancy* 4:451–473.
- Bar-Lev, Moshe E., and Danny Fox. 2023. On fatal competition and the nature of distributive inferences. *Natural Language Semantics* 31:315–348.
- Buccola, Brian, and Andreas Haida. 2019. Obligatory Irrelevance and the Computation of Ignorance Inferences. *Journal of Semantics* 36:583–616.
- Feinmann, Diego. 2023. Ignorance is a problem. *Snippets* .
- Fox, Danny. 2014. Cancelling the Maxim of Quantity: Another challenge for a Gricean theory of Scalar Implicatures. *Semantics and Pragmatics* 7.
- Fox, Danny. 2016. On why ignorance might be part of literal meaning—commentary on Marie Christine Meyer. MIT workshop on exhaustivity, September 2016.
- Lewis, David. 1988. Relevant implication. *Theoria* 54:161–174.
- Magri, Giorgio. 2009. A theory of individual-level predicates based on blind mandatory scalar implicatures. *Natural Language Semantics* 17:245–297.
- Meyer, Marie-Christine. 2013. Ignorance and Grammar. Doctoral Dissertation, MIT.
- Stalnaker, Robert. 1974. Pragmatic presuppositions. In *Semantics and Philosophy*, ed. M. K. Munitz and P. K. Unger. New York: New York University Press.