

Obligatory implicatures and the relevance of contradictions*

Moshe E. Bar-Lev

TAU

March 5, 2023

Abstract

Magri (2009a,b) proposed a generalization according to which a sentence is infelicitous whenever exhaustification over the full set of formal alternatives of the sentence leads to contextual contradiction. While Magri proposes an account of obligatory implicatures which explains some cases where this generalization expects infelicity, he does not provide a general account of this generalization. In this paper I argue for a perspective on the ‘pruning’ of alternatives which predicts this generalization, building on the counter-intuitive idea that contradictions are relevant in every context (Lewis 1988). I further argue, using disjunction in the scope of a universal quantifier as a test case, that an extension of this view to obligatory ignorance inferences provides a new perspective on the Logical Integrity Generalization put forward by Anvari (2018b), while avoiding some empirical problems for this generalization.

Keywords: implicature, exhaustification, pruning, contradiction, alternatives, ignorance inferences

1 Introduction

(1) is an infelicitous sentence given that all Italians come from the same country (Magri 2009a,b, 2011; Spector 2014; Katzir and Singh 2015; Anvari 2018b; Marty 2017, a.o.). Magri (2009a,b) argues based on the infelicity of sentences like (1) that the system of implicature computation is modular, that is, it only has access to logical entailment relations (\Rightarrow_L) and not to contextual entailment relations (\Rightarrow_C). On this view, the infelicity of (1) is due to an obligatory implicature derived by negating the alternative in (2), which leads to the overall meaning in (3), which is in turn a contextual contradiction.

*This paper benefited tremendously from endless discussions with Milica Denić, Danny Fox, Roni Katzir, and especially Amir Anvari, who collaborated with me on an earlier version of the ideas developed in this paper. I would also like to thank the audiences at Tel Aviv University and at Heinrich-Heine-Universität Düsseldorf for feedback.

- (1) #Some Italians come from a warm country.
- (2) #All Italians come from a warm country.
- (3) #Some but not all Italians come from a warm country.

Specifically, Magri assumes the grammatical theory of scalar implicatures (Chierchia et al. 2012) where an exhaustivity operator $\mathcal{E}xh$, which only cares about logical entailment relations, is responsible for the derivation of implicatures by negating certain alternatives. He further shows that, with some more or less reasonable assumptions (which we will go over in detail in §2), this theory makes the prediction in (4), where $\mathcal{E}xh_{Alt(S)}(S)$ is the result of exhaustification over S given its full set of formal alternatives $Alt(S)$ (in §3 we will properly define $\mathcal{E}xh$; for now it will suffice to simply assume that $\mathcal{E}xh_{Alt((1))}((1))$ entails the negation of (2)).¹

(4) **Narrow Magri generalization (NMG):**

A sentence S is infelicitous in context C if there is an alternative $A \in Alt(S)$ s.t.

- a. $\mathcal{E}xh_{Alt(S)}(S) \Rightarrow_L \neg A$
- b. $S \Leftrightarrow_C A$

While the NMG is successful in predicting some cases of infelicity, such as (1), it fails to predict other cases of infelicity which share a family resemblance, leading to stronger generalizations than the one predicted by Magri’s theory. I will focus in this paper on two such generalizations. The first generalization is a natural extension of (4), what I will call here the Broad Magri Generalization (BMG) in (5) (which Magri calls the “Mismatch Hypothesis”).

(5) **Broad Magri Generalization (BMG):**

If the blind strengthened meaning of a sentence S is a contradiction given common knowledge (i.e. $\mathcal{E}xh_{Alt(S)}S \Leftrightarrow_C \perp$), then sentence S sounds odd.

According to (5), whenever applying $\mathcal{E}xh$ given the full set of formal alternatives yields a contextual contradiction, the sentence is infelicitous. This is stronger than (4): If no alternative is contextually equivalent to the sentence but the joint exclusion of several alternatives leads to a contradiction, BMG expects infelicity but NMG doesn’t. Support for BMG then comes from the infelicity of (6a), (a modification of) an example due to Magri (2009a,b) which he argues has precisely this property (in contrast with (6b)). In §4.1 I will discuss this argument in detail, where it will be part of a more general discussion of the alternatives and implicatures of sentences with disjunction in the scope of a universal quantifier.

(6) *Context: A competition lasted five days, Monday through Friday; both John and Bill know that*

¹ (4) may look superficially different than how the predictions of Magri’s theory are sometimes presented; this is done in order to facilitate comparison with other generalizations.

the same person won on each of the five days. John wants to know more about this amazing person and thus asks Bill for more information; Bill knows that this person was either Mary or Sue but he doesn't know which one of them she is, so he provides the following information:

- a. #On every day, Mary or Sue won.
- b. Mary or Sue won on every day.

Another generalization which is stronger than NMG and which this paper will be concerned with is the Logical Integrity Generalization (LIG) proposed by Anvari (2018a,b), according to which a sentence is infelicitous if it contextually, but not logically, entails one of its alternatives. Anvari (2018a,b) has argued that LIG has better empirical coverage than NMG, and later in this paper we will present some novel evidence for it (see §4.2). Since this generalization does not look like anything directly connected to implicature computation (unlike BMG), I will defer a detailed discussion of it to §3, where I will argue that, despite appearances, it can also be taken to be an extension of Magri's view, much like BMG is.

There is then preliminary reason to think that both BMG and LIG have wider empirical coverage than NMG. However, unlike NMG, which follows from a particular theory of the derivation of implicatures (see §2), no explanatory account of BMG or LIG has been proposed yet. In other words, NMG follows from an explanatory account but is empirically too weak; BMG and LIG are empirically motivated, but do not follow from an explanatory account.

My goals in this paper are:

1. To propose a modification of Magri's account which predicts BMG.
2. To provide a new perspective on LIG by restating it in terms of mandatory ignorance inferences.
3. To develop a theory based on the modification of Magri's theory and the restatement of LIG which comes close to predicting LIG.
4. To argue that this theory is empirically advantageous over LIG by showing that it rules in some cases which are wrongly ruled out by LIG.

The paper is structured as follows. In §2 I present a modification of Magri's theory where BMG follows from three assumptions: (i) $\mathcal{E}xh$ applies obligatorily (following Magri); (ii) ignoring alternatives ('pruning') is only allowed if the meaning derived without pruning is irrelevant (Bar-Lev 2018); and (iii) contradictions are always relevant (Lewis 1988). In §3 I restate LIG in terms of mandatory ignorance inferences, a restatement which will make its connection with Magri's theory and BMG more transparent. Then, in §4, I provide evidence for both BMG (based on (6)) and for LIG. In §5 I propose a way in which the theory presented in §2 can be modified in order to account for cases that motivate LIG, while pointing out that it does not fully predict LIG. §6 however argues that this is a welcome result, based on empirical problems for LIG which are captured by the theory proposed in §5.

2 Proposal: Pruning and the relevance of contradictions

2.1 Magri's theory

As is known since Horn (1972), there has to be a way to ignore ('prune') some alternatives in the derivation of implicatures, and this process of pruning should depend on the context of utterance. For instance, if we are interested to know how many students came to the party, then we are likely to infer from (7) that not many students came; but if all we want to know is whether all, some, or none of the students came, (7) does not have this inference any more, but of course it still has the inference that not all students came. Whether or not we consider an alternative like *many students came* then seems highly context dependent.

(7) Some students came to the party.

What alternatives can be ignored and exactly under what conditions still remains a largely open issue (see Fox and Katzir 2011; Katzir 2014; Crnič et al. 2015). One constraint which (as far as I'm aware) has not been challenged, though, is that alternatives which are relevant to the Question Under Discussion (QUD, Roberts 1996) cannot be ignored. Magri then assumes the following constraint on pruning which avoids pruning of relevant alternatives:

(8) **Constraint on pruning** (1st version, to be revised):

$\mathcal{E}xh_{Alt'} S$ is licensed for $Alt' \subseteq Alt(S)$ given a question Q only if all alternatives in $Alt(S)$ which are relevant given Q are in Alt' .

Magri points out the following property that sentences like (1) have: Since the sentence (*some Italians come from a warm country*) and its alternative (*all Italians come from a warm country*) are contextually equivalent, there can be no question for which one of them is relevant and the other is not, as long as questions are taken to be partitions of the context set. As a result, either both the sentence and its alternative are taken to be irrelevant, or both are taken to be relevant. In the former case, the Gricean Maxim of Relevance is violated, because the speaker has uttered a sentence which is irrelevant to the QUD. In the latter case, Magri (2009a) proposes, an obligatory implicature is derived which leads to a contextual contradiction, and once again to infelicity. In order to ensure that this implicature cannot be avoided, Magri assumes that exhaustification obligatorily applies (at least at matrix position, though see Magri 2011 and later in §5 for the claim that it applies at every scope site). The predictions of this theory are then as in (9) (repeated from (4)).

(9) **Narrow Magri generalization** (NMG, repeated from (4)):

A sentence S is infelicitous in context C if there is an alternative $A \in Alt(S)$ s.t.

a. $\mathcal{E}xh_{Alt(S)}(S) \Rightarrow_L \neg A$

b. $S \Leftrightarrow_C A$

Crucially, the underlying reason for infelicity for Magri is contextual equivalence. If no alternative which is negated by $\mathcal{E}xh$ when it is in the set of alternatives is contextually equivalent to the sentence, no infelicity is expected. This is because we can consider a QUD for which the sentence would be relevant but the alternative would not; in this case it should be possible to prune the alternative and avoid deriving a contextual contradiction and infelicity.

2.2 Rethinking pruning and relevance

Magri relies on a relatively straightforward way of letting Relevance constrain pruning: We look at each alternative in isolation and ask whether it's relevant to the QUD or not: If it is, it must be in the set of alternatives $\mathcal{E}xh$ takes as argument, otherwise it doesn't. But one can think of different ways of making the connection between Relevance and pruning. My goal in the rest of this section is to argue that adopting an alternative perspective on this connection, proposed in Bar-Lev (2018, 2021), has the advantage of predicting BMG (when combined with Lewis's observation that contradictions are relevant, which we will discuss in §2.3). On this view, in order to decide what choices of pruning are licit we don't look at each alternative in isolation and ask whether it's relevant, but rather we look at different pruning choices and ask whether they yield relevant propositions when we apply $\mathcal{E}xh$. This allows for restricting pruning to no more than necessary for deriving a relevant meaning, as in (10).²

(10) **Constraint on pruning** (revised from (8)):

$\mathcal{E}xh_{Alt'S}$ is licensed for $Alt' \subseteq Alt(S)$ given a question Q only if Alt' is a maximal subset of $Alt(S)$, s.t. $\mathcal{E}xh_{Alt'S}$ is relevant given Q .

The constraint in (10) says that we can prune alternatives in order to get a relevant meaning, but no more than that; if there is a set Alt'' such that $Alt' \subset Alt'' \subseteq Alt(S)$, and $\mathcal{E}xh_{Alt''S}$ yields a relevant meaning, then $\mathcal{E}xh_{Alt'S}$ is not licensed.³ What matters for our purposes is the following consequence of (10):

² The constraint in (10) does not rule out pruning choices that lead to 'symmetry breaking' (see Fox and Katzir 2011; Katzir 2014; Crnič et al. 2015), but it can be amended following Bar-Lev so that it does, by integrating Crnič et al.'s constraint on pruning according to which pruning can only lead to weakening, as in (i) (see Crnič et al. 2015 for why this blocks symmetry breaking):

- (i) $\mathcal{E}xh_{Alt'S}$ is licensed for $Alt' \subseteq Alt(S)$ given a question Q only if Alt' is a maximal subset of $Alt(S)$, s.t.
- a. $\mathcal{E}xh_{Alt'S}$ is relevant given Q , and
 - b. $\mathcal{E}xh_{Alt(S)S} \Rightarrow \mathcal{E}xh_{Alt'S}$

Once this amendment is adopted, one can think of this constraint as aiming to maximize informativity without reaching irrelevance. For ease of exposition I ignore pruning choices which involve symmetry breaking in this paper.

³ Bar-Lev's motivation for the constraint in (10) comes from considerations having to do with so-called non-maximal readings of sentences containing definite plurals such as *the kids smiled*, which in certain contexts are judged true even if not all the kids smiled. He argues for a theory where non-maximality is the result of pruning; (10) aims to avoid having non-maximal readings in cases where stronger readings are relevant.

(11) If $\mathcal{E}x_{Alt(S)}S$ is relevant given Q , then $\mathcal{E}x_{Alt'}S$ is licensed only if $Alt' = Alt(S)$.

That is, we can only prune alternatives when the result of exhaustification with the full set of alternatives $Alt(S)$ yields an irrelevant meaning; if it yields a relevant meaning, then no pruning is allowed and the only set of alternatives that can be chosen is one which contains all the alternatives in $Alt(S)$. In what follows I will demonstrate how this constraint predicts BMG when taken together with a surprising result pointed out by Lewis (1988) according to which contradictions are relevant.

2.3 The relevance of contradictions

What does it mean for a proposition to be relevant given a question, where a question is a partition of the context set? I will assume a standard definition following Lewis (1988) who defines his notion of Aboutness, equated here with Relevance, as follows:

When we think of subject matters as partitions, we can say that P is about M [= relevant given M] iff each cell of M either implies or contradicts P (Lewis 1988: p. 163).

An equivalent definition of Relevance given a question is as follows:

(12) **Relevance of propositions given Q :**

A proposition p is relevant given a partition Q iff $\exists Q' \subseteq Q [p = \bigcup Q']$

Lewis points out that this simple view of Relevance has the immediate consequence that contradictions (as well as tautologies) are always relevant:

(13) **The relevance of contradictions:**

Contradictions are relevant to any subject matter whatsoever.

The reason for this is that no matter how the question Q partitions the context set, a contradiction will contradict each cell in Q (and a tautology will be entailed by each cell). Put differently, having the definition in (12) in mind, it will always be possible to find a subset of Q whose union is a contradiction, given that the empty set must be a subset of Q . As Lewis himself notes, this result is admittedly counter-intuitive: “It is a surprising consequence, no doubt” (Lewis 1988: p. 165).

Of course, this counter-intuitive result could be avoided by redefining Relevance so that contradictions are made irrelevant.⁴ Lewis (1988) however argues on conceptual grounds that his notion of relevance should be maintained even though it has the apparently problematic consequence in

⁴ For instance, one could revise (12) as follows:

(i) A proposition p is relevant given a partition Q iff $\exists Q' \subseteq Q [Q' \neq \emptyset \wedge p = \bigcup Q']$

(13).⁵ In the next subsection I will argue on empirical grounds that the consequence in (13) is in fact a desired one, by demonstrating that it predicts BMG when taken together with the constraint on pruning in (10).

2.4 BMG follows

Let us now show that assuming with Magri that exhaustification applies obligatorily, BMG (repeated in (14)) follows directly from the two assumptions made above: (i) that pruning more than necessary to achieve a relevant meaning is impossible ((10)), and (ii) that contradictions are relevant ((13)).

(14) **Broad Magri Generalization** (BMG, repeated from (5)):

If the blind strengthened meaning of a sentence S is a contradiction given common knowledge (i.e. $\mathcal{E}xh_{Alt(S)}S \Leftrightarrow_C \perp$), then sentence S sounds odd.

To see why BMG follows from these assumptions, consider what happens in situations where $\mathcal{E}xh_{Alt(S)}S \Leftrightarrow_C \perp$. In such cases, the meaning with the full set of alternatives is contextually equivalent to a contradiction, which, given (13), renders it relevant (no matter what the QUD is). Since the meaning exhaustification with the full set of alternatives yields is relevant, the constraint on pruning in (10) blocks pruning of any alternative, and only licenses $\mathcal{E}xh_{Alt'}S$ if $Alt' = Alt(S)$ (recall the consequence in (11)), ensuring that no meaning which is weaker than a contradiction can be derived if $\mathcal{E}xh$ applies. And assuming (following Magri) that a parse with no $\mathcal{E}xh$ is unavailable, the sentence will only have a contextually contradictory meaning, and thus it is predicted to be infelicitous.

2.5 Interim summary

In this section I proposed a theory where Magri's constraint on pruning which considers the relevance of each alternative in isolation is replaced with a more holistic constraint on pruning which blocks pruning if the meaning we get by exhaustification over the full set of alternatives is relevant. I have shown that this theory predicts BMG when combined with Lewis's observation that contradictions are relevant and Magri's assumption that applying $\mathcal{E}xh$ is obligatory at matrix position.

3 Logical Integrity and mandatory ignorance inferences

So far I focused on BMG, and proposed an account which predicts it. Recall from §1 that another prominent generalization has been discussed in the literature in recent years: The Logical Integrity

⁵ He suggests the following reasoning: Tautologies should be relevant because they do not provide any irrelevant information; and if something is relevant, its negation should be relevant as well. So if tautologies are relevant, then contradictions should be relevant as well.

Generalization (LIG) due to Anvari (2018a,b). The rest of this paper is dedicated to examining the empirical and theoretical connection between BMG and LIG, and to showing how the proposal made in the previous section can make sense of cases that motivate LIG, while avoiding its shortcomings.

Consider NMG, repeated in (15), and Anvari’s LIG in (16).

(15) **Narrow Magri generalization** (NMG, repeated from (4)):

A sentence S is infelicitous in context C if there is an alternative $A \in Alt(S)$ s.t.

- a. $\mathcal{E}xh_{Alt(S)}(S) \Rightarrow_L \neg A$
- b. $S \Leftrightarrow_C A$

(16) **Logical Integrity** (LIG):

A sentence S is infelicitous in context C if there is an alternative A in $Alt(S)$ s.t.

- a. $S \not\Rightarrow_L A$
- b. $S \Rightarrow_C A$

LIG does not immediately look like a natural extension of NMG in any way, unlike BMG; in fact it may appear to be entirely disconnected from theories of implicature computation. As I will aim to show in this section, however, LIG has a rather close connection to implicature computation despite appearances, as it can be restated in terms of mandatory ignorance inferences. This would also make the connection between BMG and LIG pretty obvious: LIG will follow from a slightly different version of the extension of Magri’s theory proposed in the previous section (which, as I have shown, predicts BMG).

A first attempt at characterizing the connection between LIG and theories of implicatures could go as follows: If an alternative A is not logically entailed by S , then it will be excluded by exhaustification, and as a result, in cases where A is contextually entailed by S , the result would be a contextual contradiction. While this reasoning may work for many cases, it ignores the possibility that an alternative would not be logically entailed but would still not be excluded by exhaustification. A case in point is simple disjunction: A isn’t logically entailed by $A \vee B$, but exhaustification cannot exclude it: *Mary was singing or dancing* can never have the inference that Mary was not singing. In other words, there are cases where alternatives are not logically entailed but are still not Innocently Excludable, using Fox’s (2007) terminology. In such cases, if A is contextually entailed by S , LIG would predict infelicity but BMG won’t. For concreteness, then, let us define $\mathcal{E}xh$ so that it only negates Innocently Excludable alternatives, as in Fox (2007).

(17) Definition of $\mathcal{E}xh$:

- a. $\llbracket \mathcal{E}xh \rrbracket(C)(p)(w) = 1$ iff $p(w) = 1 \wedge \forall q \in IE(C, p)[q(w) = 0]$
- b. $IE(C, p) = \bigcap \{C' : C' \text{ is a maximal subset of } C \text{ s.t. } p \wedge \bigwedge \{\neg q : q \in C'\} \Leftrightarrow_L \perp\}$

Indeed, one of the arguments for LIG in Anvari (2018a) comes from Hurford disjunctions such as

(18), which have exactly this property: the alternative *Mary lives in France* is contextually but not logically entailed by (18), but it cannot be excluded by $\mathcal{E}xh$ due to the general inability of $\mathcal{E}xh$ to exclude A given a sentence like $A \vee B$, or, in other words, because A is not Innocently Excludable relative to $A \vee B$.

(18) #Mary lives in France or in Paris.

Blind exhaustification over the full set of alternatives then does not entail that Mary doesn't live in France, and as a result no contextual contradiction is derived and BMG does not account for the infelicity of (18). LIG, on the other hand, captures its infelicity, because *Mary lives in France* is contextually but not logically entailed by (18) (in fact it is contextually equivalent to (18)).

But this is not to say that the infelicity of (18) cannot be accounted for within theories of implicature computation. Singh (2010) and Meyer (2014) derived the infelicity of (18) based on ignorance inferences: While sentences of the form $A \vee B$ do not have the inference that A is false, they do have the inference that the speaker is ignorant about A . Applied to (18), this derives the inference that the speaker is ignorant about whether *Mary lives in France*, which results in a contextual contradiction, because a speaker cannot be certain that (18) is true without also being certain that *Mary lives in France* is true.

There is much more to say about the source of the infelicity of Hurford disjunctions, and both approaches we just outlined—based on LIG and based on ignorance inferences—face difficult problems (see especially Marty and Romoli 2022; Anvari 2022) which I will not aim to resolve here.⁶ I would however like to point out the close connection between the two approaches to Hurford disjunctions: In both, the infelicity is due to the existence of an alternative which is contextually but not logically entailed. This is not just an accidental similarity; I would like to show now that within Meyer's (2013) grammatical theory of ignorance inferences, we can state a generalization which is very similar to the BMG and which turns out to entail LIG. This will not be an entirely good result because of arguments that I will present against LIG in §6; it will however serve as an intermediate step which will help us see the connection between LIG and theories of exhaustification.

On Meyer's (2013) theory of ignorance inferences, they are derived by applying $\mathcal{E}xh$ above a universal K (knowledge) operator, which is a universal modal paraphrasable as 'the speaker is certain that. . .' (see Fox 2016; Buccola and Haida 2019 for conceptual motivation). In order to keep things simple, let us make the simplifying assumption that the K operator applies to a sentence S whenever S is uttered (an assumption which we will revise in §6). Since (as we assumed in §2) $\mathcal{E}xh$ must apply in matrix position, it follows that any sentence S will have the parse in (19).

⁶ The infelicity of (18) may well have more than one source. Among other things, its infelicity has been attributed to considerations of redundancy, namely the fact that the same information could be conveyed with a syntactically simpler alternative (an idea playing an important role in Meyer's analysis). The data I will consider later in this section (in §4 and §6) do not have this property; as a result, using them to test the empirical status of BMG and LIG has the advantage of avoiding some of the noise associated with Hurford disjunctions like (18).

$$(19) \quad \mathcal{E}xh_{Alt'}(K(S))$$

Of course, we could also have a parse where $\mathcal{E}xh$ applies both above and below K ; in §5 I will in fact propose that this parse is the only one available (when K applies). For expository reasons, however, let us ignore for now the possibility of having another $\mathcal{E}xh$ applying below K since this will facilitate the restatement of LIG in terms of mandatory ignorance inferences.⁷

Let us assume that Alt' is the full set of alternatives of the sister of $\mathcal{E}xh$, that is $Alt(K(S))$, and (as standard) that this set consists of all the alternatives $K(A)$ where A is an alternative to S :

$$(20) \quad Alt(K(S)) = \{K(A) : A \in Alt(S)\}$$

Since K is a universal quantifier (which has no alternative), $\mathcal{E}xh$ will exclude every alternative of the form $K(A)$ if A is not logically entailed by S , and the result of excluding all of these alternatives together will be consistent; this will then be the output of exhaustification, as in (21) (see Fox and Katzir 2011 for other cases where the introduction of a universal operator allows $\mathcal{E}xh$ to negate all alternatives that are not logically entailed by the sentence).⁸

$$(21) \quad \mathcal{E}xh_{Alt(K(S))}K(S) \Leftrightarrow K(S) \wedge \bigwedge \{-K(A) : A \in Alt(S) \wedge S \Rightarrow_L A\}$$

Returning to the issue of infelicity, let us now consider a sentence S which is ruled out by LIG. Because S is ruled out by LIG, we must be able to find some alternative A so that A is contextually but not logically entailed by S . Now consider what (21) delivers in this case: Since A is not logically entailed by S , $\mathcal{E}xh_{Alt(K(S))}K(S)$ entails $K(S)$ as well as the negation of $K(A)$; but this is a contextual

⁷ Note that the parse in (19) will not derive scalar implicatures due to the absence of $\mathcal{E}xh$ below K , and, as a result, the BMG will no longer follow if it is an available parse. Assuming $\mathcal{E}xh$ above and below K in §5 will restore the prediction that the BMG should hold and will further make some intricate predictions which I will argue are borne out.

⁸ Let me explain why (21) holds in three steps: First, showing that (ia) (that is, the right hand-side of (21)) is non-contradictory; second, showing that (ia) is equivalent to (ib) and hence (ib) is also non-contradictory; and third, showing that, because (ib) is non-contradictory, (ib) (\Leftrightarrow (ia)) will be the output of exhaustification over $K(S)$ given $Alt(K(S))$.

- (i) a. $K(S) \wedge \bigwedge \{-K(A) : A \in Alt(S) \wedge S \Rightarrow_L A\}$
 b. $K(S) \wedge \bigwedge \{-A : A \in Alt(K(S)) \wedge K(S) \Rightarrow_L K(A)\}$

First, (ia) is non-contradictory because it's possible to find a set of worlds $W' \subseteq W$ such that S is true in all worlds in W' , and for every alternative $A \in NW$ (where NW is the set of all non-weaker alternatives of S , that is $\{A : A \in Alt(S) \wedge S \Rightarrow_L A\}$) there is at least one world in W' where A is false; if W' is the set of worlds compatible with the speaker's beliefs, (ia) is true. We can construct W' as follows. For every alternative A in NW , we take the set of worlds in W where S is true and A is false; this set is guaranteed to be non-empty, given that $S \Rightarrow_L A$. We can then take W' to be the union of all these sets of worlds, that is, $W' = \bigcup \{w : S(w) = 1 \wedge A(w) = 0\} : A \in NW\}$. This set of worlds has the desired property: In all worlds in this set S is true, and for each alternative $A \in NW$, there is a world in this set where A is false. Second, (ia) is equivalent to (ib), because $K(S) \Rightarrow_L K(A)$ iff $S \Rightarrow_L A$, and $K(A) \in Alt(K(S))$ iff $A \in Alt(S)$. It follows then that (ia) and (ib) are equivalent and non-contradictory. Third, because (ib) is non-contradictory, this would be result of exhaustification over $K(S)$ given $Alt(K(S))$: $\{A : A \in Alt(K(S)) \wedge K(S) \Rightarrow_L K(A)\}$ would be the the only maximal set of alternatives in $Alt(K(S))$ whose joint negation is consistent with $K(S)$, and consequently this would be the set of IE alternatives. And since we have shown that (ia) and (ib) are equivalent, (21) follows.

contradiction: one cannot possibly be certain that S is true while not being certain that A is also true, given that A is contextually entailed by S .⁹ We have shown then that whenever LIG rules out S , $\mathcal{E}xh_{Alt(K(S))}K(S)$ is a contradiction. We can then have a generalization which entails LIG and which is based on exhaustification:

(22) **Mandatory Ignorance Inferences Generalization (MIIG):**

If the blind strengthened meaning of $K(S)$ is a contextual contradiction (i.e. $\mathcal{E}xh_{Alt(K(S))}K(S) \Leftrightarrow_C \perp$), then sentence S sounds odd.

MIIG entails LIG, but they are not equivalent. In cases where S has two alternatives, A and A' , and $K(S)$ entails neither $K(A)$ nor $K(A')$ but it does entail their disjunction, LIG does not expect infelicity but MIIG does. A simple illustration of such a situation is one where a speaker is assumed to be completely knowledgeable about each of the alternatives; for example, a person is usually assumed to know where they live, which is presumably why (23) is odd (barring cases where the speaker is not expected to deliver their epistemic state, see §6):

(23) # I live in Paris or London.

MIIG then makes the correct prediction that (23) should be odd. It is not entirely clear however that this can be taken as a clear argument in favor of MIIG, since LIG might resort to other explanations of the infelicity of (23) based on the general robustness of ignorance inferences, even in cases where the speaker is not assumed to be completely knowledgeable about the alternatives. Even though MIIG and LIG are not equivalent, as we have just seen, I will treat them from now on as if they were and refer to them as the single generalization LIG/MIIG, since for all cases we will discuss in this paper they make identical predictions.

The characterization of LIG in terms of MIIG is obviously parallel to BMG, repeated once again

⁹ This holds as long as all alternatives denote bivalent propositions; whether the same result applies with trivalent ones eventually depends on how $\mathcal{E}xh$ is defined. If it excludes alternatives by assigning them non-truth rather than falsity (as entertained by Spector and Sudo 2017), the same results will apply even in the trivalent case. Such a view will have the benefit of making the theory I will propose predict the infelicity of both (i) and (ii) (due to Gajewski and Sharvit 2012; Spector and Sudo 2017) which were used by Anvari as arguments for LIG. Both of these cases are captured by LIG, because (i) contextually (but not logically) entails that John is unaware that all the students smoke, and (ii) contextually (but not logically) entails that John is unaware that Mary lives in Paris.

(i) *Context: All students smoke.*
John is unaware that some students smoke.

(ii) *Context: Mary lives in Paris.*
John is unaware that Mary lives in Paris or London.

I however note that Spector and Sudo (2017) provided arguments against letting $\mathcal{E}xh$ exclude alternatives by assigning them non-truth; any attempt to explain the infelicity of (ii) along these lines would then have to face these arguments, something which is beyond the scope of this paper. I thus leave this issue for future work, hoping that a resolution will be possible. For further relevant complications with LIG, mandatory ignorance inferences and presuppositions see Marty (2017); Marty and Romoli (2021).

in (24). The only difference between BMG and MIIG is in the kind of structure which leads to a contextual contradiction: whether it is one where $\mathcal{E}xh$ applies directly to the sentence (as in BMG), or one where K intervenes between $\mathcal{E}xh$ and the sentence (as in MIIG).

(24) **Broad Magri Generalization** (BMG, repeated from (5)):

If the blind strengthened meaning of a sentence S is a contradiction given common knowledge (i.e. $\mathcal{E}xh_{Alt(S)}S \Leftrightarrow_C \perp$), then sentence S sounds odd.

It is not difficult to see then that our proposal in §2 would predict BMG if only structures without K (or with $\mathcal{E}xh$ applying below K) were possible, and would predict LIG/MIIG if only structures with $\mathcal{E}xh$ applying above K (but not below K) were possible. Which one should it be then? Before we can answer this question (something we will do in §5), we should first examine the empirical status of these generalizations, something we haven't done so far in detail.

4 Test case: Disjunction in the scope of a universal quantifier

In many cases, the predictions of BMG and LIG/MIIG converge. Take for example (1), repeated here as (25):

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

Both BMG and LIG/MIIG predict infelicity here, simply because when S and A are contextually equivalent and A is an Innocently Excludable alternative of S , $K(S) \wedge \neg K(A)$ (which results from exhaustifying above K) is a contextual contradiction, just like $S \wedge \neg A$ (which results from exhaustifying below K).

In order to pit BMG and LIG/MIIG against each other, it would be very helpful to find cases where $\mathcal{E}xh_{Alt(S)}(S)$ yields a contradiction and $\mathcal{E}xh_{Alt(K(S))}(K(S))$ does not, and vice versa. I would like to show now that looking at sentences with disjunction in the scope of a universal quantifier is helpful for this purpose, because the meaning with $\mathcal{E}xh$ below K and the one with $\mathcal{E}xh$ above K are incompatible with each other. As a result, it will be easy to check what happens in cases where one of them is a contextual contradiction and the other is not.

Sentences with disjunction in the scope of a universal quantifier such as (26) are known to give rise to so-called Distributive Inferences (DIs) like those in (27) (see Crnič et al. 2015; Denić 2021; Bar-Lev and Fox 2023, a.o.).

(26) Every one of the students is French or Spanish. $\forall x(Px \vee Qx)$

(27) **Distributive inferences** (DIs):

a. \leadsto Some student is French $\exists xPx$

b. \leadsto Some student is Spanish $\exists xQx$

While DIs are very robust inferences (see Marty et al. 2023), (26) can also be uttered when the speaker is ignorant about whether or not there are any French students and about whether or not there are any Spanish students, as in (28).¹⁰

(28) **Ignorance inferences** (IIs):

- | | | |
|----|--|---------------------|
| a. | \sim The speaker is not certain that some student is French | $\neg K\exists xPx$ |
| b. | \sim The speaker is not certain that some student is Spanish | $\neg K\exists xQx$ |

Note that DIs and IIs contradict each other: One cannot possibly infer from (26) both that the speaker is certain that some student is French (as in (27)) and that the speaker is not certain that some student is French (as in (28)) while maintaining speaker consistency.

How do the two readings arise? Let me sketch a simple picture, setting aside almost all the details of the derivation of DIs and IIs.¹¹ On this view, the two kinds of inferences follow from two different parses: the DIs in (27) are derived (like other scalar implicatures) from exhaustification below K , while the IIs in (28) are derived (like other ignorance inferences) from exhaustification above K .¹² The picture that emerges is then as follows: DIs are derived when $\mathcal{E}xh$ applies directly to the sentence in (26), as in (29a), and IIs are derived when $\mathcal{E}xh$ applies above K , as in (29b) .

- (29) a. $\mathcal{E}xh_{Alt(\forall x(Px \vee Qx))} \forall x(Px \vee Qx) \Rightarrow \exists xPx \wedge \exists xQx$
 b. $\mathcal{E}xh_{Alt(K(\forall x(Px \vee Qx)))} K(\forall x(Px \vee Qx)) \Rightarrow \neg K\exists xPx \wedge \neg K\exists xQx$

Disjunction in the scope of a universal quantifier then provides an interesting testing ground for teasing apart the prediction of BMG from those of LIG/MIIG, because DIs (the result of applying $\mathcal{E}xh$ below K) and IIs (the result of applying $\mathcal{E}xh$ above K) contradict each other. This property will allow us to find cases where DIs (derived with exhaustification below K) lead to a contextual contradiction but IIs (derived with exhaustification above K) do not; such cases are predicted to be infelicitous by BMG but not by LIG/MIIG. Conversely, this property will also allow us to find cases where IIs (derived with exhaustification above K) lead to a contextual contradiction but DIs (derived with exhaustification below K) do not; such cases are predicted to be infelicitous by LIG/MIIG but not by BMG. In the next two subsections I will argue that, at least prima facie, this strategy provides arguments in favor of both BMG and LIG/MIIG.

Before we do that, one note is in order. While as I said I will not go into the details of how DIs

¹⁰ To see that (26) can indeed be interpreted with IIs, it may be helpful to think of situations where the number of students is relatively small, for instance if there are just 3 students. In this case, it is arguably more natural to interpret (26) with IIs than with DIs (this claim has been made by Denić 2021).

¹¹ A detailed view along these lines can be found in Bar-Lev and Fox (2023). A similar view has been advocated in Denić (2021), but for her the ability to derive DIs depends on pruning all the existential alternatives, which makes the picture more complicated.

¹² For Bar-Lev and Fox (2023), the derivation of DIs requires applying $\mathcal{E}xh$ recursively below K . $\mathcal{E}xh$ here should then be taken to either stand for two applications of Bar-Lev and Fox's $\mathcal{E}xh$, or to one application of an exhaustivity operator which operates recursively until a fixed point is reached.

are derived, one assumption that will be important to the discussion concerns the question what the formal alternatives of sentences of the form $\forall x(Px \vee Qx)$ are. I will assume that these are the alternatives predicted by [Katzir's \(2007\)](#) theory of alternative generation, namely all the possible combinations of replacements of \forall with \exists and $Px \vee Qx$ with Px , with Qx , and with $Px \wedge Qx$. While this assumption may look almost trivial to the reader unfamiliar with the recent literature on distributive inferences, it has not always been taken for granted that alternatives where \forall is replaced with \exists should be generated (see [Fox 2007](#); [Crnič et al. 2015](#)). It has however been argued (by [Bar-Lev and Fox 2020](#) and even more explicitly by [Bar-Lev and Fox 2023](#)) that on top of being motivated on conceptual grounds, generating these alternatives is also empirically motivated. To the extent that my analysis below is on the right track, it can serve as yet another reason to think that alternatives where \forall is replaced with \exists should be generated.

4.1 An argument in favor of BMG ([Magri 2009a,b](#))

Let us first look at a case where DIs are contextually contradictory, in (30) (repeated from (6)), which is a slightly modified version of an example by [Magri](#).¹³

- (30) *Context: A competition lasted five days, Monday through Friday; both John and Bill know that the same person won on each of the five days. John wants to know more about this amazing person and thus asks Bill for more information; Bill knows that this person was either Mary or Sue but he doesn't know which one of them she is, so he provides the following information:*
- a. #On every day, Mary or Sue won.
 - b. Mary or Sue won on every day.

While (30b) is perfectly natural, (30a) is infelicitous. This infelicity disappears if we slightly change the context so that it's no longer common ground that the same person won on each day. Following [Magri \(2009a\)](#), we can explain the contrast if we assume that the universal quantifier must take scope above disjunction in (30a) but not in (30b), and that DIs are derived with this scope configuration: Given that the context entails that the same person won on each day, the DIs of (30a) according to which Mary won on some day and Sue won on some day are contextually contradictory.

- (31) **Interpretation of (30a) with DIs is a contextual contradiction:**
- a. On every day, Mary or Sue won,
 - b. on some day Mary won, and
 - c. on some day Sue won.

¹³ [Magri's](#) original example involved an indefinite rather than disjunction. [Magri](#) is aware that his theory does not predict the infelicity of (30) while BMG does; see [Magri \(2009a: pp. 37–38\)](#). The BMG is also crucial for his analysis of differences between individual level vs. stage level predicates, providing yet another argument for BMG.

Infelicity is then predicted by BMG when combined with a theory where $\mathcal{E}xh$ derives DIs. It is however not derived by NMG, because there is no alternative which is contextually equivalent to the sentence. Note furthermore that if deriving IIs with $\mathcal{E}xh$ above K were possible, the sentence shouldn't have been odd, because IIs in this case are not a contextual contradiction; the following putative interpretation of (30a) would be entirely consistent with the context:

(32) **Interpretation of (30a) with IIs is not a contextual contradiction:**

- a. The speaker is certain that on every day, Mary or Sue won,
- b. the speaker isn't certain that on some day Mary won, and
- c. the speaker isn't certain that on some day Sue won.

Using LIG's terms, no alternative is contextually entailed by the sentence, and as a result LIG/MIIG don't predict infelicity.¹⁴ (30a) is then a case of infelicity which is captured by BMG but not by LIG/MIIG.

4.2 An argument in favor of LIG/MIIG (Denić 2021)

While Anvari (2018a,b) presents several cases of infelicity in support for LIG over NMG, I would like to focus on a novel argument which seems to me to be simpler, based on an observation due to Denić (2021).¹⁵ Denić points out that the sentences in (33) are rather terrible:¹⁶

¹⁴ One may consider ways in which the set of alternatives could be enriched in order to make NMG and LIG/MIIG account for it. One possibility, entertained by Magri, is to assume that (30b) is an alternative to (30a); since (30b) is logically stronger but they are contextually equivalent, NMG and LIG/MIIG would expect infelicity (for NMG this will further require assuming that there are no existential alternatives in the set of alternatives, otherwise (30b) won't be innocently excludable). However, as Magri (2009a) points out, the assumption that (30b) is a formal alternative to (30a) does not accord well with standard theories of alternative generation such as Katzir (2007), and furthermore, he shows that this assumption has problematic empirical consequence.

An alternative idea could be to generate the set of alternatives as in Katzir (2007), and then close it under disjunction. This will result in having an alternative for (30a) which is logically equivalent to (30b), that is:

- (i) (On every day Mary won) \vee (on every day Sue won)

While I see no empirical reason why making this assumption would be problematic, it is entirely ad hoc. In particular, it would raise the question why sets of alternatives should be closed under disjunction but not under conjunction, given that non-closure under conjunction is needed if one wants to have an implicature theory of free choice inferences (Fox 2007).

¹⁵ Some of Anvari's arguments in favor of LIG over NMG come from cases of infelicity which involve the interaction between implicature computation and presuppositions, which I aimed to avoid as much as possible in this paper in order to keep things simple; though see fn. 9 for a way in which such cases can be accounted for by the theory I proposed. Another argument made in Anvari (2018a) concerns so-called Hurford disjunctions, an issue which I briefly touched on in §3 (see especially fn. 6).

¹⁶ Denić uses the contrast between (33) and (i) as evidence that the issue has to do with IIs: The only apparent difference between them is that in (i) the IIs are not contradictory.

- (i) a. Every one of these three girls is called Sima, Rina or Dina.
- b. Every one of these three students read Anna Karenina, Germinal or Harry Potter.

- (33) a. #Every one of these three girls is Sima, Rina or Dina.
 b. #Every one of these three authors wrote Anna Karenina, Germinal or Harry Potter.

LIG captures the infelicity of (33a), for instance, because of the contextual but non-logical entailment in (34): If every one of these girls is Sima, Rina, or Dina, then it must be the case that one of them is Sima, given that it is (contextually) impossible for more than one person to be Sima (and likewise for Rina and Dina). Similar reasoning applies in the case of (33b).

- (34) Every one of these three girls is Sima, Rina or Dina \Rightarrow_C At least one of these girls is Sima

Given that we have shown that LIG is predicted by MIIG, it should come as no surprise that, as *Denić* pointed out, if IIs are derived the result is a contextual contradiction. In other words, MIIG predicts infelicity here, because the IIs are contradictory: If the speaker is certain that every one of these girls is Sima, Rina, or Dina, then it must be the case that she is also certain that one of them is Sima.¹⁷

- (35) **Interpretation of (33a) with IIs is a contextual contradiction:**
- a. The speaker is certain that each of these girls is Sima, Rina, or Dina,
 - b. the speaker isn't certain that at least one of these girls is Sima,
 - c. the speaker isn't certain that at least one of these girls is Rina, and
 - d. the speaker isn't certain that at least one of these girls is Dina.

Note however that BMG does not expect infelicity here: If $\mathcal{E}xh$ could apply below K and derive DIs, we would get an entirely consistent interpretation:

- (36) **Interpretation of (33a) with DIs is not a contextual contradiction:**
- a. Every one of these three girls is Sima, Rina or Dina,
 - b. at least one of these girls is Sima,

¹⁷ A difficult issue for a view where IIs are responsible for the infelicity of the sentence in (33) (like the one in *Denić* 2021 and the one we will adopt in §5) concerns their embedability. As *Denić* (2021) points out, (33a) does not seem to improve too much when embedded under negation (though she claims that the judgments of infelicity in (i) are not as clear as in (33a)):

- (i) ??It is not the case that every one of these three girls is Sima, Rina or Dina.

However, explaining infelicity in (i) along the same lines as (33a) will require assuming that IIs are derived at an embedded position below negation, an assumption which does not seem very appealing. This problem seems very general and applies to any ignorance-based account of infelicity: See *Anvari* (2018b) who argues against an ignorance-based account of infelicity in the context of the interaction between presuppositions and implicature computation (mentioned in fns. 15 and 9), and see *Marty and Romoli* (2022) who argue against an ignorance-based account of the infelicity of Hurford disjunctions (mentioned in §3). While this may seem to undermine the ignorance-based view of the infelicity of (33a) which I will adopt (following *Denić*), in §6 I will argue that this view has some clear advantages over views that are not based on ignorance, such as LIG (especially the felicity of (47)). I will have to leave the question of the status and source of infelicity in (i) for future work.

- c. at least one of these girls is Rina, and
- d. at least one of these girls is Dina.

Completely parallel considerations apply to (33b); (33a) and (33b) are then cases of infelicity which are captured by LIG/MIIG but not by BMG.

5 Refining the proposal

Let us take stock. In the previous section we have seen that there are both cases of infelicity predicted by BMG but not by LIG/MIIG ((30a)), as well as cases of infelicity predicted by LIG/MIIG but not by BMG ((33)). In this section I will aim to make sense of these findings along the lines of the theory of BMG proposed in §2. As we will see, however, the theory that will emerge does not predict LIG/MIIG in full generality. In the next section I will argue that this is a feature of the theory rather than a bug.

Recall that we expect the BMG to hold if what we care about is what happens when $\mathcal{E}xh$ applies below K , and we expect the MIIG to hold if what we care about is what happens when $\mathcal{E}xh$ applies above K . Our evidence so far may seem to indicate that we should care about both: whenever one of these structures yields a contextual contradiction, the sentence is infelicitous. This however looks like a conceptually strange idea: When a sentence can have two different parses, we do not normally choose a contextually contradictory parse if it has one. A simple minded evidence for this comes from (37).

(37) I saw the woman with the hat.

This sentence has a contextually contradictory parse where it means that the instrument I was using in order to see the woman was the hat, yet the sentence does not sound odd: the contextually contradictory reading is obviously not the only possible reading, and is in fact not readily available. Why should the situation be any different when the ambiguity at stake has to do with the position of $\mathcal{E}xh$ relative to K ?

One may think that exhaustification is special in some sense, so that as long as one choice of $\mathcal{E}xh$ insertion leads to a contextual contradiction this must be the chosen parse. However, even with exhaustification we do not seem to always choose contradictory parses: *Mary was smiling and dancing* does not look like a contradiction, even though it should have a parse which is contradictory, for instance one where $\mathcal{E}xh$ applies to the first conjunct.¹⁸

A more appealing direction is to assume that there is no ambiguity to begin with: The only possible parse is one where $\mathcal{E}xh$ applies both below and above K , as expected if one adopts Magri's

¹⁸ It is in principle possible to distinguish this case from the cases we were concerned with above, since this parse will lead to a logical (rather than merely contextual) contradiction. As I will argue in §6, however, there are cases where one choice of $\mathcal{E}xh$ insertion leads to a contextual contradiction but the sentence is still felicitous.

(2011) assumption that $\mathcal{E}xh$ must apply at every scope site. In that case, the only parse we should consider for a sentence S is the one in (38) (at least as long as K applies; I will elaborate on this in §6):

$$(38) \quad \mathcal{E}xh_{Alt_1}(K(\mathcal{E}xh_{Alt_2}(S)))$$

An immediate objection to this idea may seem to come from the evidence we have seen for LIG/MIIG in (33): With (38), we apparently expect DIs to be derived at the level of the lower $\mathcal{E}xh$; and once that is achieved, the higher $\mathcal{E}xh$ will no longer derive IIs (because DIs and IIs logically contradict each other, and $\mathcal{E}xh$ maintains logical consistency). And if we derive DIs rather than IIs, we no longer have an explanation for the infelicity of (33).

However, if there were independent reasons why one cannot derive DIs for (33), we might have a path towards an explanation, if we make sure that whenever the lower $\mathcal{E}xh$ cannot derive distributive inferences, the higher $\mathcal{E}xh$ will (obligatorily) derive IIs if they are contradictory.¹⁹ Indeed, *Denić*

¹⁹ In what follows I will assume that this is indeed what happens when the alternatives responsible for the derivation of DIs are pruned from the domain of the lower $\mathcal{E}xh$. Making sure that it does is not entirely trivial though, and it requires some auxiliary assumptions. IIs will be derived if the higher $\mathcal{E}xh$ in (i) negates both propositions in (ii).

$$(i) \quad \mathcal{E}xh_{Alt_1}(K(\mathcal{E}xh_{Alt_2}(\forall x(Px \vee Qx))))$$

$$(ii) \quad \begin{array}{l} \text{a. } K(\exists x Px) \\ \text{b. } K(\exists x Qx) \end{array}$$

On standard assumptions though, the alternatives in Alt_1 will not be simply the ones in (ii), but rather those in (iii) (unless alternatives where $\mathcal{E}xh$ is deleted are generated, a possibility which will make things simpler here but which may however be problematic elsewhere, especially when recursive $\mathcal{E}xh$ is at stake as in *Bar-Lev and Fox's* derivation of DIs).

$$(iii) \quad \begin{array}{l} \text{a. } K(\mathcal{E}xh_{Alt_2}(\exists x Px)) \\ \text{b. } K(\mathcal{E}xh_{Alt_2}(\exists x Qx)) \end{array}$$

We should then make sure that once the alternatives responsible for deriving the distributive inferences are pruned from the domain of the lower $\mathcal{E}xh$ in (i), i.e., as long as they aren't in Alt_2 , negating the alternatives in (iii) will indeed entail the negation of (ii).

This derivation is not straightforward though: a sentence S is usually taken to be an alternative of itself, namely a member of $Alt(S)$; thus $\forall x(Px \vee Qx)$ will also be a member of Alt_2 (as long as pathological choices of pruning aren't considered). As a result, the alternatives in (iii) will entail the propositions in (iv) respectively, and, consequently, excluding these alternatives would give us nothing: Since they contradict $\forall x(Px \vee Qx)$, their exclusion will be vacuous.

$$(iv) \quad \begin{array}{l} \text{a. } K((\exists x Px) \wedge (\neg \forall x(Px \vee Qx))) \\ \text{b. } K((\exists x Qx) \wedge (\neg \forall x(Px \vee Qx))) \end{array}$$

There are at least two ways in which we can ensure that IIs are still derived (on top of the one mentioned above where $\mathcal{E}xh$ can be deleted). First, we can assume that a sentence S is not in fact a member of $Alt(S)$, in which case the problem disappears. Second, we can assume that the alternatives for each alternative are generated independently of the sentence they are taken to be alternatives of; that is, since $\exists x Px$ does not usually have the more complex $\forall x(Px \vee Qx)$ as one of its formal alternatives, it will not have it as an alternative even when it is considered as an alternative to that sentence. Instead of Alt_2 in (iii), then, we will have a subset of Alt_2 which only contains alternatives that are at most as complex as $\exists x Px$ and $\exists x Qx$. Both of these options will ensure that negating the alternatives in (iii) will yield an equivalent result to negating the alternatives in (ii), and will ensure that IIs are derived when the alternatives

(2021) has argued that we generally disprefer deriving DIs when the number of disjuncts is no smaller than the number of individuals in the domain of quantification, as is the case in (33). And Bar-Lev and Fox (2023) argued for an even stronger generalization (what they call the categorical Denić generalization, in (39)), according to which whenever this is the case DIs are impossible to derive. Of course, one would like to know what explains this generalization; for our purposes, however, we will simply assume that it holds.²⁰

(39) **The categorical Denić generalization (CDG):**

In sentences of the form $\forall x \in D(P_1x \vee \dots \vee P_nx)$, distributive inferences can only be derived if $n < |D|$.

If the parse always involves $\mathcal{E}xh$ above and below K as in (38), the only way to respect the generalization in (39) in the interpretation of (33) is to prune the alternatives responsible for the derivation of DIs from the domain of the lower $\mathcal{E}xh$, that is, they should not be members of Alt_2 .²¹ But once DIs are not derived by the lower $\mathcal{E}xh$, the higher $\mathcal{E}xh$ will derive IIs; and we will not be able to prune the alternatives and avoid the IIs, because (as we have seen) the IIs are contextually contradictory, and so $\mathcal{E}xh$ over the full set of alternatives leads to a contradiction and, consequently, the theory in §2 will block pruning.

A theory where $\mathcal{E}xh$ applies both above and below K is then able to account for all the data we have seen so far, when taken together with the theory of pruning from §2 and when the categorical Denić generalization is taken into account. This theory also maintains the prediction from §2 that the BMG should hold, because if the lower $\mathcal{E}xh$ derives a contradiction with the full set of alternatives nothing can be pruned from its domain, and the end result will be contradictory. It however does not entirely predict LIG/MIIG; in the next section I will consider two cases where the theory proposed here and LIG/MIIG make different predictions, and argue that they favor the theory I proposed. This theory will then turn out to be not only more explanatory than LIG (or the stipulative view entertained above where MIIG rules out a sentence even if it has a non-contradictory parse), but also empirically advantageous over LIG/MIIG.

responsible for the derivation of DIs are pruned from the domain of the lower $\mathcal{E}xh$ in (i).

²⁰ I refer the reader to Denić (2021) and Bar-Lev and Fox (2023) for independent evidence for (39), as well as for proposals as to why it should be true. Bar-Lev and Fox (2023) aim to account for it within a general perspective on cases of infelicity like (1), but their account does not predict BMG (specifically, it fails to explain the infelicity of (30a)), and as far as I can see it cannot explain both the case of infelicity that we used to motivate LIG/MIIG from §4.2 together with the exceptions to LIG/MIIG which we will discuss in §6. One may hope for a general perspective which will make sense of all these cases as well as the Categorical Denić Generalization; I will have to leave this task to future work.

²¹ Note that our theory of pruning from §2 does not prevent pruning here, since $\mathcal{E}xh$ over the full set of alternatives does not lead to a contradiction; recall that if we were able to derive DIs in this case the result would not be a contextual contradiction.

6 Predictions diverging from LIG/MIIG

6.1 Prediction #1: Felicity when DIs can be derived and IIs are contradictory

The theory proposed in the previous section predicts IIs to lead to infelicity as expected by LIG/MIIG only to the extent that the lower $\mathcal{E}xh$ in (38) cannot derive DIs for some reason, e.g., due to the categorical *Denić* generalization. If it can, then even if the IIs are contextually contradictory, the ability of the lower $\mathcal{E}xh$ to derive DIs (thus blocking a derivation of IIs by the higher $\mathcal{E}xh$) will rescue from infelicity.

(40) Prediction diverging from LIG/MIIG (#1):

If it is possible to derive scalar implicatures (e.g., DIs) with $\mathcal{E}xh$ below K so that it is impossible to derive contradictory IIs when another $\mathcal{E}xh$ applies above K , then S can be felicitous even if $\mathcal{E}xh_{Alt(K(S))}K(S) \Leftrightarrow_C \perp$.

In order to test this prediction we should look for cases where DIs are not ruled out by the categorical *Denić* generalization, but the IIs are contradictory. While LIG/MIIG predict infelicity in such a case, the theory presented in the previous section does not; deriving DIs will block the derivation of the contradictory IIs. With this in mind, consider (41).

- (41) a. Each person in this group of four people has a parent or a child in the group.
b. #Both people in this group (of two people) have a parent or a child in the group.

Note first that LIG wrongly expects both (41a) and (41b) to be infelicitous, and the number of individuals in the group should hardly matter. This is so because of the contextual (but non-logical) entailment in (42): If everyone in the group has a parent or a child in the group, there must be at least one parent-child pair in the group, and consequently there must be at least one person who has a child in the group and at least one person who has a parent in the group.²²

- (42) Each person in this group has a parent or a child in the group \Rightarrow_C Some person in the group has a parent in the group.

As a result, if IIs were derived, then both sentences in (41) should have been a contextual contradiction: A speaker cannot be certain that everyone in the group has a parent or a child in the group without also being certain that some person in the group has a parent in the group and likewise for

²² (i) is another example which can be used to make the same point (courtesy of Danny Fox, p.c.): If either sentence in (i) is true, there has to be one grade above the average grade and one grade below the average grade, so DIs are contextually entailed.

- (i) a. Every one of my grades is either above or bellow my average grade.
b. #Both of my grades are either above or bellow my average grade.

someone having a child in the group. So IIs in this case would yield a contextual contradiction:

(43) **Interpretation of (41a)/(41b) with IIs is a contextual contradiction:**

- a. The speaker is certain that each person in this group has a parent or a child in the group,
- b. the speaker isn't certain that some person in this group has a parent in the group, and
- c. the speaker isn't certain that some person in this group has a child in the group.

LIG/MIIG then wrongly predicts infelicity for both examples in (41). The theory proposed here however does not: Because nothing prevents the lower $\mathcal{E}xh$ from deriving DIs for (41a), it can derive them and block the derivation of IIs by the higher $\mathcal{E}xh$ operator, thus leading to a consistent meaning:

(44) **Interpretation of (41a)/(41b) with DIs is not a contextual contradiction:**

- a. Each person in this group has a parent or a child in the group,
- b. some person in this group has a parent in the group, and
- c. some person in this group has a child in the group.

This rescue operation by the lower $\mathcal{E}xh$ is not possible in the case of (41b), however, because the derivation of DIs is blocked by the categorical *Denić* generalization. The theory proposed in §5 which relies on applying $\mathcal{E}xh$ both above and below K then correctly predicts the contrast between (41a) and (41b) when the categorical *Denić* generalization is taken into account; LIG/MIIG, in contrast, wrongly predict (41a) to be infelicitous.

6.2 Prediction #2: Felicity when the Maxim of Quantity is inactive

So far, I assumed that K always applies, which, given that $\mathcal{E}xh$ applies at every scope site, means that ignorance inferences will always be derived (unless blocked by a lower exhaustivity operator, as we have just seen). There are however cases where ignorance inferences shouldn't be derived. As *Grice (1978)* pointed out, this happens when the speaker is not expected to convey all the information they have; in the context of a treasure hunt, the following sentence does not give rise to the ignorance inferences normally associated with disjunction:

(45) The prize is either in the garden or the attic (but I'm not telling you which).

The disappearance of ignorance inferences in treasure hunt cases has been attributed to the Maxim of Quantity being inactive: In such a situation, the speaker is not expected to deliver all the relevant information they have as the Maxim of Quantity normally dictates (see *Fox 2014*). How should we think about these examples in terms of a theory where ignorance inferences are derived with $\mathcal{E}xh$ applying above a K operator? I will assume that K has to apply whenever the Maxim of Quantity is

active (see Fox 2016; Buccola and Haida 2019 for a conceptual motivation behind this), but when it isn't, then *K* doesn't have to apply.²³ And in the absence of *K*, there is of course no room for *Exh* above *K*, and consequently no ignorance inferences are derived in the grammar.

Once we assume that *K* doesn't have to apply when the Maxim of Quantity is inactive, we make the following prediction: Whenever the Maxim of Quantity is inactive, IIs will not be derived even if deriving them leads to a contextual contradiction; in other words, one way to avoid infelicity due to contradictory IIs would be to have a situation where IIs are not expected to arise to begin with, that is, in the absence of *K*.

(46) **Prediction diverging from LIG/MIIG (#2):**

If the Maxim of Quantity is inactive, then a sentence can be felicitous even if IIs are otherwise contextually contradictory.

We can now look at what happens with examples like (33a) when the Maxim of Quantity is inactive:

- (47) *Context: There are three girls behind masks and you have to identify which one is which. I know exactly who's behind each mask and I say:*
(All I can tell you is that) each of these three girls is Sima, Rina or Dina

The sentence in (47) is much better in this context than when uttered out of the blue (as in (33a)), and suddenly looks entirely felicitous. This is surprising for LIG, where the issue of whether the Maxim of Quantity is active or not isn't expected to play any role, and as a result (47) should be odd for the same reasons that it's odd when uttered out the blue (as in (33a)). This is similarly surprising for MIIG, which does not take into account the possibility of having a parse without *K* at all. To sum up, in this section I discussed cases where LIG/MIIG wrongly predict infelicity while the theory I proposed in §5 correctly predicts no infelicity, lending it further support.²⁴

²³ Another option would be to say that when the Maxim of Quantity is inactive, *K* still applies but *Exh* doesn't have to apply above it. I avoid this direction simply because I cannot see how it can be made compatible with the assumption that *Exh* applies at every scope site.

²⁴ Another case where LIG has been argued to be too strong is discussed in Marty and Romoli (2022), based on felicitous Hurford disjunctions like (i). Their argument relies on the assumption that (i) has the alternative in (ii), which is contextually but not logically entailed by (i); under this assumption, LIG incorrectly predicts (i) to be infelicitous.

(i) John studied in Paris or somewhere else in France.

(ii) John studied in France.

While I do not aim to provide a theory of Hurford disjunctions in this paper (see fn. 6), I should note that the theory proposed here makes the same problematic prediction that LIG does: Applying *Exh* above *K* to (i) should entail the exclusion of *K*((ii)), yielding a contextual contradiction (this concern applies to Meyer's theory of Hurford disjunctions just as well, a point which seems to be overlooked in Marty and Romoli 2022). Of course, this would not be a problem if we assumed that (ii) is not an alternative of (i); however, as Marty and Romoli (2022: fn. 6) point out, this assumption does not seem to be independently motivated and is at odds with the theory of alternative generation in Katzir (2007).

	(1)	(30a)	(33)	(41a)	(47)
NMG (Magri’s account)	✓	✗	✗	✓	✓
BMG (Proposal in §2)	✓	✓	✗	✓	✓
LIG/MIIG	✓	✗	✓	✗	✗
Revised proposal in §5	✓	✓	✓	✓	✓

Table 1 A summary of how the predictions of the various generalizations and theories discussed in this paper fare with respect to the main data points.

7 Conclusion

In this paper I proposed an account of the Broad Magri Generalization (BMG; Magri’s “Mismatch Hypothesis”) based on a theory which only allows pruning when exhaustification leads to an irrelevant meaning otherwise, and on Lewis’s observation about the relevance of contradictions. I have further shown that Anvari’s Logical Integrity Generalization (LIG) can also be partly made sense of within this system in terms of mandatory ignorance inferences, while the theory that emerges correctly predicts some notable exceptions to LIG. Table 1 summarizes how the various generalizations and theories we discussed fare with respect to the main data points presented in this paper.

References

- Anvari, A. (2018a). Logical integrity. In *Semantics and Linguistic Theory*, Volume 28, pp. 711–726.
- Anvari, A. (2018b). Logical integrity: from maximize presupposition! to mismatching implicatures. Manuscript.
- Anvari, A. (2022). Atoms and oddness. Talk given at Rutgers University, April 1, 2022.
- Bar-Lev, M. E. (2018). *Free Choice, Homogeneity, and Innocent Inclusion*. Ph. D. thesis, Hebrew University of Jerusalem.
- Bar-Lev, M. E. (2021). An implicature account of homogeneity and non-maximality. *Linguistics and Philosophy* 44, 1045—1097.
- Bar-Lev, M. E. and D. Fox (2020). Free choice, simplification, and innocent inclusion. *Natural Language Semantics* 28, 175–223.
- Bar-Lev, M. E. and D. Fox (2023). On fatal competition and the nature of distributive inferences. Manuscript, TAU and MIT.
- Buccola, B. and A. Haida (2019). Obligatory irrelevance and the computation of ignorance inferences. *Journal of Semantics* 36(4), 583–616.
- Chierchia, G., D. Fox, and B. Spector (2012). Scalar implicature as a grammatical phenomenon. In C. Maienborn, K. von Stechow, and P. Portner (Eds.), *Semantics: An International Handbook of Natural Language Meaning*, Volume 3, pp. 2297–2331. Mouton de Gruyter.
- Crnič, L., E. Chemla, and D. Fox (2015, Dec). Scalar implicatures of embedded disjunction. *Natural Language Semantics* 23(4), 271–305.

- Denić, M. (2021). Probabilistic informativeness in implicature computation: the case of embedded disjunctions. Manuscript.
- Fox, D. (2007). Free choice and the theory of scalar implicatures. In U. Sauerland and P. Stateva (Eds.), *Presupposition and Implicature in Compositional Semantics*, pp. 71–120. London: Palgrave Macmillan UK.
- Fox, D. (2014). Cancelling the maxim of quantity: Another challenge for a Gricean theory of scalar implicatures. *Semantics and Pragmatics* 7(5), 1–20.
- Fox, D. (2016). On why ignorance might be part of literal meaning—commentary on Marie Christine Meyer. Talk given at the MIT workshop on exhaustivity.
- Fox, D. and R. Katzir (2011, Mar). On the characterization of alternatives. *Natural Language Semantics* 19(1), 87–107.
- Gajewski, J. and Y. Sharvit (2012). In defense of the grammatical approach to local implicatures. *Natural language semantics* 20(1), 31–57.
- Grice, H. P. (1978). Further notes on logic and conversation. In *Pragmatics*, pp. 113–127. Brill.
- Horn, L. R. (1972). *On the semantic properties of logical operators in English*. Ph. D. thesis, University of California, Los Angeles.
- Katzir, R. (2007). Structurally-defined alternatives. *Linguistics and Philosophy* 30(6), 669–690.
- Katzir, R. (2014). On the roles of markedness and contradiction in the use of alternatives. In S. P. Reda (Ed.), *Pragmatics, Semantics and the Case of Scalar Implicatures*, pp. 40–71. Springer.
- Katzir, R. and R. Singh (2015). Economy of structure and information: Oddness, questions, and answers. In E. Csipak and H. Zeijlstra (Eds.), *Proceedings of Sinn und Bedeutung*, Volume 19, pp. 302–319.
- Lewis, D. (1988). Relevant implication. *Theoria* 54(3), 161–174.
- Magri, G. (2009a). *A theory of individual-level predicates based on blind mandatory implicatures: constraint promotion for optimality theory*. Ph. D. thesis, Massachusetts Institute of Technology.
- Magri, G. (2009b). A theory of individual-level predicates based on blind mandatory scalar implicatures. *Natural language semantics* 17(3), 245–297.
- Magri, G. (2011). Another argument for embedded scalar implicatures based on oddness in downward entailing environments. *Semantics and Pragmatics* 4(6), 1–51.
- Marty, P. and J. Romoli (2021). Presuppositions, implicatures, and contextual equivalence. *Natural Language Semantics*.
- Marty, P. and J. Romoli (2022). Varieties of Hurford disjunctions. *Semantics and Pragmatics* 15, 3.
- Marty, P., J. Romoli, Y. Sudo, and R. Breheny (2023). What makes an inference robust? Manuscript.
- Marty, P. P. (2017). *Implicatures in the DP domain*. Ph. D. thesis, Massachusetts Institute of Technology.
- Meyer, M.-C. (2013). *Ignorance and grammar*. Ph. D. thesis, Massachusetts Institute of Technology.
- Meyer, M.-C. (2014). Deriving Hurford’s constraint. In *Semantics and Linguistic Theory*, Volume 24, pp. 577–596.
- Roberts, C. (1996). Information structure in discourse: Toward a unified theory of formal pragmatics. *Ohio State University Working Papers in Linguistics* 49, 91–136.
- Singh, R. (2010). Oddness and ignorance inferences. Ms., Handout presented at Modularity Reading Group at MIT.
- Spector, B. (2014). Scalar implicatures, blindness and common knowledge: Comments on Magri (2011). In

Pragmatics, Semantics and the Case of Scalar Implicatures, pp. 146–169. Springer.

Spector, B. and Y. Sudo (2017). Presupposed ignorance and exhaustification: how scalar implicatures and presuppositions interact. *Linguistics and Philosophy* 40(5), 473–517.