[Manuscript, 13 September 2023]

Is Hyperbole a Scalar Inference?

Diego Feinmann

Polish Academy of Sciences, Institute of Computer Science, Poland.

Abstract. This article is concerned with the interpretation of hyperbolic statements. In the first part, I contend that Kao, Wu, Bergen, and Goodman (2014)—to my knowledge, the first (and only) attempt to derive hyperbole interpretations using formal methods—is unsuccessful. In the second part, I put forth a conjecture—roughly, a sentence S can be interpreted hyperbolically as meaning S^-, where S^- is weaker than S, only if S^- is a scalar alternative of S. I substantiate this conjecture by a series of novel empirical observations.♥♣

1 Introduction

A hyperbole is a false statement—or an almost certainly false statement—that isn't meant literally and is used for effect (typically, to convey affect). For example, I may utter 'Chomsky wrote 10,000,000 books', despite being common ground that Chomsky didn't write 10,000,000 books, to communicate that he wrote a lot of books and that I am positively impressed by this. To the best of my knowledge, Kao, Wu, Bergen, and Goodman (2014) stands alone as the only attempt to derive hyperbole interpretations using formal methods. In this article, I argue that this attempt is unsuccessful. In addition, I advance a conjecture that, if proven true, would establish a fundamental connection between the interpretation of hyperbolic statements and the computation of scalar implicatures.

2 Kao, Wu, Bergen, and Goodman (2014)

Kao et al. (2014), which is couched within the Rational Speech Act (RSA) framework (Frank and Goodman 2012), has two distinctive features. First, it assumes that utterance interpretation

♥ This work benefited enormously from conversations with Benjamin Spector, Leon Bergen, Moshe E. Bar-Lev, Emmanuel Chemla, Manuel Križ, Yasu Sudo, Amir Anvari, Nathan Klinedinst, Paul Egré, Matthew Mandelkern, Daniel Rothschild, Salvador Mascarenhas, Philippe Schlenker, Louise McNally, Janek Guerrini, Tadeusz Ciecierski, Jeremy Kuhn, Dariusz Kalociński, Émile Enguehard, and Keny Chatain.

♣ Earlier versions of this paper were presented in the following workshops: Linguae Seminar (ENS/IJN Paris), 3 June 2021; The New York Philosophy of Language Workshop, 20 Sept 2021; S-Babble (UC San Diego), 5 October 2021; Semantics Research Seminar (University College London), 15 Oct 2021; LingLunch (CNRS/Paris 7), 4 Nov 2021. I am indebted to these audiences for valuable comments and criticisms.

operates along two dimensions (formalised as QUDs)—namely, the state-of-the-world dimension (e.g. *What's the weather like?*, *How much money does Chris owe?*, *Is John coming the party tomorrow?*) and the speaker-affect dimension (*Does the speaker have affect*?);[1] in addition, it assumes that the speaker chooses her utterances to maximise the probability of accomplishing her goals, which can be either communicating information along the state-of-the-world dimension, communicating information along the speaker-affect dimension, or communicating both.

Bergen (2016: 145) provides a model simulation—using as input sentence 'Bob owes me \$1,000,000'—that illustrates the workings of this account.[2] In Bergen's (2016) simulation, there are 3 possible money states (\$10, \$100, and \$1,000,000) corresponding to the amount of money that Bob owes the speaker and 2 possible affect states (⊤, which indicates negative affect, and ⊥, which indicates the absence of affect). There is a total of 6 worlds: each world consists of a money-affect pair ($M = x$, $A = y$), where x stands for the amount of money that Bob owes the speaker, while y stands for the speaker's affect. The speaker is assumed to know exactly how much money Bob owes her as well as what her affective state is. The listener's prior (which is also the common prior[3]) is assumed to satisfy the following constraints: First, $\mathrm{P}(A = \top \mid M = 10) < \mathrm{P}(A = \top \mid \mathrm{M} = 100) < \mathrm{P}(A = \top \mid M = 1{,}000{,}000)$ (i.e. the greater the amount of money that Bob owes the speaker, the likelier it is that the speaker is unhappy). Second, $\mathrm{P}(M = 1{,}000{,}000) < \mathrm{P}(M = 100) \leq \mathrm{P}(M = 10)$ (i.e. Bob is less likely to owe a million dollars than he is to owe smaller amounts of money). The values of the prior that Bergen uses in his simulation are given below.

x	$\mathrm{P}(M = x)$	$\mathrm{P}(A = \top \mid M = x)$
\$10	0.495	0.1
\$100	0.495	0.33
\$1000000	0.01	0.9

Table 1. The prior distribution. The column labelled $\mathrm{P}(M = x)$ provides the marginal probability of each possible money state, while the column labelled $\mathrm{P}(A = \top \mid M = x)$ provides the conditional probability of the speaker having affect given each possible money state.

[1] This distinction raises a conceptual issue (and issue which I will ignore here): in what sense is information relative to the speaker's affective state not information about the world?
[2] Kao et al. (2014) is also presented as an account of 'pragmatic halo' effects in natural language (Lasersohn 1999). In this article, my sole focus is on Kao et al. (2014) as an account of hyperbole.
[3] Like most game-theoretic models of communication, Kao et al. (2014) exploits the *common prior assumption*—for a an overview, see Morris (1995).

Bergen (2016) assumes that speaker is in world (100, ⊤), wants to communicate her affect (i.e. the speaker wants to address the QUD *Does the speaker have affect?*) and, to do so, can utter either '\$10', '\$100', or '\$1,000,000'. According to Table 1, '\$1,000,000' is the best that the speaker can do to accomplish this (because she knows that, upon hearing '\$1,000,000', the listener will infer that the probability of her having affect is 0.9). Let's now look at what happens on the listener's side. Upon hearing '\$1,000,000', the listener will consider different reasons why the speaker might have uttered such a sentence. Typically, the listener will conclude that, if the speaker uttered S, then she believes, and wants to communicate, that S is true. The core idea behind Kao et al.'s (2014) model is that, in the event of a hyperbolic statement, the listener will not draw this conclusion, the reason for this being two-fold: first, it is not very likely that the speaker believes that '\$1,000,000' is true (as shown in Table 1, it is common ground that '(Bob owes the speaker) \$1,000,000' is almost certainly false); second, there's an alternative, and more likely, explanation for why the speaker has chosen to utter '\$1,000,000'—namely, she wants to communicate her affect. Given these considerations, the listener will conclude that the speaker uttered S not because she thinks S is true but because she is affected (in this case, negatively).

This example illustrates the general idea that underlies Kao et al.'s (2014) account: To communicate her affect, a speaker can utter a declarative sentence S—even if she knows that S isn't true—provided that (i) S is very likely to be false according to the common prior and (ii) does a good job at communicating information along the speaker-affect dimension. In such conditions, upon hearing S, the listener will not update her prior with ⟦S⟧; instead, she will conclude that the speaker is affected (positively or negatively).

Typically, the listener will also extract information about the state of the world from hyperbolic statements; this, as Bergen (2016: 149) points out, 'is because of the listener's *a priori* knowledge that the state of the world and the speaker's affect are likely to be linked.' In the example discussed here, if the speaker has affect ⊤ (namely, negative affect), then it is a lot more likely that she is in a \$100-world than in a \$10-world[4]—that is, it's a lot more likely that

[4] For example, if $P(A = \top) = 0.5$, then $P(M = 100 \mid A = \top) = 0.32$ while $P(M = 10 \mid A = \top) = 0.099$. (These conditional probabilities can be calculated using Bayes' theorem).

she is in a world in which Bob owes her a significant amount of money than in a world in which Bob owes her little money.

2.1 A critical assessment of Kao et al.'s (2014) account

Kao et al.'s (2014) work represents a significant achievement in the field of pragmatics: they managed to put together a formally explicit model of a phenomenon that had historically resisted formal analysis. However, upon examining the model's predictions, at least three problems become apparent. In what follows, I discuss these problems.

2.1.1 Problem I

In plain English, Kao et al. (2014)—applied to 'Bob owes me $1,000,000'—could be summarised as follows:

> (i) The speaker utters Bob owes me $1,000,000. (ii) 'This can't be true!', the listener says to herself. 'The speaker is surely trying to communicate something else other than state-of-the-world information.' (iii) 'If the speaker was in a world in which Bob owes me $1,000,000 was true', the listener reasons, 'she'd be, almost certainly, unhappy.' (iv) 'Voilà', the listener concludes, 'the speaker is trying to tell me she's very unhappy (i.e. the speaker is trying to communicate her affect).' (v) 'And, surely, if the speaker is very unhappy, then Bob must owe her a significant amount.'

Note that (v) is a non-sequitur: it does not follow from the fact that the speaker is unhappy that Bob is likely to owe her a significant amount of money. Of course, if the only possible reason for the speaker being unhappy is Bob owing her money, then the listener will end up inferring—provided that the listener's prior satisfies the constraints imposed in Bergen's (2016) simulation—that Bob is likely to owe the speaker a significant amount of money. In real life, however, the listener won't be able to infer this from the fact that the speaker has negative affect: this is because, in real life, there is more than just one reason why someone may have negative affect (failing an exam, breaking an arm, etc.).

What this means is that, as soon as a realistic set of worlds is assumed, the actual interpretation that the reasoning in (i)-(v) will derive for 'Bob owes me $1,000,000'—or any other hyperbolic statement—is the following: ‘the speaker has (positive/negative) affect and something that justifies the speaker having (positive/negative) affect is the case.’ This is obviously too weak.

2.1.2 Problem II

Kao et al.’s (2014) account is global in nature: as far as the uttered sentence goes, it sees nothing else but its literal meaning (i.e. a proposition). On this account, therefore, the following contrast can’t be explained:

(1) a. This exercise is impossible to solve. (H ✓)
b. This exercise doesn’t have a solution. (H ✗)

(1)a and (1)b are (arguably) truth-conditionally equivalent; however, (1)a, but not (1)b, supports hyperbole.

Furthermore, if one adopts Kao et al.’s (2014) perspective, it’s not clear how to account for the fact that (2)a, but not (2)b, works as an exaggeration ((2)b is just a plain lie).

(2) a. Chomsky wrote thousands of books. (H ✓)
b. Chomsky won the Nobel Prize. (H ✗)

Indeed, according to Kao et al. (2014), it should be possible to interpret (2)b hyperbolically—like (2)a, it is extremely unlikely to be true and evokes high-affect worlds.

2.1.3 Problem III

Probabilistic conditioning is at the heart of Kao et al. (2014): the listener learns about the speaker’s affect state by conditioning the common prior on the proposition that the uttered sentence expresses. The following utterances, therefore, constitute a problem for this approach:

(3) a. He broke an unbreakable chair. (≈ *He broke a chair that is very hard to break*.)
b. He solved an unsolvable problem. (≈ *He solved a problem that is very hard to solve*.)

Indeed, (3)a-b, on any reasonable semantics, are necessary falsehoods (in probability talk, $P([\![(3)a]\!]) = P([\![(3)b]\!]) = 0$). Thus, if fed with (3)a-b, Kao et al.’s (2014) model will fail to derive the attested interpretations: it's just not possible to condition P on p in cases in which $P(p) = 0$ (in such cases, $P(\cdot|p)$ is undefined).

✾

Given the problems noted, I don't think that Kao et al.'s (2014) model can be regarded as a successful theory of hyperbole. In the next section, I'll advance a theoretical conjecture—one that links the phenomenon of hyperbole with that of scalar implicatures—that may lead to development of an alternative theory.

3 Is hyperbole a scalar inference?

Hyperbole can be viewed, at least in part, as a *weakening problem*,[5] e.g. how come that from 'Julia ate all of the cookies in the jar' we can infer, under the right epistemic conditions, *Julia ate most of the cookies in the jar* and not, for example, *Julia ate all of the chocolate cookies in the jar*? (both interpretations are weaker than the original sentence, yet only the former is attested). Tellingly, scalar implicatures (SIs) give rise to a structurally analogous problem (only that, in this case, it's a *strengthening problem*), e.g. how come that from 'Julia ate some of the cookies in the jar' we can infer *Julia ate some of the cookies in the jar, but it's not the case that she ate all of the cookies in the jar* and not, for example, *Julia ate some of the cookies in the jar, but it's not the case that she ate all of the chocolate cookies in the jar* (both interpretations are stronger than the original sentence, yet only the former is attested).

Most, if not all, attempts at solving the problem that SIs give rise to[6] have been guided by the following conjecture:

(4) A sentence S can be interpreted as meaning $S \wedge \neg S^+$, where S^+ is stronger (more informative) than S, only if S^+ is a scalar alternative of S.

From this perspective, the reason why 'Julia ate some of the cookies in the jar' cannot be interpreted as meaning 'Julia ate some of the cookies in the jar, but it's not the case that she ate all of the chocolate cookies in the jar' but can be interpreted as meaning 'Julia ate some of the cookies in the jar, but it's not the case that she ate all of the cookies in the jar' lies in the fact

[5] The other part of the problem involves explaining hyperbole's affective component.
[6] A particular way of framing this problem came to be known as the 'symmetry problem' (Kroch 1972; Fox 2007; Katzir 2007, among others).

that 'she ate all of the chocolate cookies in the jar', unlike 'she ate all of the cookies in the jar', doesn't qualify as a scalar alternative of the uttered sentence.

Here I'd like to propose an analogous conjecture concerning the intepretation of hyperbolic statements:

(5) A sentence S can be interpreted hyperbolically as meaning S^{-}, where S^{-} is weaker (less informative) than S, only if S^{-} (or a sentence contextually equivalent to S^{-})[7] is a scalar alternative of S.[8]

(5) enables us to answer the question raised at the outset of this section (how come that from 'Julia ate all of the cookies in the jar' we can infer, under the right epistemic conditions, *Julia ate most of the cookies in the jar* and not, for example, *Julia ate all of the chocolate cookies in the jar*?) in a straightforward manner: 'Julia ate most of the cookies in the jar' is a scalar alternative of 'Julia ate all of the cookies in the jar'; 'Julia ate all of the chocolate cookies in the jar', by contrast, isn't.

In addition, (5) makes two interesting predictions. Indeed, according to (5), if S can be interpreted hyperbolically as meaning S^{-}, then S must have a scalar implicature α that is weaker (less informative) than S and contextually equivalent to S^{-}. This leads us to expect the following: (i) it should be possible to paraphrase the hyperbolic interpretation that S gives rise to by using α (because α is contextually equivalent to S^{-}); (ii) if S is also a scalar alternative of α,[9] an utterance of α would be expected to have the scalar implicature ¬S (because α is weaker than S). These predictions appear to be born out (see Figure 1 below).[10]

[7] The parentheticals protect the conjecture from 'silly' counterexamples—for example, 'No student is interested in the life of Dante' can be intepreted hyperbolically as meaning 'Few students are interested in the life of the author of the Divine Comedy'. The latter sentence, under standard assumptions, isn't a scalar alternative of 'No student is interested in the life of Dante'; however, 'Few students are interested in the life of the author of the Divine Comedy' is contextually equivalent to 'Few students are interested in the life of Dante', which is a scalar alternative of the uttered sentence.

[8] I intend the expression *scalar alternative* to be understood in a theory-neutral way: an alternative that needs to be generated by some mechanism or other to derive so-called scalar implicatures.

[9] In a framework such as Horn's (1972), this is always the case, as the scalar-alternative relation is symmetric.

[10] Of course, to arrive at a definite conclusion, a much larger dataset would need to be tested. Figure 1 is only meant as a preliminary assessment of the proposed conjecture.

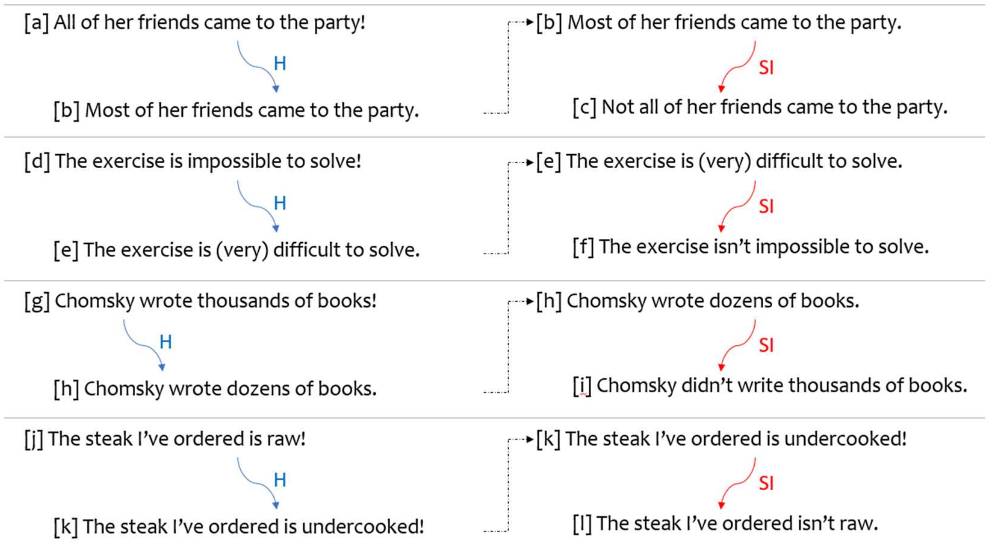

Figure 1. The 'H' arrow pairs a (false) statement S with α, a scalar alternative of S, whose meaning can be viewed as a possible hyperbolic interpretation of S. The 'SI' arrow pairs α with ¬S.

More generally, if (5) holds, it's reasonable to expect parallel behaviour between hyperbole and scalar implication, at least within certain paradigms. In what follows, I show that important parallelisms between these two phenomena can indeed be found.

3.1 Parallelism I (Numerals)

On neo-Gricean approaches to number interpretation, 'Three boys *blah*' is (typically) analysed as having weak truth-conditions (*Three or more boys blah*) and the upper-bounded component of the meaning (*No more than three boys blah*) as a scalar implicature (e.g. Horn 1972, Schulz and van Rooij 2006). From this perspective, the difference between (6)a and (6)b is that the former implicates that Chomsky didn't write more than 10,000,000 books, whereas the latter entails it. (In what follows, I'll assume that this analysis is correct.)

(6) a. Chomsky wrote 10,000,000 books.
$\rightsquigarrow_{SI}$ Chomsky didn't write more than10,000,000 books.
b. Chomsky wrote exactly 10,000,000 books.
$\not\rightsquigarrow_{SI}$ Chomsky didn't write more than 10,000,000 books.
$\Rightarrow$ Chomsky didn't write more than 10,000,000 books.

Now, consider the sentences in (7) ('*n*' stands for a high yet plausible number in the context of book writing, e.g., 40 books): whereas (7)a supports hyperbole, (7)b doesn't.

(7) a. Chomsky wrote 10,000,000 books.
$\rightsquigarrow_{H}$ Chomsky wrote *n* books or more.
b. Chomsky wrote exactly 10,000,000 books.
$\not\rightsquigarrow_{H}$ Chomsky wrote *n* books or more.

(6) and (7) can therefore be said to exhibit parallel patterns: ‘Chomsky wrote 10,000,000 books’, in its non-hyperbolic interpretation, implicates that Chomsky didn’t write more than 10,000,000 books; ‘Chomsky wrote exactly 10,000,000 books’, by contrast, doesn’t—indeed, this sentence entails, and hence doesn’t implicate, that Chomsky didn’t write more than 10,000,000 books. Similarly, ‘Chomsky wrote 10,000,000 books’ can be interpreted hyperbolically as meaning ‘Chomsky wrote *n* books [or more]’, where *n* is a high yet plausible number of books; ‘Chomsky wrote exactly 10,000,000 books’, by contrast, can’t be interpreted in such a way. If one accepts the proposed conjecture and, furthermore, believes that the scalar alternatives of ‘Chomsky wrote (exactly) 10,000,000 books’ can be derived only by replacing ‘10,000,000’ with another bare numeral, then the noted parallelism is unsurprising: ‘Chomsky wrote exactly 10,000,000 books’, unlike ‘Chomsky wrote 10,000,000 books’, has neither stronger nor weaker alternatives; hence ‘Chomsky wrote exactly 10,000,000 books’, unlike ‘Chomsky wrote 10,000,000 books’, is neither expected to be strengthened via a SI nor is it expected to support hyperbole.

‘Chomsky wrote *n* books [or more]’, it should be noted, isn’t the most natural paraphrase of (7)a’s hyperbolic intepretation—as opposed to, say, ‘Chomsky wrote a lot books’, which feels more appropriate. Under the proposed conjecture, this is unexpected: ‘Chomsky wrote *n* books [or more]’ is a scalar alternative of (7)a; ‘Chomsky wrote a lot books’, by contrast, is not a scalar alternative of (7)a (at least not under standard assumptions), nor is it contextually equivalent to a scalar alternative of (7)a.[11]

Does this observation threaten the proposed conjecture? I don’t think it does, at least not in a fundamental way. If one turns the proposed conjecture into a procedural account of how the listener recovers the message that the speaker intends to communicate by means of a hyperbolic statement, one obtains something like the following: interpreting a sentence S hyperbolically amounts to identifying S’s semantic content not with $[\![S]\!]$ but with $[\![S^{-}]\!]$, where S^{-} stands for a

[11] Indeed, unless one picks a highly contrived context, ‘a lot of books’, unlike ‘*n* or more’, would be expected to have a vague lower-bound.

scalar alternative of S that is weaker than S (and plausibly true given common knowledge). Upon hearing 'Chomsky wrote 10,000,000 books', however, interpreters are bound to have *a number* of possible landing sites for the hyperbole inference, each of which corresponding to a scalar alternative of the uttered sentence (e.g. 'Chomsky wrote 40 books [or more]', 'Chomsky wrote 41 books [or more]', 'Chomsky wrote 42 books [or more]', etc.). But which one to pick? In a realistic model of communication, the hyperbolic interpretation of a sentence like (7)a would need to be identified not with a scalar-alterative meaning (a proposition) but with a *distribution* over possible scalar-alternative meanings. From this perspective, it does make sense that, to describe the observed interpretation of (7)a, one feels compelled to use a *vague* sentence (e.g. 'Chomsky wrote a lot books', 'Chomsky wrote a vast amount of books'): these sentences manage to convey what an individual scalar alternative wouldn't be able to: the fact that the listener, after interpreting (7)a hyperbolically, can't tell with certainty what the lowest number of books that Chomsky definitely wrote is.[12]

3.2 Parallelism II (Hurford Disjunctions)

According to Hurford's (1974) well-known generalisation, often referred to as Hurford's constraint (HC), a sentence that contains a disjunctive phrase is odd if one of the disjuncts entails the other.[13] This generalisation correctly predicts (8)a-b to be odd; however, it fails to make sense of the fact that (8)c is oddness-free ('all of the books' entails 'some of the books').

(8) a. # John bought a computer or a laptop.
    b. # John either killed or murdered Paul.
    c. Jane read either some of the books or all of the books.

Chierchia, Fox, and Spector (2008) argued that (8)c isn't a counterexample to HC: according to them, (8)c and other apparent violations of this generalisation involve the presence of an implicature-computing operator *exh* within the first disjunct that 'breaks' the entailment between the disjuncts—as a result of this implicature, (8)b ends up having the reading 'Jane read either some *but not all* of the books or all of the books'. Now, one may ask: why isn't (8)a and (8)b also rescued by *exh*? This question is hardly ever raised; however, the assumption that most linguists appear to make is that, whereas certain words (e.g. 'some' and 'all') are lexical

---

[12] For the connection between vague sentences and the communication of probabilistic information, see Égré et al. (2023).
[13] I note that Hurford's generalization is, or at least appears to be, too weak (e.g. Singh 2008).

alternatives of each other (or, in Horn's (1972) terminology, form a *lexical scale*), others (e.g. 'computer' and 'laptop', 'kill' and 'murder') aren't.[14]

Let's assume that the account of the contrast between (8)a-b and (8)c just outlined is correct: then, if (5), the proposed conjecture, is also correct, the following prediction can be made: for any 'bad' Hurford disjunction *H*, it is never possible to use *H*'s strong disjunct to hyperbolically convey that *H*'s weak disjunct is the case. This prediction, as illustrated in Figure 2, is borne out.

Bad Hurford disjunctions	**Hyperbole unavailable**
[a] # Has John eaten a piece of fruit or an apple?	[b] John has eaten an apple. $\nrightarrow_{H}$ [c] John has eaten a piece of fruit.
[d] # John bought a computer or a laptop.	[e] John bought a laptop. $\nrightarrow_{H}$ [f] John bought a computer.
[g] # John either killed or murdered Paul.	[h] John murdered Paul. $\nrightarrow_{H}$ [i] John killed Paul.

Figure 2. Testing the following prediction: for any 'bad' Hurford disjunction *H*, it is never possible to use *H*'s strong disjunct to hyperbolically convey that *H*'s weak disjunct is the case.

It should be noted that, if one takes 'good' Hurford disjunctions as reference point, the data aren't entirely parallel: though it's possible to find good Hurford disjunctions whose strong disjunct can be used to hyperbolically to convey the weak disjunct, it's also possible to find good Hurford disjunctions that don't permit this, as illustrated below:

[14] Recall Gazdar (1979:58) famous statement: "scales are, in some sense, 'given to us'".

Good Hurford disjunctions	Hyperbole available
[a] ✓ Is this exercise (very) difficult or impossible?	[b] This exercise is impossible. $\searrow_H$ [c] This exercise is (very) difficult.
[d] ✓ Don't eat beef that is (very) undercook or raw.	[e] The beef is raw. $\searrow_H$ [f] The beef is (very) undercooked.
[g] ✓ John read most or all the books.	[h] John read all the books. $\searrow_H$ [i] John read most of the books.
	Hyperbole unavailable
[j] ✓ John read 1 or 2 books.	[k] John read 2 books. $\not\searrow_H$ [l] John read 1 book.
[m] ✓ John read the first two chapters or all the chapters.	[n] John read all the chapters. $\not\searrow_H$ [o] John read the first two chapters.

Figure 3. 'Good' Hurford disjunctions whose stronger disjunct gives rise to hyperbole ([a], [d], [g]) and 'good' Hurford disjunctions whose stronger disjunct does not give rise to hyperbole ([j], [m]).

The fact that this is so, it should be noted, doesn't threaten (5), according to which a sentence S can be interpreted hyperbolically as meaning S^{-} *only if* (and not *if and only if*) S^{-}—or a sentence contextually equivalent to S^{-}—is a scalar alternative of S. The kind of example that would threat (5) is a 'bad' Hurford disjunction such that its strong disjunct can be used hyperbolically to convey that its weak disjunct is the case—such cases, as far as I can tell, do not exist (see Figure 2).

3.3 Parallelism III (Indirect Implicatures)

As discussed in §2.1.2., (1)a—repeated as (9)a below—can be used/interpreted hyperbolically; (1)b—repeated as (9)b below—, by contrast, cannot.

(9) a. This exercise is impossible to solve.
$\rightsquigarrow_H$ This exercise is (very) difficult to solve.
b. This exercise doesn't have a solution.
$\not\rightsquigarrow_H$ This exercise is (very) difficult to solve.

Now, consider (10)a-b—the negations of (9)a-b, respectively:

(10) a. This exercise isn't impossible to solve.
$\rightsquigarrow_{SI}$ This exercise is (very) difficult to solve.
b. This exercise has a solution.
$\not\rightsquigarrow_{SI}$ This exercise is (very) difficult to solve.

(10)a carries the (indirect) implicature that the exercise is difficult (or very difficult) to solve; (10)b, by contrast, doesn't. The observed parallelism, if viewed through the lens of the proposed conjecture, is not at all surprising: (10)a, and not (10)b, is interpreted as meaning 'The exercise isn't impossible but it's (very) difficult to solve' because (10)a, unlike (10)b, has the alternative 'The exercise isn't (very) difficult to solve'; likewise, (9)a, and not (9)b, is interpreted hyperbolically as meaning 'The exercise is (very) difficult to solve' because (9)a, unlike (9)b, has the alternative 'The exercise is (very) difficult to solve'.

At this point, a careful reader may feel compelled to make the following point:

> **Careful reader's point**: Arguably, the most accurate paraphrase of the hyperbolic intepretation of (10)a is 'The exercise is *very* difficult to solve' (or 'The exercise is *extremely* difficult to solve') and not the plain 'The exercise is difficult to solve'. The author writes as if 'The exercise is *very* difficult to solve' was (or could be) a scalar alternative of 'The exercise is impossible to solve'; this assumption, however, is non-kosher: under standard assumptions, alternatives cannot be generated by replacing a lexical item (e.g. 'impossible') by a non-lexical constituent (e.g. 'very difficult')—cf. Horn (1972) and Katzir (2007). One can of course reject this principle; but, if one does, then something needs to be said about the symmetry problem: indeed, if alternatives can be generated by replacing lexical items with non-lexical constituents, then how come that a sentence of the form '<u>some</u> *blah*' appears to have '<u>all</u> *blah*', but not '<u>some but not all</u> *blah*', as an alternative?

One can approach this challenge in at least two ways. One option is to bite the bullet: if the proposed conjecture is correct, (9)a can be interpreted hyperbolically as meaning 'The exercise is difficult to solve' (whatever this means in the context in which (9)a is uttered[15]) but not as 'The exercise is very difficult to solve'. How to account, then, for the fact that an utterance of (9)a, at least intuitively, conveys not just that the exercise is difficult but that it is very (or extremely) difficult? This, it seems to me, is a possible answer:

(i) A hyperbolic utterance does give rise to the inference that the speaker is affected by the situation that she describes—after hearing (9)a, therefore, the listener will conclude that the speaker is (negatively) affected by the exercise's degree of difficulty (how to derive this inference is an open issue; see § 3.4).

[15] As is known, relative adjectives like 'difficult', 'tall' or 'rich' express properties only relative to contextually-provided thresholds or comparison classes.

(ii) If the speaker is negatively affected by the exercise's degree of difficulty, she most likely wasn't able to solve it (or she did solve it but struggled in the process).

(iii) Thus, it's more likely than not that the speaker deems the exercise to be very difficult as opposed to difficult but not too difficult.

Another option would be to reject the idea that 'The exercise is very difficult to solve' can't be a scalar alternative of (9)a. Now, if one does this, something needs to be said—as the careful reader points out—about the so-called 'symmetry problem'. And something can be said, I think.

In recent years, the case has been made that non-monotonic expressions are, in a sense, more complex than monotonic ones: on the experimental front, it has been shown that non-monotonic quantifiers are harder to learn than monotonic ones (e.g. Steinert-Threlkeld and Szymanik 2017; Chemla et al. 2019); on the theoretical front, ways of measuring semantic complexity have been identified under which nonmonotonic quantifiers come out more complex than monotonic ones (e.g. Carcassi, Schouwstra, and Kirby 2019; van de Pol et al. 2023).

In the light of these findings, it seems reasonable to consider the following solution to the symmetry problem: '<u>some</u> *blah*' has '<u>all</u> *blah*' as an alternative because 'some' and 'all' are of the same *semantic* complexity (in the sense that they are both monotonic); likewise, '<u>some</u> *blah*' doesn't have '<u>some but not all</u> *blah*' as an alternative because 'some' and 'some but not all' are not of the same complexity (the former is monotonic while the latter is non-monotonic).[16] Notably, if one adopts this solution, 'The exercise is very difficult to solve' can be (safely) characterised as a scalar alternative of (9)a: indeed, the expressions 'difficult' and

[16] This, I believe, is an attractive idea, an idea which Buccola, Križ, and Chemla (2022) appear to hint at. There's a growing consensus that 'breaking symmetry' via expression complexity (i.e. 'some but not all', unlike 'all', has more words than 'some') faces non-trivial empirical challenges—see, for example, Breheny et al. (2018) and Buccola, Križ, and Chemla (2022).

'very difficult', despite differing in word count, are both monotonic,[17] and can thus be considered equal in (semantic) complexity.[18]

3.4 Caveats and limitations

The first thing I'd like to note is this: the proposed conjecture does help to solve what I have termed 'hyperbole's weaking problem', but it does so only *partially*. Take, for example,

(11) Julia ate all of the cookies in the jar, and she watched her favourite movie.

(11) can be hyperbolically intepreted as meaning (i) 'Julia ate *most* of the cookies in the jar (, and she watched her favourite movie)' but not as meaning (ii) 'Julia ate all of the cookies in the jar *or* she had a cup of coffee'. Note that both (i) and (ii) are scalar alternatives of (11), and are weaker that (11); the question raised in §3 reemerges: why is it that (11) can be intepreted hyperbolically as meaning (i) but not as meaning (ii)? On this question, the proposed conjecture is mute.

It seems clear that not all alternatives can serve as landing site for the hyperbole inference (e.g. 'p and q' cannot be hyperbolically intepreted as meaning 'p or q'); in other words, being an alternative doesn't seem to be sufficient to qualify as a possible interpretation of a hyperbolic statement. This, as I already remarked in relation to the hybrid results obtained in Table 1, doesn't pose a threat to the proposed conjecture: (5) is compatible with the observation that not all scalar alternatives are hyperbole-friendly. It is left for further research to determine what other constraint(s)—in addition to the one articulated in (5)—is/are needed to provide a conclusive solution to hyperbole's weakening problem.

---

[17] If an object x falls in the extension of '(very) difficult', then any object y such that y is equally or more difficult than x, will also fall in the extension of '(very) difficult'. Because of this, '(very) difficult' can be said to denote a monotonically increasing function.

[18] Note that the account jut outlined immediately makes sense of the contrast between (i) and (ii):

(i) Is this exercise difficult, or is it very difficult?
(It reads as 'Is this exercise difficult but not very difficult, or is it very difficult?' This is consistent with the view that 'very difficult' is a scalar alternative of 'difficult'.)

(ii) # Is this exercise difficult, or is it difficult but not very difficult?
(It doesn't read as 'Is this exercise very difficult, or is it difficult but not very difficult?' This is the reading that one would expect if 'difficult but not very difficult' was a scalar alternative of 'difficult'.)

It's also worth noting that a theory of hyperbole should do more than just solving hyperbole's weakening problem: indeed, a theory of hyperbole should explain why a hyperbolic statement is weakened the way that it is, as well as why it manages to transmit information about the speaker's affective state. This is what Kao et al. (2014) tried to do—and they should be praised for trying to do that (ultimately, that is what is needed). Here I have attempted something considerably less ambitious than them: a (partial) solution to the hyperbole's weakening problem. Quite clearly, a lot more needs to be done.

Finally, I would like to draw attention to one possible counterexample to the proposed conjecture. Take, for example, the sentence in (12).

(12) John is a tower.

It seems uncontroversial that (12) can be intepreted as meaning that John is quite a lot taller than average. It's not clear to me, however, that (12) does have a scalar alternative that could be recruited to generate such an interpretation. Let's assume that as a matter of fact it doesn't: then, should we view (12), or similar examples, e.g., 'John is a giraffe', 'John is a pyramid', as counterexamples to the proposed conjecture?

I'm not persuaded that we should. Indeed, there are at least two coherent ways of thinking about (12) that are (or can be made) compatible with (5). One possibility is that 'John is a tower' is a metaphor, and hence a manifestation of a different empirical phenomenon. Another possibility is that 'John is a tower' is both a metaphor and a hyperbole: 'John is a tower', a category mistake if John is taken to be a person, is first intepreted metaphorically—to simplify things, assume that the metaphor-computing mechanism involves replacing 'a tower' by '(at least) 30m tall'; then, 'John is (at least) 30m tall', which is almost certain to be false, can be interpreted hyperbolically as meaning, for example, 'John is (at least) 2m tall', which, under standard assumptions, is a scalar alternative of 'John is (at least) 30m tall'. (See §3.1 for some thoughts as to why 'John is (very) tall' would typically be preferred over 'John is (at least) 2m tall' as a paraphrase of (12)'s hyperbolic intepretation).

Of course, these two ways of looking at (12) make sense provided that that the 'classical' picture, according to which hyperbole and metaphor are distinct phenomena, is on the right track. If this picture turns out not to be correct, then (12) could be regarded as a problem for the proposed conjecture.

4 Final thoughts

Is hyperbole a scalar inference? It depends, I guess, on one's working definition of 'scalar inference'. What seems likely is that scalar implicatures and hyperbolic statements share a common element—namely, reliance on scalar alternatives.

In this paper, I have tried to avoid as much as possible the issue of what it is that makes a sentence a scalar alternative of another sentence—suffice is to say that this is a hard problem that remains unsolved (for discussion, see Breheny et al. 2018). The fact that this problem hasn't yet been solved, it's worth noting, doesn't prevent us from agreeing on whether a given a sentence should count (or shouldn't count) as a scalar alternative of another—indeed, given 'All *blah*', it can be confidently said that 'Most *blah*' counts as one of its scalar alternatives because 'Most *blah*' has the scalar implicature 'Not all *blah*'—and, if the conjecture advanced here is correct, because 'All *blah*' can be interpreted hyperbolically as meaning 'Most *blah*'.

Moving forwards, the research question that I find most exciting is the following: assuming that the proposed conjecture is correct, why is it that that not all scalar alternatives can serve as landing site for the hyperbole inference—for example, why is it that that '*p* or *q*' has the implicature 'not(*p* or *q*)', but '*p* and *q*' can't be interpreted hyperbolically as meaning '*p* or *q*'? Answering this question in a principled way, I believe, may contribute not only to better understanding how hyperbole works, but also to shedding light on the elusive nature of alternatives.

References

Bergen, Leon. 2016. "Joint Inference in Pragmatic Reasoning." PhD Thesis, Massachusetts Institute of Technology.

Breheny, Richard, Nathan Klinedinst, Jacopo Romoli, and Yasutada Sudo. 2018. "The Symmetry Problem: Current Theories and Prospects." *Natural Language Semantics* 26 (2): 85–110. https://doi.org/10.1007/s11050-017-9141-z.

Buccola, Brian, Manuel Križ, and Emmanuel Chemla. 2022. "Conceptual Alternatives." *Linguistics and Philosophy* 45 (2): 265–91. https://doi.org/10.1007/s10988-021-09327-w.

Carcassi, Fausto, Marieke Schouwstra, and Simon Kirby. 2019. "The Evolution of Adjectival Monotonicity." In *Proceedings of Sinn Und Bedeutung*, 23:219–30.

Chierchia, Gennaro, Danny Fox, and Benjamin Spector. 2008. "The Grammatical View of Scalar Implicatures and the Relationship between Semantics and Pragmatics." In *Handbook of Semantics*, edited by Paul Portner, Claudia Maienborn, and Klaus von Heusinger. New York: Mouton de Gruyter.

Égré, Paul, Benjamin Spector, Adèle Mortier, and Steven Verheyen. 2023. "On the Optimality of Vagueness: 'Around', 'between' and the Gricean Maxims." *Linguistics and Philosophy* 46 (5): 1075–1130. https://doi.org/10.1007/s10988-022-09379-6.

Fox, Danny. 2007. "Free Choice and the Theory of Scalar Implicatures." In *Presupposition and Implicature in Compositional Semantics*, edited by Uli Sauerland and Penka Stateva, 71–120. Palgrave Studies in Pragmatics, Language and Cognition. London: Palgrave Macmillan UK. https://doi.org/10.1057/9780230210752_4.

Frank, Michael C., and Noah D. Goodman. 2012. "Predicting Pragmatic Reasoning in Language Games." *Science* 336 (6084): 998–998. https://doi.org/10.1126/science.1218633.

Goodman, Noah D., and Andreas Stuhlmüller. 2013. "Knowledge and Implicature: Modeling Language Understanding as Social Cognition." *Topics in Cognitive Science* 5 (1): 173–84. https://doi.org/10.1111/tops.12007.

Horn, Laurence. 1972. "The Semantics of Logical Operators in English." PhD Thesis, UCLA.

Kao, Justine T., Jean Y. Wu, Leon Bergen, and Noah D. Goodman. 2014. "Nonliteral Understanding of Number Words." *Proceedings of the National Academy of Sciences* 111 (33): 12002–7. https://doi.org/10.1073/pnas.1407479111.

Katzir, Roni. 2007. "Structurally-Defined Alternatives." *Linguistics and Philosophy* 30 (6): 669–90. https://doi.org/10.1007/s10988-008-9029-y.

Kroch, Anthony. 1972. "Lexical and Inferred Meanings for Some Time Adverbs." *Quarterly Progress Report of the Research Lab. of Electronics 104, MIT.*, 260–67.

Lasersohn, Peter. 1999. "Pragmatic Halos." *Language* 75 (3): 522–51. https://doi.org/10.2307/417059.

Morris, Stephen. 1995. "The Common Prior Assumption in Economic Theory." *Economics & Philosophy* 11 (2): 227–53. https://doi.org/10.1017/S0266267100003382.

Pol, Iris van de, Paul Lodder, Leendert van Maanen, Shane Steinert-Threlkeld, and Jakub Szymanik. 2023. "Quantifiers Satisfying Semantic Universals Have Shorter Minimal Description Length." *Cognition* 232 (March): 105150. https://doi.org/10.1016/j.cognition.2022.105150.

Schulz, Katrin, and Robert van Rooij. 2006. "Pragmatic Meaning and Non-Monotonic Reasoning: The Case of Exhaustive Interpretation." *Linguistics and Philosophy* 29 (2): 205–50.

Singh, Raj. 2008. "On the Interpretation of Disjunction: Asymmetric, Incremental, and Eager for Inconsistency." *Linguistics and Philosophy* 31 (2): 245–60. https://doi.org/10.1007/s10988-008-9038-x.