



## < ‘PAPER,’ PAPER, *PAPER* PAPER >; HORN SCALES AND LEXICAL EXPRESSIONS

---

HYUN PARK, *Yale University*<sup>1</sup>

The main goal of this paper is to examine Horn scales using the scales created with certain lexical expressions and follow the logical consequences of positive and negative scales for a relevant variable — that is, in the same way we understand [some] to mean [not all], we give lexical expressions an upper bound in terms of meaning. The lexical expressions we will focus on are lexical clones, scare quotes (which I will note using the quotation mark sign ‘\_’), negation, and the neutral form. The scales formed by these lexical expressions are the following: < ‘coffee,’ coffee, *coffee* coffee > for the positive scale and < ‘coffee,’ not coffee > for the negative scale. Though ‘coffee’ exists on both the positive and negative scales for coffeeeness, I will offer an explanation for why it, and other words that exist on both the positive and negative scales of a given feature, cannot be used to form a composite scale, such as # < not coffee, ‘coffee,’ coffee, *coffee* coffee > or # < cold, lukewarm, hot >. Finally, I will argue that saying ‘coffee’ when you really think of the referent as a different lexical expression of coffee could amount to a lie, depending on the lexical expression, thereby blurring even further the boundary between pragmatic and semantic meaning.

### I. HOW TRUE IS TRUE ENOUGH?

Children learn from a very young age that they should tell the truth. However, as we acquire more lexical entries and experience a more diverse range of social situations, the gray areas between truth and falsity grows larger, and the permissibility of statements that aren’t quite lies but aren’t quite true either becomes obfuscated. This paper aims to explore a certain kind of not-quite lies that can be triggered by Horn scales, a set of lexical entries that measure a certain salient variable on a spectrum. I won’t be examining any one specific scale but a scale that can apply to any lexical item based on certain types of modifications, such as lexical cloning, scare quotes (noted using the single quote, ‘\_’), negation, and the neutral form. The scales formed by these lexical expressions will hereby be referred to by variants of coffee, and they are ordered in the following manner: < ‘coffee,’ coffee, *coffee* coffee > for the positive scale and ‘*coffee*,’ not *coffee* for the negative scale. I will examine the expression [‘coffee’] and other scalemates that I will refer to as double-crossing scalemates, and why they cannot be the bridge between positive and the negative

---

1. Hyun Park is a junior majoring in philosophy and linguistics at Yale University. Her academic interests include philosophy of language, metaphysics, and semantics/pragmatics.

scale, despite the fact they are found on both scales. Though not all incongruities of specificity in what is said and what is actually thought will amount to a lie, I will argue that certain instances would qualify as such.

## II. THE TAXONOMY OF MEANING: ENTAILMENT AND IMPLICATURE

To understand not-quite truths, we need to understand what distinguishes not-quite truths from truths in the first place: meaning. Broadly speaking, meaning can be expressed at two different “levels,” the levels corresponding each to a subfield of linguistics. These subfields are called semantics and pragmatics.

Meaning conveyed at the semantic level is referred to as entailment. Entailment deals with what is literally said. We derive truth conditions, or the way the world must be for a sentence to be true, from what is “literally” said. From here, we judge the sentence as true or false by comparing these truth conditions to the actual state of our world. In semantics, truth and falsity is a mutually exclusive and exhaustive binary. For example, the sentence “The Eiffel Tower is in Incheon,” is true if and only if the Eiffel Tower is located somewhere in Incheon. If the Eiffel Tower is located somewhere that is not Incheon (say, perhaps, Oslo), this sentence is false. This is the relation between truth and meaning that is measured in logic, and it is not affected by changes in context. That is, neither the person who utters it nor the time at which it is uttered changes whether “The Eiffel Tower is in Incheon” is true or false, assuming that the Eiffel Tower’s location is fixed and that this sentence is uttered at points in time at which the Eiffel Tower exists.

Implicature corresponds to meaning expressed at the pragmatic level. It refers to meaning that is communicated intentionally, but not literally. There are two types of implicatures: conversational, those that can be canceled, and conventional, those that cannot. In this paper, I will focus on conversational implicature. For example, if someone asked me if I would recommend a class I have taken and I answered with “Well, I liked the lighting of the room it was held in,” it implicates that I would not recommend it. According to Paul Grice, conversational implicatures are created when a conversational maxim is violated, of which there are four: quantity, quality, relation, and manner.<sup>2</sup> The earlier example is a violation of the maxim of relation, because I deliberately refrained from saying something that would be more salient to my recommendation of a course. What I said seems unrelated to what was asked of me if we look only at the literal meanings of words, but my intentions were clear. When a speaker violates a maxim clearly and obviously, the listener looks for the reason they have answered in such a way, and this is where conversational implicatures find their meaning.

A different way to describe the rules that guide our intuitions on implicatures are the Q and R principles. The Q principle states that you should say as much as you can, and the R principle states that you should not say any more than necessary.<sup>3</sup> If I

2. Paul Grice, “Logic and Conversation,” in *Studies in the Way of Words* (Harvard University Press, 1989), 27.

3. Laurence Horn, “Neo-Gricean Pragmatics: A Manichaeian Manifesto,” in *Pragmatics*, ed. Noel

told my friend that I beat a small animal in reference to playing a game of chess with my younger cousin, my friend would be confused. Though I did not technically tell a lie, I was not as specific as I was expected to be. I did not say as much as I could have based on the knowledge I had, and therefore made an utterance where misinterpretation was not only possible but probable. In a normal exchange, saying “small animal” creates a Q-based implicature — because I could have chosen to be more specific about the identity of the small animal, I implicated that the animal I was referring to was not human. In the same way, if I say “I broke a bone yesterday,” my interlocutor will assume I meant my own bone. Because people are usually referring to their own bone when they speak in this manner, they will assume that “a bone” refers to a bone that belongs to me, and not my brother. This is called an R-based implicature, where my utterance is assumed to be more specific than is explicitly stated.

### III. HORN SCALES

Q-based and R-based implicatures are dependent on the existence of an underlying Horn scale. The classic example of a Horn scale is < some, many, most, all >, in which every scalemate to the right entails whatever is meant by the scalemate or scalemates to the left — for example, if I have many cats, it must be the case that I also have some cats. However, I am more likely to say “I have many cats” than to say “I have some cats” if it is the case I have many cats because I want to say as much as I can, according to the Q principle. Q-based and R-based implicatures, therefore, change the strength and/or specificity of a certain word that exists on a certain scale.

One test to determine whether there is a scale present looks like this: “it’s not just  $x$ , it’s  $y$ .” Examples include: “it’s not just a finger, it’s a thumb”; “it’s not just some cats, it’s all cats”; and more. The word of note here is “just” — this means all instances of  $y$  must fall under the set of all  $x$ ’s but not vice versa. There is a salient variable that only increases in intensity. A given salient variable can have up to two scales associated with it: positive and negative. The positive scale for temperature, for example, looks something like this: < lukewarm, warm, hot >. The negative scale looks like this: < lukewarm, cool, cold >. Both these scales share a salient variable — temperature — but they measure it in two different directions. The relationship between warm and hot is an increase in magnitude, the same as the relationship between cool and cold. We see that the magnitude of the words used decreases from hot to warm to lukewarm, but then increases from lukewarm to cool to cold. The “absolute value,” so to speak, increases in either direction from lukewarm, in the same way the literal absolute value increases in integers, the further one gets from zero. An increase in magnitude in two directions is one way to test for the presence of a positive and negative scale, given that the first scalemate of both scales are the same word. I will elaborate more on this phenomenon later.

The salient variable that we examine here is coffeeness. All of the lexical expressions of coffee bring into focus the coffeeness of a particular object, coffeeness often being a comment on the quality of the coffee but only insofar as it references a

closeness to some Platonic ideal of coffee. Obviously, such a coffee would be of a high quality, but being of high quality does not seem to be the only thing that gives coffee its coffeeness, since a bottle of expensive wine is of high quality but fails to be coffee. I'll comment briefly on the variants [not coffee], ['coffee'] and [*coffee* coffee].

I will start with the negation, which may not sound as though it should exist on a scale for coffeeness at all. Everything that is not coffee is not coffee — this is a tautology. However, this does not mean that everything that is not coffee is referred to as [not coffee]. Think about how odd it would be for me to refer to New York City as [not coffee]. For the word [coffee] to be used when referring to an object at all requires some relevance to coffee in a situation. If a barista brought tea rather than coffee to my table, it would make perfect sense for me to say “this is not coffee,” because tea has in common with coffee the fact I could have plausibly asked for it at a cafe. In a similar way, if a barista brought me a particularly disappointing coffee (perhaps the beans were roasted for too long, or there is too much sugar in it), I may whisper to my friend the very same “this is not coffee.” To elaborate on the earlier tautology, everything that isn't coffee is not coffee, but not everything that is called [not coffee] is not coffee, *e.g.*, bad coffee which both is coffee but one can say of it, “this is not coffee.”

[Coffee coffee] is a type of lexical expression produced by lexical cloning, or repeating a particular word *W* to bring attention to it. According to Laurence Horn, these clones are “a psychologically or perceptually salient exemplar or central category member ... not an outlier or peripheral category member.”<sup>4</sup> When spoken aloud, the prosodic focus (the accented or emphasized unit, so to speak) is placed on the first instance of *W*—*WW*. Lexical cloning can have many different uses depending on the context, but we'll use them to refer to the prototypical, high quality instance of *W*.

['Coffee'] is a trickier case. In his 2003 paper “Scare Quotes and Their Relation to Other Semantic Issues,” Stefano Predelli refers to scare quotes as a type of attachment trigger, or “an expression responsible for the presence of an attachment,” where attachment is defined as follows: “an additional piece of information is to be associated with that sentence, determined (at least in part) on the basis of the contributions offered by the expressions they enclose.”<sup>5</sup> One example of an attachment that can be triggered by scare quotes is an acknowledgement that the enclosed word or phrase is not perfectly suitable to a given situation, despite the fact that it was utilized by the utterer.<sup>6</sup> Now, a word can fail to be perfectly suitable for many reasons, but for our purposes, it will fail to be suitable because the referent is lower in coffeeness than [coffee].<sup>7</sup>

4. Laurence Horn, “The Lexical Clone: Pragmatics, Prototypes, Productivity,” in *Exact Repetition in Grammar and Discourse*, ed. Rita Finkbeiner and Ulrike Freywald (De Gruyter, 2018), 242.

5. Stefano Predelli, “Scare Quotes and Their Relation to Other Semantic Issues,” *Linguistics and Philosophy* 26, no. 1 (2003): 7.

6. Predelli, 2.

7. In the process of editing this paper, someone pointed out to me that Predelli describes attachment triggers as a type of conventional implicature, not conversational implicature. However, words that trigger conventional implicatures can be part of scales that create conversational implicatures. One good example is < and, but >. [But] conventionally implicates that the items it conjuncts together are


Given this background, the lexical expressions of coffee easily pass the test for Horn scales. “It’s not just ‘coffee,’ it’s coffee”; “it’s not just coffee, it’s *coffee* coffee”; “it’s not just ‘coffee,’ it’s *coffee* coffee” are all grammatical sentences whose meaning is easily grasped by English speakers. “It’s not just ‘coffee,’ it’s not coffee” is the same.

The scale of lexical expressions can be evaluated subjectively. For example, coffee from Starbucks is often referred to as [not coffee] or [‘coffee’] when people think it compares negatively to other coffees they have had. However, others refer to it as [*coffee* coffee], when they believe it is a good instance of coffee.

#### IV. DOUBLE-CROSSING SCALEMATES

Let’s return to the temperature scales. One may ask why we cannot collapse the positive and negative temperature scales into one scale. Despite the fact that these scales have a scalemate in common, we cannot use the “it’s not just *x*, it’s *y*” construction in a way that would combine the two scales — neither #“it’s not just warm, it’s cold” nor #“it’s not just cold, it’s hot” make much sense. But why exactly can’t a scale such as # < cold, cool, lukewarm, warm, hot > or its reverse be formed?

The answer can be easily gleaned from a graphical representation of the successful temperature scales and the unsuccessful composite scale:<sup>8</sup>



From here, we can see why the scale in Figure 1 is successful while the scale in Figure 2 is not — the slope of the first graph is just positive, whereas the slope of the second is negative then positive. In other words, the intensity of the words goes down from cold to lukewarm and then back up to hot. A fundamental feature of

unexpected or contrast with each other in some way, since this meaning is not truth conditionally relevant nor cancellable. However, it can be placed on a scale with the scalemate [and]: “he didn’t just quit his job and make more money, he quit his job *but* made more money.” When placed in contrast with [but], [and] conversationally implicates that there is no contrast between the two conjuncted items. The situation isn’t quite parallel, since [but] itself doesn’t conversationally implicate anything, but the conventional implicature does not preclude it from existing on a Horn scale, which are reliably associated with conversational implicature. In the same way, the meaning of [‘coffee’] having conventionally implicated meaning does not preclude the possibility of [‘coffee’] also having conversationally implicated meaning.

8. Somewhat ironically, none of the figures used in this paper will be to scale.

Horn scales is that they are listed from weakest to strongest (or at least, not strong-weak-strong) — if the composite of the temperature scales does not allow for a scale that goes from weakest to strongest, it cannot be referred to as a Horn scale. For an illustration of why #“it’s not just cold, it’s hot” and vice versa is unsuccessful, Figure 2 shows that the relevant property of the scale, intensity of temperature, does not change from cold to hot. This is another reason this fails to be a scale. On the scale  $< \text{some, many, most, all} >$ , all scalemates entail all scalemates to the left. If I had all of my mother’s cats in my room, it is necessarily true that I also have most, many, and some of my mother’s cats in my room. The  $\# < \text{cold, cool, lukewarm, warm, hot} >$  scale entails that everything that is hot is cold or that everything that is warm is cool. This is incoherent. When forming a Horn scale, the relevant variable must indeed be variable and have only a positive rate of change in order for the words to exist on the same scale. Because the composite temperature scale does not meet these qualifications, it cannot be a Horn scale.<sup>9</sup> Words like [lukewarm] are a special type of scalemate due to the fact they can be placed on both the positive and negative scales of the same variable, which, as we’ve seen, is not often possible. I will refer to scalemates of this kind that can exist on both positive and negative scales as double-crossing scalemates. To my knowledge, double-crossing scalemates have not yet been discussed in previous academic literature.

The scale of [coffee] works in a similar way to the temperature scales. [‘Coffee’] is the weakest scalemate in both the positive and negative scales, but cannot be the bridge between two scales, forming the scale  $\# < \text{not coffee, ‘coffee,’ coffee, coffee coffee} >$  or its reverse,  $\# < \text{coffee coffee, coffee, ‘coffee,’ not coffee} >$ , though #“it’s not just not coffee, it’s *coffee coffee*” sounds less egregious than previous examples.

It’s possible that there is more than one double-crossing scalemate for a given variable. In the case of lexical expressions, one example would be the double negation, or [not not coffee]. [Not not coffee] has no trouble passing the test given above, for both scales. “It’s not just not not coffee, it’s coffee,” “it’s not just not not coffee, it’s *coffee coffee*,” “it’s not just not not coffee, it’s not coffee” are all perfectly fine sentences, which is evidence that it can exist on both the positive and negative scales.<sup>10</sup>

However, an issue arises when [not not coffee] is compared with [‘coffee’] — not only do we not have an intuition for what #“it’s not just not not coffee, it’s ‘coffee’” could mean, the same is true for its reverse, #“it’s not just ‘coffee,’ it’s not not coffee.” This would mean that a scale containing both [not not coffee] and [‘coffee’] (such as  $\# < \text{not not coffee, ‘coffee,’ coffee, coffee coffee} >$  or  $\# < \text{not coffee, ‘coffee,’ not not coffee} >$ ) cannot exist. To generalize from this example, this seems to suggest that two double-crossing scalemates cannot occupy the same scale. If

9. However, if the relevant property of the scale was the presence of particle movement in the object, then a sentence like [it’s not just cold, it’s hot] would not at all be out of place. This is the goal of the Kelvin temperature scale — by constructing an “absolute zero” as the starting point from which you definitionally cannot have less particle movement, you prevent the possibility of a negative temperature scale. This means only a positive scale can exist, necessarily containing both hot and cold (since both are levels of particle movement), with the amount of particle movement increasing as you move up the scale.

10. For more discussion on the double negation, see Laurence Horn, “Lie-toe-tease: double negatives and unexcluded middles,” *Philosophical Studies* 174, no. 1 (2017): 79–102.

this is the case, then why? The answer is present in an old adage: you can't double cross a double crosser. The generalization is correct: it is not possible to place a double-crossing scalemate in the same scale as another double-crossing scalemate. There is not enough of a difference in meaning between ['coffee'] and [not not coffee] to justify such a comparison. In compliance with the R principle, we are not to say more than is necessary, and the distinction between ['coffee'] and [not not coffee] is never necessary, because the difference in their closeness (or distance) to the Platonic ideal of coffee is miniscule. The principle applies to other scales with double-crossing scalemates as well. Just as we do not say #“My dog is a golden retriever, and I have a dog” because the second clause adds nothing new to the common ground, we do not say #“it's not just lukewarm, it's tepid” or vice versa, because tepidness adds very little information, if any information at all, that hasn't already been expressed by lukewarmness.

## V. POLITENESS AND SCALAR SPECIFICITY

So far, we have only discussed cases where the lexical expressions of coffee have been consistent with the beliefs and attitudes of the speakers. But what happens when people speak in a manner that is not consistent with their beliefs and attitudes? How do we correctly categorize the misleading use of lexical expressions? In other words, if I say ['coffee'] when I really mean [*coffee* coffee], could I be lying? Well, it's not the case that using a scalemate that is not perfectly precise always results in a lie. One example is the politeness principle. Let's say that I attend a concert where my friend Kang-Joh is a performer. Inspired by the moment, Kang-Joh attempts a difficult technique in his keytar solo that goes conspicuously wrong. Feeling insecure, he asks me after the show whether anyone liked the keytar solo. I can respond in two ways:

- (A) Some people liked your keytar solo.
- (B) Some people disliked your keytar solo.

The difference between (A) and (B) lies only in the verb used. If Kang-Joh hears (A), he is likely to believe that not everyone liked the keytar solo. [Some] is taken to mean not [all] — [some] rules out the meaning of [all], because I could have just said [all] had I meant it. If Kang-Joh hears (B), he can't use the same reasoning as given above to believe that not everyone didn't like the keytar solo. [Some] does not rule out the meaning of [all]. Bonnefon et al. propose that “the inference from ‘some X-ed’ to ‘not all X-ed’ is made less available when X threatens ... a sense of positive identity and public self-esteem that all humans project and are motivated to support in social interactions.”<sup>11</sup> To apply it in this situation, Kang-Joh's inference from (B) to “Not all people disliked your keytar solo” is made less available because people disliking the keytar solo would threaten his public image. I could even say (B) with the knowledge that all people disliked his keytar solo, because knowing

11. Jean-François Bonnefon, Aidan Feeny, and Gaëlle Villejoubert, “When Some Is Actually All: Scalar Inferences in Face-Threatening Context,” *Cognition* 112, no. 2 (2009): 250.

everyone disliked the keytar solo is about as useful as knowing some people disliking the keytar solo — either way, you'd probably want to give up on becoming a keytar rockstar or practice significantly more.<sup>12</sup>

But what's stopping us from going even further and saying that everyone enjoyed the keytar solo? Wouldn't it be even more polite to say exactly what the soloist wants to hear? The answer is no — most people will refrain from saying “everyone liked your keytar solo” if they believe no one liked the keytar solo. Though it would be impolite to imply that no one liked the keytar solo, it would be even worse to lie outright.<sup>13</sup> As previously established, [some] is entailed by [all] — given a non-empty set, if [all] is true, [some] is necessarily true — so, “some people disliked your keytar solo” is not a lie in cases where “everyone disliked your keytar solo” is true.

Now, let's apply this logic to coffee. If my friend had made a cup of coffee that I judged to be low in coffeeness despite the fact that it was made with roasted coffee beans, I would probably mentally categorize it as [‘coffee’] and maybe even [not coffee]. However, it would be rude to tell them [it's ‘coffee’], even though [‘coffee’] does exist on the positive scale of coffeeness. This suggests that the unmarked version of ‘coffee’ is one that is on the negative scale of coffeeness, meaning that without additional context, [‘coffee’] is seen as referring to something that is bad at being coffee, rather than something that is less good at being coffee. Politeness allows us to use weaker scalemates when a stronger scalemate is more appropriate, and is an established method of breaking the Q principle that is likely recognized by your interlocutors.<sup>14</sup>

## VI. WHAT IS A LIE?

However, this still doesn't give us an account of lexical expressions and lies. We'll stick with coffee for now. Let's say that my mother hands me a cup of coffee that I judge to be absolutely phenomenal — it is the perfect specimen of [coffee coffee]. However, at this particular moment in time, I am angry with my mother. So, though I believe it to be false that this cup of coffee is [‘coffee’], when she asks me how I like the coffee, I tell her “I mean, it's ‘coffee,’” and I intend on deceiving her into thinking I disliked her coffee. For the sake of this example, let's say my mother is an award-winning barista and that her coffees are world-renowned for being

12. Bonnefon, Feeny, and Villejoubert, “When Some Is Actually All: Scalar Inferences in Face-Threatening Context,” 251.

13. The exact nature of what has gone wrong in cases of clear-cut falsehood is unclear. Grice originally classified it as a violation of the maxim of quality, making it the same type of “wrong” as failing to be relevant or failing to be concise. In Grice, “Logic and Conversation,” he himself admits that there's something potentially worse about violations of quality than violations of quantity, relation, and manner, and that they may need to be classified separately. In a similar way, falsehoods seem to be worse than other deviations from polite behavior.

14. Politeness is, of course, not the only established method of breaking the Q principle. Another method would be humility. Let's say I got a 94% on a midterm I took, and my friend asks me how I did. I probably wouldn't say “I did amazing — I basically got all the questions right.” Rather, I would say something like “I did okay, I did get some of the questions wrong though.” Politeness and humility seem to be first-person and second-person versions of the same phenomenon, where specificity is removed in favor of your interlocutor's feelings. But I have neither the time nor space to elaborate here.



consistently excellent. Did I tell a lie? To answer this question, we must look at what has traditionally constituted a lie. In “Telling it slant: Toward a taxonomy of deception,” Laurence Horn lists four conditions that have historically constituted the definition of lying:<sup>15</sup>


- (C1) *S* says/asserts that *p*
- (C2) *S* believes that *p* is false
- (C3) *p* is false
- (C4) *S* intends to deceive *H*

In contrast, misleading has not inspired any type of strict demarcation. It can be intentional or unintentional, it can be direct or indirect (though when you say “he misled me” you implicate that you don’t mean the more informative “he lied to me”), and it seems like whether something is categorized as misleading or not is more context-dependent than lying is.

In order for “this is ‘coffee’” to qualify as an undisputable lie, I must assert that this is [‘coffee’] as distinct from [coffee] or [*coffee* coffee], I must believe that it is false that the coffee is actually [‘coffee’], it must be false that this is [‘coffee’], and I must be intending to deceive my hearer. Given that I constructed a scenario in which (C2), (C3), and (C4) are all true, this boils down to the propositional content of “this is ‘coffee.’” As shown by the keytar examples, if there is an entailment relationship between “this is ‘coffee’” and “this is *coffee* coffee,” no lie is present, regardless of whether it is the proposition of the utterance that entails what is believed or the proposition of what is believed entailing what is uttered. Now, our initially nebulous question has become easier to answer: is there any entailment relationship between [‘coffee’] and [*coffee* coffee]? If yes, that is, if [‘coffee’] entails [*coffee* coffee] or vice versa, I merely mislead. If not, I lie.

My intuition reports that this is a lie, and I can explain why this is the case with the use of Horn scales. We established earlier from the politeness example that unmarked [‘coffee’] is associated with the negative scale of coffeeeness. [Coffee coffee], of course, is on the positive scale. Though it may seem insignificant, opposing scales are crucial to understanding why my utterance was a lie. “This is ‘coffee’” contains a proposition not entailed by “this is *coffee* coffee,” because using a word from the opposite scale of what you mean escapes the entailment relation.

One way to represent what I mean is a number line with scalemates *A*, *B*, and *C*, where *A*, *B*, and *C* are scalemates on a positive or negative scale that measure the same variable, with *A*, *B*, and *C* on the positive scale and *Y* and *Z* on the negative:



15. Laurence Horn, “Telling It Slant: Toward a Taxonomy of Deception,” in *The Pragmatic Turn in Law — Mouton Series in Pragmatics 18*, ed. Janet Giltrow and Dieter Stein (De Gruyter, 2017), 25.

There is a lot of difference in meaning between *A* and *C*, but there is an entailment relation from *C* to *A*, since they're both on the positive scale. Therefore, it would not be a lie to utter *A* when you actually believe *C* to be the case and vice versa, because *A* and *C* are both on the positive scale. However, saying *Z* when you actually believe *B* would result in different propositions and, therefore, would be a lie, even though it is a "shorter" distance than the distance between *A* and *C*.

This is reflected in our intuitions about temperature. For example, if I say "it's cool" in reference to a bucket of water that I actually believe is warm, this would be a lie. However, if I said "it's cool" when I think the bucket water is frigid, this would not be a lie, though the difference between cool (- 80°F) and frigid (- 32°F) is much more than the difference between cool and warm (- 90°F). I may have been intending to mislead in both, but only one of these situations qualifies as a lie. This shows that two propositions are not distinguished by the difference in meaning between two assertions but by the crossing from the positive scale to negative scale or vice versa.

The underlying Horn scales explain why this is a lie, but do not show us the mental calculus that goes into determining which proposition is contained in an assertion. This can be done using the Q and R principles. Let's return to the bucket of water. Let's say I tell my friend "it's lukewarm" when the water is actually cold. If I meant to use a word on the negative or positive scale of temperature, I would likely have been more specific than lukewarm — because I could have been more specific, there is a Q-based implicature that says I did not mean the stronger scalemates of [cold] or [hot]. This is how the speaker, who is lying, would want the hearer to interpret their word choice. In addition, the R-principle states that I ought not to speak more than necessary, and it seems ineffective and superfluous to mislead by saying something is "cold, just not very cold" or "hot, just not very hot," as would be the case if lukewarm were on the same scale as what I actually believed. This is how the speaker actually chooses their words in their attempt to mislead — not only is the speaker trying to mislead, but they are trying to mislead in a manner that leads the listener to think something significantly different from what is actually the case, to expand the (fabricated) common ground. With this context, it seems clear that we can categorize the proposition belonging to "it's 'coffee'" such that it is a lie when you believe "it's *coffee* coffee" to be the case.

If it is true that the utterance of ['coffee'] when one means [*coffee* coffee] is borrowing from the negative scale, making it a lie as I've argued, then the gesture of scare quotes must have semantic content, as it changes the truth conditions of a sentence.<sup>16</sup> It would not be a lie to say [coffee] when you mean [*coffee* coffee], but the addition of a gesture changes what is asserted from a word on the positive scale of coffeeness to a word on the negative scale of coffeeness. This shouldn't come as a major shock, since we already know of gestures that contain truth-conditional content: nodding and shaking one's head. Nodding and shaking one's head, of course, is somewhat different due to the fact these gestures are not attached to

---

16. Though he argues in a completely different manner, Predelli also concludes that scare quotes contain semantic information. For more, see Predelli, "Scare Quotes and Their Relation to Other Semantic Issues."

another lexical unit in the same way scare quotes are. The former roughly translate to a yes or no even without much context, whereas I'm not sure how you would translate scare quotes unless they were accompanied by a spoken word. In any case, it does blur even further the dichotomy between semantics and pragmatics, making the distinction between the two fields appear somewhat more arbitrary, useful though it may be. The situation becomes even more interesting as you realize that replacing [*coffee* coffee] with ['coffee'] is not a lie in every situation. The meaning conveyed with both "This isn't what I'd call 'coffee'" and "This isn't what I'd call *coffee* coffee" is that a certain instance of coffee does not live up to some established standard. There is not even an instance of misleading here, let alone a lie.

## VII. CONCLUSION

Using what I called double-crossing scalemates, I demonstrated that no negative and positive scale can be collapsed into one composite scale, even if they measure the same variable. I focused on the scales formed by lexical expressions (< 'coffee,' not coffee > and < 'coffee,' coffee, *coffee* coffee >) and used the variant ['coffee'] to show that there is semantic content associated with scare quotes, which is a gesture. Though I wouldn't commit to the claim that there is no such thing as the semantic-pragmatic distinction, this certainly broadens the gray area. I'd be interested to see if there are any other lexical expressions that can be placed on this scale, especially ones that are variants due to a gesture, or any other instances of truth-conditional content in gestures.

An interesting avenue to explore in regards to the double-crossing scalemates is the role of intention. I've argued that depending on the context given, a double-crossing scalemate can be interpreted as belonging to either the positive or negative scale for a given variable. The context-dependency of double-crossing scalemates are reminiscent of the context-dependency of indexical meaning. Perhaps an analysis of double-crossing scalemates as a type of indexical, with its character being its "location" between the positive and negative scales and its content being whether or not it is placed on the positive or negative scale at a given point in time would be worth pursuing.

## REFERENCES

- Bonnefon, Jean-François, Aidan Feeny, and Gaelle Villejoubert. "When Some Is Actually All: Scalar Inferences in Face-Threatening Context." *Cognition* 112, no. 2 (2009): 249–258.
- Grice, Paul. "Logic and Conversation." In *Studies in the Way of Words*, 23–40. Harvard University Press, 1989.
- Horn, Laurence. "Lie-toe-tease: double negatives and unexcluded middles." *Philosophical Studies* 174, no. 1 (2017): 79–102.
- . "Neo-Gricean Pragmatics: A Manichaeian Manifesto." In *Pragmatics*, edited by Noel Burton-Roberts, 158–183. Palgrave Macmillan, 2007.

- Horn, Laurence. "Telling It Slant: Toward a Taxonomy of Deception." In *The Pragmatic Turn in Law — Mouton Series in Pragmatics 18*, edited by Janet Giltrow and Dieter Stein, 23–55. De Gruyter, 2017.
- . "The Lexical Clone: Pragmatics, Prototypes, Productivity." In *Exact Repetition in Grammar and Discourse*, edited by Rita Finkbeiner and Ulrike Freywald, 233–264. De Gruyter, 2018.
- Predelli, Stefano. "Scare Quotes and Their Relation to Other Semantic Issues." *Linguistics and Philosophy* 26, no. 1 (2003): 1–28.