A surface-scope analysis of authoritative readings of modified numerals[1]

Brian BUCCOLA — *The Hebrew University of Jerusalem*
Andreas HAIDA — *The Hebrew University of Jerusalem*

Abstract. A sentence like *You're allowed to draw at most three cards* has a so-called authoritative reading characterized by two kinds of inference: an upper-bound inference (you're not allowed to draw more than three cards) and a free-choice inference (you're allowed to draw any number of cards in the range $[0,3]$). We show that authoritative readings are available for a variety of expressions beyond just *at most*, and we provide the first (as far as we know) surface-scope account of such readings, building on the recursive exhaustification account of free-choice disjunction proposed by Fox (2007).

Keywords: modified numerals, free-choice inferences, exhaustivity, recursive exhaustification

1. The puzzle

Sentence (1) has two readings, which Büring (2008) calls a *speaker insecurity reading* (or ignorance reading) and an *authoritative reading*. On its speaker insecurity reading, the sentence conveys that the maximum number of cards you're allowed to draw is either three or fewer than three, and the speaker is ignorant about which is true. This reading can be brought out by following the sentence with *but I don't know exactly how many*.

(1) You're allowed to draw at most three cards.

We will be concerned in this paper with the authoritative reading of (1), which is characterized by two kinds of inference: (i) an *upper-bound* (UB) inference, viz. that you're not allowed to draw four or more cards, and (ii) a *free-choice* (FC) inference, viz. that you're allowed to draw any number of cards in the range $[0,3]$.[2] Authoritative readings constitute a well-known puzzle for the semantics of modified numerals (Geurts and Nouwen, 2007; Büring, 2008; Penka, 2014), which we describe in two parts.

First, on standard assumptions about the meanings of *allow* and *at most three*, a surface-scope analysis of (1) predicts a weak literal meaning, 'There is a permissible world in which you draw three cards or fewer', notated henceforth as $\diamond[\leq 3]$. Neither an UB inference ($\neg\diamond[\geq 4]$) nor a FC inference ($\diamond[=3] \wedge \diamond[=2] \wedge \ldots$) logically follows from this interpretation. Moreover, an

1. For many helpful comments, questions, and discussions, we would like to thank Luis Alonso-Ovalle, Moysh Bar-Lev, Emmanuel Chemla, Luka Crnič, Danny Fox, Doris Penka, Bernhard Schwarz, Stephanie Solt, Benjamin Spector, and Tue Trinh, as well as the audiences at the Institut Jean Nicod (École normale supérieure, LANGUAGE seminar series, July 6, 2016), at *Sinn und Bedeutung* 21 (University of Edinburgh, September 4–6, 2016), at *Logic in Language and Conversation* (*LogiCon*) (Utrecht University, September 19–20, 2016), and at ZAS (*Semantikzirkel*, September 21, 2016). This research was supported by the Israel Science Foundation (ISF 1926/14) and by the German-Israeli Foundation for Scientific Research (GIF 2353).

2. Our empirical claim regarding FC inferences deviates from that of Penka (2014), who assumes only partial, rather than total, FC: (1) conveys only that you may draw exactly three cards and may draw fewer than three cards, not that you may draw any number of cards in the range $[0,3]$ (though the sentence is compatible with such a state of affairs). This partial view is also in line with claims about ignorance inferences, which have been argued to be partial, not total (Schwarz, 2016). We take a different stance, which is that these inferences are normally total, but in certain contexts can be partial. The system we propose will be able to account for such weakening (by 'pruning' certain alternatives from consideration), as we discuss at the end of the paper, in §4.4.

inverse-scope analysis (*at most three* > *allow*) derives only the speaker insecurity reading, and the reason for this is that when *at most three* takes widest scope, a speaker insecurity reading obligatorily emerges. For instance, (2) is incompatible with the speaker knowing exactly how many cards Ann drew; he knows it's either three or fewer than three, but is ignorant about which is true (Geurts and Nouwen, 2007; Büring, 2008; Nouwen, 2010, 2015; Schwarz, 2016).[3]

(2) Ann drew at most three cards.

Second, the antonym of *at most*, *at least*, does not give rise to authoritative readings. For example, (3) cannot be used to convey that you're allowed to draw three or more cards, with FC in the range $[3,n]$ (for some contextually specified UB n, e.g. 52, the number of cards in a standard deck), and a lower bound (LB) that prohibits drawing two cards or fewer. To see this, notice the oddity of embedding (3) under a expression like *The rules state that ...*, or of a game master uttering (3).[4]

(3) You're allowed to draw at least three cards.

To summarize, the puzzle has two parts: why does (1) have an authoritative reading, and why does (3) not have an authoritative reading? The most prominent and successful solution to this puzzle, as far as we are aware, is due to Penka (2014). She takes *at most* to be the oddity in this puzzle, and she solves it by decomposing *at most* into a negative component, *ANT*, plus *at least*, and giving (1) a split-scope analysis: *ANT three* scopes above *allow*, which in turn scopes above *at least* (hence, the scope of *at most three* is 'split').

(4) a. *at most three* $\rightsquigarrow$ *ANT* + *three* + *at least* (Penka, 2014)
b. ANT three [λn [allowed [at least n [λm you draw m cards]]]]
c. $\forall n \,.\, n > 3 \rightarrow \neg \diamondsuit [\geq n]$ $(\equiv \neg \diamondsuit [\geq 4])$

On this account, the literal meaning of (1) immediately entails an UB; FC then follows from neo-Gricean reasoning, given certain assumptions about the Horn scales responsible for generating alternatives. Together, these ingredients solve the first half of the puzzle. The second half is solved simply because *at least* has no analogous decomposition: parsing (3) with surface scope (*allow* > *at least three*) or with inverse scope (*at least three* > *allow*) leads only to an ignorance reading, not to a LB authoritative reading. (See Penka's paper for details.)

2. New observations

Penka's (2014) proposal attributes the contrast between *at most*, which gives rise to authoritative readings, and *at least*, which doesn't, to a special property of *at most*, namely that it's composed of a negative part, *ANT three*, that can take scope separately from its remainder, *at least*. We

3. There are a number of proposals for how to derive ignorance readings (e.g. Geurts and Nouwen, 2007; Büring, 2008; Nouwen, 2010; Coppock and Brochhagen, 2013; Nouwen, 2015; Schwarz, 2016). Our point here is that whichever one of these accounts one adopts, the prediction is that parsing (1) with *at most three* taking scope above *allow* will necessarily derive only an ignorance reading (not an authoritative reading) by whatever mechanism the account assumes is responsible for (2).

4. One might think that an authoritative reading of (3) is somehow 'blocked' due to competition with the more natural and straightforward sentence *You're required to draw at least three cards*. We discuss (and provide arguments against) this idea in §4.3.

now show that a variety of expressions, beyond just *at most*, give rise to authoritative readings, including both members of certain other antonym pairs.[5]

To start, (5) has a natural authoritative reading characterized by a LB inference (the catering premises may not open earlier than 5:00 AM) and a FC inference (they may open at 5:00 AM and at any time later than that).

(5) The catering premises may open at the earliest at 5:00 AM.

Similarly, (6) has a natural authoritative reading characterized by an UB inference (deductions may not occur later than the time of submission) and a FC inference (they may occur at the time of submission and at any time earlier than that).

(6) Deductions may occur at the latest at the time of submission.

Thus, unlike *at least* and *at most*, both of *at the earliest* and *at the latest* give rise to authoritative readings.[6]

Furthermore, other numeral modifiers besides *at most* can have authoritative construals. For example, (7), with *between three and seven*, has both a LB inference (the Speaker is not allowed to appoint fewer than three MPs, maybe because that one or those two MPs would hold too much power) and an UB inference (the Speaker is not allowed to appoint more than seven MPs, maybe because that would give, overall, too much power to too many MPs), as well as a FC inference (the Speaker is allowed to appoint any number of MPs in the range $[3,7]$).

(7) The Speaker is allowed to appoint between three and seven MPs to exercise his powers to issue recess writs when he is out of the country.

What's more, we find cases where *at least* actually *does* give rise to a LB authoritative reading, e.g. when conjoined with *at most*. For instance, (8c) conveys that the guild may not have fewer than three members, nor more than 100, and that it may have any number of members in the range $[3,100]$.

(8) a. You're allowed to nominate at least two and at most four authors this week!
b. Each bidder is allowed to bid for at least five lots and at most fifteen lots.
c. The guild may have at least three and at most 100 members.

Finally, we note the robust contrast in (9). In particular, while *at least* in (9a) cannot be construed authoritatively, in (9b) it can be: if you choose to write a paper, then you're not allowed to write fewer than three pages (LB inference), and you're allowed to write exactly three, four, ... pages (FC inference).

5. The examples in (5–8) come from Google and the Wikipedia corpus.
6. One potentially important difference between *at least/most* and *at the earliest/latest* is the presence of the definite determiner *the* in the latter. We're unsure exactly what role *the* plays here, but we observe that adding *the* to *at least* surprisingly improves sentences intended to be interpreted authoritatively: *(The rules state that) you're allowed to draw three cards at the least*, though not perfect, sounds much better to our ears than (3).

(9) The syllabus states that ...
  a. #You're allowed to write at least three pages.
  b. You're allowed to (either) give a presentation or (else) write at least three pages.

The LB authoritative uses of *at least* noted above are prima facie evidence against a decompositional account of the puzzle, i.e. an account that solves the puzzle by attributing a special property to *at most*.

Moreover, the availability of authoritative readings for both members of antonym pairs like *at the earliest* and *at the latest* casts doubt on a decompositional explanation of the asymmetry between *at least* and *at most*. For example, while it's possible to decompose *at the earliest* into something like 'not earlier than' and to decompose *at the latest* into something like 'not later than', doing so would undermine the explanation of the contrast between *at least* and *at most*: why not decompose *at least* into 'not fewer than'?[7]

Finally, extending the decompositional account to non-superlative expressions would require some rather ad hoc syntactic assumptions: *between three and seven* would need to decompose into something like 'not fewer than three and not more than seven'. It's unclear to us how such an approach would proceed, nor how conceptually appealing it would be.

Before concluding this section, it's important to establish that, just like for (1), an inverse-scope analysis of the above examples would not account for their authoritative readings. The reason is because, just like for *at most*, when *at the earliest*, *at the latest*, *between*, and *at least ... and at most* do not occur in the scope of an operator like *allow*, as in the examples in (10), ignorance readings obligatorily emerge, and consequently, an inverse-scope analysis would predict ignorance inferences across the board.[8]

7. To be more precise, the meaning of *ANT*, for Penka (2014), encodes the relation >: *ANT d D* (*D* a predicate of degrees) means $\neg\exists d'.d' > d \wedge d' \in D$. Presumably, this could be extended to *at the latest*, where > would refer not to the 'greater than' ordering over degrees, but to the 'later than' ordering over points (or intervals) of time: *ANT t T* (*T* a predicate of points of time) would mean $\neg\exists t'.t' > t \wedge t' \in T$. However, using *ANT* to decompose *at the earliest* wouldn't work; for that, one would need a different expression that involves the 'earlier than' relation, <. Call it $ANT_<$. The point now is that one would need a principled explanation why *at least* does not decompose into something involving $ANT_<$, which would lead to LB authoritative readings for sentences like (3). In other words, one would need to explain why *at least* is exceptional in this regard.

8. That *between* obligatorily implies ignorance (like superlative modifiers) is debatable. For instance, in his influential article on modified numerals, Nouwen (2010) classifies *at least* and *at most* as class B modifiers (they obligatorily imply ignorance when unembedded), but he classifies *between* (in contrast to *from ... to*) as class A. However, Nouwen does not devote any detailed discussion to *between*, and we think that it patterns a lot like superlative modifiers, e.g., when it comes to Nouwen's hexagon example:

(i) A hexagon has $\left\{\begin{array}{c}\text{more than three/fewer than eleven}\\ \text{\#at least four/at most ten}\\ \text{\#between four and ten}\end{array}\right\}$ sides.

Nevertheless, as Benjamin Spector points out to us, there may still be some truth to the notion that the potential ignorance inferences associated with *between* are more fragile than those associated with superlative modifiers. If so, then an inverse-scope analysis of examples like (7) might be possible after all. However, one would still need to account for the rest of our new observations. Instead, we hope to show that all the observations can be subsumed under a single account.

(10) a. Bill arrived at the earliest (latest) at 8:00 AM.
b. The Speaker appointed between three and seven MPs to excercise his powers to issue recess writs when he is out of the country.
c. Cindy nominated at least two and at most four authors.

The take-home message is that it's not *at most* which is the oddity in this puzzle, nor is it *at least*, per se. Rather, it's more specifically *at least* in just some sentences, like (3), but not (8c) or (9b). We therefore want to try to analyze all the above cases uniformly, without proposing that *at most* or any other expression has any special (non-standard) morphosyntax.

3. Proposal

3.1. Free-choice disjunction

Our starting point is to highlight some striking similarities between authoritative readings of *at most* in the scope of *allow* and FC readings of disjunction in the scope of *allow*.[9] Specifically, disjunction in the scope of an existential modal like *allow* licenses similar FC and 'bound' inferences. For example, in a context where the relevant desserts are cake, gelato, and pie, (11) licenses two inferences: (i) you're not allowed to have pie (a kind of 'bound' inference, in the sense that, as far as desserts go, you're bound to the set {cake, gelato}), and (ii) you're allowed to have cake, and you're allowed to have gelato (a FC inference) (von Wright, 1969; Kamp, 1973).[10] However, similar to what we saw for (1), a surface-scope analysis of (11) predicts a weak literal meaning, 'There is a permissible world in which you have cake or gelato' ($\diamondsuit c \vee g$), from which neither a 'bound' inference ($\neg \diamondsuit p$) nor a FC inference ($\diamondsuit c \wedge \diamondsuit g$) logically follows.

(11) You're allowed to have cake or gelato.

An influential account of FC disjunction is due to Fox (2007), who proposes that the meaning of (11) is a strengthened version of the weak surface-scope meaning $\diamondsuit(c \vee g)$. Strengthening is due to the presence of a grammatical device, *exh*, a covert analog of *only* that is responsible for scalar implicatures (Chierchia et al., 2012). What *exh S* means is that *S* is true and that each 'innocently excludable' member of $\text{alt}(S)$ (the set of alternatives of *S*) is false, where, intuitively, a proposition $q \in \text{alt}(S)$ is innocently excludable (relative to $[\![S]\!]$) just in case the negation of q doesn't contradict $[\![S]\!]$ and also doesn't force any disjunction of members of $\text{alt}(S)$ to be true (unless that disjunction is already entailed by $[\![S]\!]$).[11]

(12) $[\![\text{exh } S]\!] = [\![\text{exh}]\!](\text{alt}(S))([\![S]\!])$,
where $[\![\text{exh}]\!](A)(p) = p \wedge \bigwedge\{\neg q : q \in \text{IE}(p,A)\}$,
and $\text{IE}(p,A) = \bigcap\{A' : A' \text{ is a maximal subset of } A \text{ s.t. } \{p\} \cup \{\neg q : q \in A'\} \text{ is consistent}\}$

9. Comparing superlative modifiers with disjunction is by no means a novel idea. Büring's (2008) landmark paper established an exciting line of inquiry into how to analyze the ignorance inferences associated with *at least* and *at most* by building on well-understood accounts of the ignorance inferences associated with *or*. However, to our knowledge, no one has extended this line of inquiry to analyze authoritative construals of *allow . . . at most* along the same lines as that of FC construals of *allow . . . or*.
10. There is, potentially, a third inference, viz. that you're not allowed to have both cake and gelato ($\neg \diamondsuit (c \wedge g)$). We ignore this exclusivity inference henceforth, as there is no detectable analog to it when it comes to modified numerals (it's logically impossible to, e.g., draw exactly two cards and exactly three cards).
11. We write 'q is IE (relative to p and A)' to mean $q \in \text{IE}(p,A)$, and we omit 'relative to p and A' when it's clear from context.

For example, suppose that alt(John ate cake or gelato) = $\{c \vee g,\ c,\ g,\ c \wedge g\}$ (Sauerland, 2004; Fox, 2007). Then, relative to $c \vee g$, neither c nor g is IE: negating c would force g to be true, and conversely, negating g would force c to be true. (While it's consistent with $c \vee g$ to negate c or to negate g, the intuition is that we can't arbitrarily pick which of c or g to negate; thus, we don't negate either one.) However, $c \wedge g$ is IE, and so $[\![$exh [John ate cake or gelato]$]\!]$ = $(c \vee g) \wedge \neg(c \wedge g)$, i.e. the proposition that John ate cake or gelato but not both, which corresponds to the exclusive interpretation of ordinary disjunction.

Crucially, recursive application of *exh* is predicted to be possible. In a simple case like above, additional applications are vacuous, but when disjunction occurs in the scope of *allow*, recursive exhaustification is not vacuous and in fact delivers precisely the attested FC reading. To see this, consider the LF in (13). Let's start with the constituent *exh S*, and assume that $\text{alt}(S) = \{\diamond(c \vee g),\ \diamond c,\ \diamond g,\ \diamond p\}$.[12] Then, relative to $[\![S]\!]$ and alt(S), we see that $\diamond p$ is IE, but neither $\diamond c$ nor $\diamond g$ is IE: negating $\diamond c$ would force $\diamond g$ to be true, and vice versa. Thus, $[\![\text{exh } S]\!] = \diamond(c \vee g) \wedge \neg \diamond p$, the proposition that there is a permissible world in which you have cake or gelato, and no permissible world in which you have pie. Thus, the inner *exh* immediately derives the 'bound' inference (pie is not allowed), but not the FC inferences.

(13) exh [exh [$_S$ allowed [you have cake or gelato]]]

Let's move now from *exh S* to the entire sentence, *exh exh S*. Assume, along with Fox (2007), that alt(exh S) is the set of strengthened alternatives to S.

(14) $\text{alt}(\text{exh } S) = \{[\![\text{exh}]\!](\text{alt}(S))(p) : p \in \text{alt}(S)\}$

Computing alt(exh S) requires computing $[\![\text{exh}]\!](\text{alt}(S))(p)$ for each $p \in \text{alt}(S)$. This is given below.

(15) $$\text{alt}(\text{exh } S) = \left\{ \begin{array}{l} [\![\text{exh}]\!](\text{alt}(S))(\diamond(c \vee g)), \\ [\![\text{exh}]\!](\text{alt}(S))(\diamond c), \\ [\![\text{exh}]\!](\text{alt}(S))(\diamond g), \\ [\![\text{exh}]\!](\text{alt}(S))(\diamond p) \end{array} \right\} = \left\{ \begin{array}{l} \diamond(c \vee g) \wedge \neg \diamond p, \\ \diamond c \wedge \neg \diamond g \wedge \neg \diamond p, \\ \diamond g \wedge \neg \diamond c \wedge \neg \diamond p, \\ \diamond p \wedge \neg \diamond c \wedge \neg \diamond g \end{array} \right\}$$

Having computed alt(exh S), we are now in a position to compute $[\![\text{exh}]\!](\text{alt}(\text{exh } S))([\![\text{exh } S]\!])$, i.e. $[\![\text{exh } [\text{exh } S]]\!]$. It's the conjunction of $[\![\text{exh } S]\!]$ (computed above as $\diamond(c \vee g) \wedge \neg \diamond p$) and the negation of all innocently excludable alternatives of *exh S*, i.e. the negation of the members of some subset of (15). Which subset? The answer is $\text{alt}(\text{exh } S) - \{[\![\text{exh } S]\!]\}$, i.e. every proposition except the strengthened meaning of S is IE. Thus, the overall meaning is $[\![\text{exh } S]\!]$ conjoined with the negation of each of these alternatives. We give the negations of all the propositions in (15) below, but for clarity's sake, we write them as material implications.[13]

12. We could also include the alternative $\diamond(c \wedge g)$, which would turn out to be IE, and its negation would deliver the exclusivity inference mentioned in fn. 10.

13. This is licit because $\neg(p \wedge \neg q \wedge \neg r) \equiv p \rightarrow (q \vee r)$.

(16) a. $\diamond(c \vee g) \rightarrow \diamond p$ (contradicts $[\![\text{exh } S]\!]$)
b. $\diamond c \rightarrow (\diamond g \vee \diamond p)$
c. $\diamond g \rightarrow (\diamond c \vee \diamond p)$
d. $\diamond p \rightarrow (\diamond c \vee \diamond g)$ (already entailed by $[\![\text{exh } S]\!]$)

Putting all the pieces together, the meaning of *exh exh S* is $\diamond(c \vee g) \wedge \neg \diamond p$ (the strengthened meaning of *S*, i.e. $[\![\text{exh } S]\!]$), plus the conjunction of the material implications $\diamond c \rightarrow (\diamond g \vee \diamond p)$ and $\diamond g \rightarrow (\diamond c \vee \diamond p)$. Some reflection reveals that these implications, together with $[\![\text{exh } S]\!]$, are equivalent to $\diamond c$ and $\diamond g$.[14] Thus, the overall meaning derived is the following, which is precisely the reading we wanted to derive.

(17) $[\![$exh [exh [$_S$ allowed [you have cake or gelato]]]$]\!]$

$$= \underbrace{\underbrace{\underbrace{\diamond(c \vee g)}_{\text{weak meaning}} \wedge \underbrace{\neg \diamond p}_{\text{`bound'}}}_{\text{strengthened meaning}} \wedge \underbrace{\diamond c \wedge \diamond g}_{\text{FC}}}_{\text{recursively strengthened meaning}} \qquad (\equiv \diamond c \wedge \diamond g \wedge \neg \diamond p)$$

Our proposal in a nutshell. Given the inferential (and syntactic) similarity between (1) (*You're allowed to draw at most three cards*) and (11) (*You're allowed to have cake or gelato*), our goal is to provide a surface-scope analysis of (1), whereby its weak meaning is recursively strengthened as we just saw for FC disjunction. More precisely, we propose that (1) is parsed with two applications of *exh*, which results in an authoritative reading, as schematized in (18).

(18) $[\![$exh [exh [$_S$ allowed [you draw at most three cards]]]$]\!]$

$$= \underbrace{\underbrace{\underbrace{\diamond[\leq 3]}_{\text{weak meaning}} \wedge \underbrace{\neg \diamond [\geq 4]}_{\text{UB}}}_{\text{strengthened meaning}} \wedge \underbrace{\diamond[=3] \wedge \diamond[=2] \wedge \ldots}_{\text{FC}}}_{\text{recursively strengthened meaning}}$$

The clear similarities between (17) and (18) may give the impression that a solution to (the first half of) the puzzle is straightforward. However, as we hope to show, although there is a solution in sight, it ultimately requires a conservative, but non-trivial amendment to the theory of exhaustification. We therefore proceed in three steps: first, we present what seems like a straightforward account and show the problem it faces; second, we propose an ad hoc amendment that fixes the problem; finally, we show that our amendment can be generalized in a non-stipulative way.

14. More precisely:
$$\begin{aligned} & \diamond(c \vee g) \wedge \neg \diamond p \wedge (\diamond c \rightarrow (\diamond g \vee \diamond p)) \wedge (\diamond g \rightarrow (\diamond c \vee \diamond p)) \\ \equiv\ & (\diamond c \vee \diamond g) \wedge \neg \diamond p \wedge (\diamond c \rightarrow (\diamond g \vee \diamond p)) \wedge (\diamond g \rightarrow (\diamond c \vee \diamond p)) \\ \equiv\ & (\diamond c \vee \diamond g) \wedge \neg \diamond p \wedge (\diamond c \rightarrow \diamond g) \wedge (\diamond g \rightarrow \diamond c) \\ \equiv\ & (\diamond c \vee \diamond g) \wedge \neg \diamond p \wedge \diamond c \wedge \diamond g \\ \equiv\ & \diamond c \wedge \diamond g \wedge \neg \diamond p \end{aligned}$$

3.2. *At most*: first attempt

To achieve our goal, it's necessary (and sufficient) to show that the sets of alternatives $\mathrm{alt}(S)$ and $\mathrm{alt}(\mathrm{exh}\ S)$ derived by standard means (e.g. Katzir, 2007; Fox and Katzir, 2011) deliver the right results. Demonstrating this is difficult because of the space of possible alternative sets; however, we believe that the following characterization is plausible: $\mathrm{alt}(S)$ is the set of propositions obtained by replacing *at most* with *fewer than*, *exactly*, *at least* or *more than*, plus those obtained by replacing *three* with any numeral, plus those obtained by doing both replacements simultaneously. Thus, we assume the following equality.[15]

(19) $$\mathrm{alt}(S) = \begin{array}{l} \{\diamond[<n] : n \in \mathbb{N}_0\} \\ \cup\{\diamond[\leq n] : n \in \mathbb{N}_0\} \\ \cup\{\diamond[=n] : n \in \mathbb{N}_0\} \\ \cup\{\diamond[\geq n] : n \in \mathbb{N}_0\} \\ \cup\{\diamond[>n] : n \in \mathbb{N}_0\} \end{array} = \begin{array}{l} \{\diamond[\leq n] : n \in \mathbb{N}_0\} \\ \cup\{\diamond[=n] : n \in \mathbb{N}_0\} \\ \cup\{\diamond[\geq n] : n \in \mathbb{N}_0\} \end{array}$$

The first step is to compute $[\![\mathrm{exh}\ S]\!] = [\![\mathrm{exh}]\!](\mathrm{alt}(S))([\![S]\!])$. All alternatives of the form $\diamond[=n]$ and $\diamond[\geq n]$, for $n \geq 4$, are IE. As a result, we negate all such alternatives, which immediately derives our UB. Moreover, alternatives that 'overlap' with (are compatible with) $\diamond[\leq 3]$, e.g. $\diamond[\geq 1]$, $\diamond[\leq 2]$, and $\diamond[=2]$, are not IE; for instance, negating $\diamond[=2]$ would entail the disjunction $\diamond[=0] \vee \diamond[=1] \vee \diamond[=3]$. So the first round of exhaustification delivers the following result.

(20) $$\begin{aligned} [\![\mathrm{exh}]\!](\mathrm{alt}(S))(\diamond[\leq 3]) &= \diamond[\leq 3] \wedge \neg\diamond[=4] \wedge \neg\diamond[\geq 4] \wedge \neg\diamond[=5] \wedge \neg\diamond[\geq 5] \wedge \ldots \\ &\equiv \underbrace{\underbrace{\diamond[\leq 3]}_{\text{weak meaning}} \wedge \underbrace{\neg\diamond[\geq 4]}_{\text{UB}}}_{\text{strengthened meaning}} \end{aligned}$$

Turning now to the recursive layer of exhaustification, let's continue to assume, with Fox (2007), that $\mathrm{alt}(\mathrm{exh}\ S)$ is the set of strengthened alternatives to S, repeated below from (14).

(14) $$\mathrm{alt}(\mathrm{exh}\ S) = \{[\![\mathrm{exh}]\!](\mathrm{alt}(S))(p) : p \in \mathrm{alt}(S)\}$$

The hope is that we now derive the desired FC inferences: $\bigwedge\{\diamond[=n] : n \in \mathbb{N}_0 \cap [0,3]\}$. However, it turns out that the FC we derive is too weak. To see why, we first need to compute $\mathrm{alt}(\mathrm{exh}\ S)$.

(21) $\mathrm{alt}(\mathrm{exh}\ S) =$

$$\begin{array}{l} \{[\![\mathrm{exh}]\!](\mathrm{alt}(S))(\diamond[=n]) : n \in \mathbb{N}_0\} \\ \cup\{[\![\mathrm{exh}]\!](\mathrm{alt}(S))(\diamond[\leq n]) : n \in \mathbb{N}_0\} \\ \cup\{[\![\mathrm{exh}]\!](\mathrm{alt}(S))(\diamond[\geq n]) : n \in \mathbb{N}_0\} \end{array} = \begin{array}{l} \{\diamond[=n] \wedge \neg\diamond[<n] \wedge \neg\diamond[>n] : n \in \mathbb{N}_0\} \\ \cup\{\diamond[\leq n] \wedge \neg\diamond[>n] : n \in \mathbb{N}_0\} \\ \cup\{\diamond[\geq n] \wedge \neg\diamond[<n] : n \in \mathbb{N}_0\} \end{array}$$

Next, we must determine which members of $\mathrm{alt}(\mathrm{exh}\ S)$ are IE relative to $[\![\mathrm{exh}\ S]\!] = \diamond[\leq 3] \wedge \neg\diamond[\geq 4]$. As before, we list all the negations as material implications.

15. We assume here that the equivalences $\diamond[<n] \equiv \diamond[\leq (n-1)]$ and $\diamond[>n] \equiv \diamond[\geq (n+1)]$ hold, i.e. that the relevant scale of numbers operative in (1) is the natural numbers. However, we believe that assuming a dense scale, as proposed by Fox and Hackl (2006), would not alter our main results.

(22) a. $\{\diamondsuit[=n] \to (\diamondsuit[<n] \vee \diamondsuit[>n]) : n \in \mathbb{N}_0\}$
b. $\{\diamondsuit[\leq n] \to \diamondsuit[>n] : n \in \mathbb{N}_0\}$ (contradicts $[\![\text{exh } S]\!]$ when $n \geq 3$)
c. $\{\diamondsuit[\geq n] \to \diamondsuit[<n] : n \in \mathbb{N}_0\}$ (contradictory when $n = 0$)

We see that all the strengthened *exactly n* alternatives are IE, all the strengthened *at most n* alternatives for $n < 3$ are IE (but not for $n \geq 3$), and all the strengthened *at least n* alternatives for $n > 0$ are IE (but for $n \geq 4$, $[\![\text{exh } S]\!]$ already entails their negation). Putting all the pieces together, the meaning of *exh exh S* is $\diamondsuit[\leq 3] \wedge \neg\diamondsuit[\geq 4]$ (the strengthened meaning of S, i.e. $[\![\text{exh } S]\!]$), plus the conjunction of the material implications below.

(23) a. $\{\diamondsuit[=n] \to (\diamondsuit[<n] \vee \diamondsuit[>n]) : n \in \mathbb{N}_0\}$
b. $\{\diamondsuit[\leq n] \to \diamondsuit[>n] : n \in \mathbb{N}_0 \cap [0,2]\}$
c. $\{\diamondsuit[\geq n] \to \diamondsuit[<n] : n \in \mathbb{N}_0 \cap [1,\infty)\}$

Some reflection reveals that these inferences, together with $\diamondsuit[\leq 3] \wedge \neg\diamondsuit[\geq 4]$, are equivalent to $\diamondsuit[=0]$ and $\diamondsuit[=3]$. In other words, the FC inference we derive for (1) is that you're allowed to draw 0 cards and you're allowed to draw exactly 3 cards, which is weaker than what we were hoping to derive: we're missing the inferences $\diamondsuit[=1]$ and $\diamondsuit[=2]$. We take it that this result is incorrect.[16]

3.3. *At most*: second attempt

Intuitively, the problem we just encountered is that the members of $\text{alt}(\text{exh } S)$ are too strong, hence their negation results in inferences that are too weak. What we need is for $\text{alt}(\text{exh } S)$ to include weaker alternatives, whose exclusion will result in stronger inferences, namely total FC.[17] Our idea is that the set of alternatives for the second *exh* includes not just all strengthened propositions taken from $\text{alt}(S)$; rather, it includes all strengthened alternatives taken from the disjunctive closure of $\text{alt}(S)$, notated as $\text{alt}(S)^{\vee}$. (For now, this is just a stipulation; in the next subsection, show that if we assume that sets of alternatives are always closed under disjunction, we retain our main result for recursive exhaustification without disrupting cases of non-recursive exhaustification.)

(24) $\text{alt}(\text{exh } S) = \{[\![\text{exh}]\!](\text{alt}(S))(p) : p \in \text{alt}(S)^{\vee}\}$ (cf. (14))

By allowing p in (24) to range over propositions taken from the disjunctive closure of $\text{alt}(S)$, we introduce weaker alternatives into $\text{alt}(\text{exh } S)$, whose exclusion turns out to derive total FC. To see this, let $p = \diamondsuit[=0] \vee \diamondsuit[=1] \vee \diamondsuit[=3]$. p is not in $\text{alt}(S)$, but it is in $\text{alt}(S)^{\vee}$. $[\![\text{exh}]\!](\text{alt}(S))(p)$, which is now in $\text{alt}(\text{exh } S)$, is equal to $p \wedge \neg\diamondsuit[=2] \wedge \neg\diamondsuit[\geq 4]$. This strengthened alternative is

16. What we derive is a kind of partial FC, which is stronger than the partial FC that Penka (2014) derives (cf. fn. 2), and weaker that the total FC that we want to derive. In the oral talk version of this material, we assumed for simplicity that only the *exactly* alternatives were active, and we showed that this derives an even weaker FC inference than what we derive here: namely, that there are (at least) two distinct numbers m and n in $[0,3]$ such that $\diamondsuit[=m]$ and $\diamondsuit[=n]$. Again, we take this result to be incorrect.
17. This method is not a guaranteed recipe for deriving stronger inferences overall: the addition of weaker alternatives could instead introduce new symmetries (in the sense of Fox (2007)), which would result in fewer inferences overall. (Thanks to Emmanuel Chemla for stressing this point to us.) However, our amendment turns out not to do this, but instead has the desired effect.

IE relative to alt(exh S) and $[\![\text{exh } S]\!]$. Its negation, which is equivalent to $p \to (\diamond[=2] \vee \diamond[\geq 4])$, and $[\![\text{exh } S]\!]$, which is $\diamond[\leq 3] \wedge \neg\diamond[\geq 4]$, jointly entail $\diamond[=2]$. Extrapolating from this example, we see that we derive $\diamond[=n]$ for every $n \in \mathbb{N}_0 \cap [0,3]$, just as desired.

Thus, the overall meaning derived for (1) is:

(25) $$[\![\text{exh}]\!]\Big(\text{alt}(\text{exh } S)\Big)\Big([\![\text{exh}]\!]\big(\text{alt}(S)\big)\big([\![S]\!]\big)\Big)$$
$$= \underbrace{\underbrace{\underbrace{\diamond[\leq 3]}_{\text{weak meaning}} \wedge \underbrace{\neg\diamond[\geq 4]}_{\text{UB}}}_{\text{strengthened meaning}} \wedge \underbrace{\diamond[=3] \wedge \diamond[=2] \wedge \diamond[=1] \wedge \diamond[=0]}_{\text{FC}}}_{\text{recursively strengthened meaning}}$$

3.4. Generalizing the amendment

Instead of stipulating something special about how alt(exh S) is computed, we now show that our amendment can be generalized in a simple way: for *any* sentence S, we take alt(S) to be closed under disjunction. Formally, if $\text{alt}'(S)$ is the set of usual alternatives of S, e.g. those derived by the algorithm proposed by Fox and Katzir (2011), then $\text{alt}(S) = \text{alt}'(S)^{\vee}$. The meaning of *exh* is the same as before, as given in (12), and in particular still refers to alt(S); it's just that alt(S) has been updated to be closed under disjunction.

Importantly, for a non-recursive occurrence of *exh*, this amendment makes no difference, in the sense that we don't derive stronger inferences than before. For instance, assume, for some sentence S, that $\text{alt}'(S)$ includes p and q but not $p \vee q$. Exhaustifying S now involves potentially excluding $p \vee q$, since this proposition is in alt(S). However, excluding $p \vee q$ is possible ($p \vee q$ is IE) if and only if excluding both p and q is possible (p and q are both IE). And in turn, excluding $p \vee q$ is logically equivalent to excluding both p and q. Thus, we don't gain anything at this level: either p and q are both IE, in which case so is $p \vee q$, and negating the former is equivalent to negating the latter; or p and q are not both IE, in which case $p \vee q$ is not IE, and not negating both p and q is equivalent to not negating $p \vee q$. (See also Spector (2017).)

The effect only surfaces for recursive *exh*, and here's why. The overall meaning is of *exh exh S* is:

(26) $$[\![\text{exh}]\!]\Big(\text{alt}(\text{exh } S)\Big)\Big([\![\text{exh}]\!]\big(\text{alt}(S)\big)\big([\![S]\!]\big)\Big)$$

Again following Fox (2007), alt(exh S) is the set of strengthened alternatives of S,[18] but now we get something different from before our pre-amendment set in (14). What we get is precisely what we stipulated earlier, in (24), except now we derive it as a consequence of closing *all* alternative sets under disjunction.

18. Whether this set, too, is closed under disjunction doesn't matter: closing it under disjunction doesn't result in any additional inferences, for the same reason that closing the usual alternative set under disjunction for non-recursive exhaustification doesn't matter. (At least, this is true for double exhaustification. It could, in principle, matter for triple exhaustification, i.e. for a sentence of the form *exh exh exh S*, but we don't pursue this line of inquiry here.)

(27) $$\begin{aligned}\mathrm{alt}(\mathrm{exh}\ S) &= \{[\![\mathrm{exh}]\!](\mathrm{alt}(S))(p) : p \in \mathrm{alt}(S)\} \\ &= \{[\![\mathrm{exh}]\!](\mathrm{alt}'(S)^{\vee})(p) : p \in \mathrm{alt}'(S)^{\vee}\} \\ &= \{[\![\mathrm{exh}]\!](\mathrm{alt}'(S))(p) : p \in \mathrm{alt}'(S)^{\vee}\} \\ &= (24)\end{aligned}$$

Finally, we note that closing the alternative set under disjunction allows us to assume that for (1), alt′ delivers just the *exactly* alternatives, and this is because any *at least* or *at most* alternative can be written as a disjunction of *exactly* alternatives.

(28) $$\mathrm{alt}(S) = \begin{pmatrix} \{\diamondplus[\leq n] : n \in \mathbb{N}_0\} \\ \cup\{\diamondplus[= n] : n \in \mathbb{N}_0\} \\ \cup\{\diamondplus[\geq n] : n \in \mathbb{N}_0\}\end{pmatrix}^{\vee} = \{\diamondplus[= n] : n \in \mathbb{N}_0\}^{\vee}$$

We don't claim this is in fact what alt′ does, but making this assumption will simplify discussion of other examples in §4.2.

4. Discussion

4.1. Free-choice disjunction revisited

Importantly, our amendment doesn't disrupt the analysis of FC disjunction like (11) (*You're allowed to have cake or gelato*). The reason is because, in this case, closing the alternative set under disjunction doesn't add any new IE alternatives. For example, if, as we assumed earlier, $\mathrm{alt}'(S) = \{\diamondplus(c \vee g), \diamondplus c, \diamondplus g, \diamondplus p\}$, then alt$(S)$ is as follows, where the underlined alternatives are those that we gain (relative to before) by closing the set of alternatives under disjunction:[19]

(29) $$\mathrm{alt}(S) = \left\{\begin{matrix} \diamondplus c, \quad \diamondplus g, \quad \diamondplus p, \\ \diamondplus(c \vee g)\end{matrix}\right\}^{\vee} = \left\{\begin{matrix} & \diamondplus c, \quad \diamondplus g, \quad \diamondplus p & \\ \diamondplus(c \vee g), & \underline{\diamondplus(c \vee p)}, & \underline{\diamondplus(g \vee p)} \\ & \underline{\diamondplus(c \vee g \vee p)} & \end{matrix}\right\}$$

For the first round of exhaustification, none of the three new alternatives are IE, and so $[\![\mathrm{exh}\ S]\!]$ is still $\diamondplus(c \vee g) \wedge \neg\diamondplus p$. For the second round, we compute alt(exh S).

(30) $$\mathrm{alt}(\mathrm{exh}\ S) = \left\{\begin{matrix} [\![\mathrm{exh}]\!](\mathrm{alt}(S))(\diamondplus c), \\ [\![\mathrm{exh}]\!](\mathrm{alt}(S))(\diamondplus g), \\ [\![\mathrm{exh}]\!](\mathrm{alt}(S))(\diamondplus p), \\ [\![\mathrm{exh}]\!](\mathrm{alt}(S))(\diamondplus(c \vee g)), \\ \underline{[\![\mathrm{exh}]\!](\mathrm{alt}(S))(\diamondplus(c \vee p))}, \\ \underline{[\![\mathrm{exh}]\!](\mathrm{alt}(S))(\diamondplus(g \vee p))}, \\ \underline{[\![\mathrm{exh}]\!](\mathrm{alt}(S))(\diamondplus(c \vee g \vee p))}\end{matrix}\right\} = \left\{\begin{matrix} \diamondplus c \wedge \neg\diamondplus g \wedge \neg\diamondplus p, \\ \diamondplus g \wedge \neg\diamondplus c \wedge \neg\diamondplus p, \\ \diamondplus p \wedge \neg\diamondplus c \wedge \neg\diamondplus g, \\ \diamondplus(c \vee g) \wedge \neg\diamondplus p, \\ \underline{\diamondplus(c \vee p) \wedge \neg\diamondplus g}, \\ \underline{\diamondplus(g \vee p) \wedge \neg\diamondplus g}, \\ \underline{\diamondplus(c \vee g \vee p)}\end{matrix}\right\}$$

The negations of all the members of alt(exh S) are provided below. As before, we write the negations of all the conjunctive alternatives as material implications.

19. Note that for any two propositions p and q, $\diamondplus p \vee \diamondplus q \equiv \diamondplus(p \vee q)$. We use the latter form in the following set.

(31) a. $\diamond c \rightarrow (\diamond g \vee \diamond p)$
b. $\diamond g \rightarrow (\diamond c \vee \diamond p)$
c. $\diamond p \rightarrow (\diamond c \vee \diamond g)$ (already entailed by $[\![\text{exh } S]\!]$)
d. $\diamond (c \vee g) \rightarrow \diamond p$ (contradicts $[\![\text{exh } S]\!]$)
e. $\underline{\diamond (c \vee p) \rightarrow \diamond g}$
f. $\underline{\diamond (g \vee p) \rightarrow \diamond c}$
g. $\underline{\neg \diamond (c \vee g \vee p)}$ (contradicts $[\![\text{exh } S]\!]$)

Previously, we saw that (a–b), together with $[\![\text{exh } S]\!]$, amount to the FC inferences $\diamond c$ and $\diamond g$. Now, in addition to those, there are three potential new inferences, (e–g); however, (e) and (f), together with $[\![\text{exh } S]\!]$, again amount to $\diamond c$ and $\diamond g$, and (g) contradicts $[\![\text{exh } S]\!]$. Overall, then, no new inferences are derived relative to what we derived earlier.

4.2. New observations revisited

Our account extends naturally to all the new data introduced in §2. We just have to assume that the relevant *exactly* alternatives are available. (Other alternatives may also be available, but recall that, since we take the set of alternatives to be closed under disjunction, the *exactly* alternatives already generate the complete set that we're interested in; therefore, we ignore those alternatives in what follows.)

For example, the alternative set of (6) (*Deductions may occur at the latest at the time of submission*) must include the (meanings of) alternatives of the form *Deductions may occur at (exactly)* t, for some t in the set of points of time T.

(32) $\text{alt}(6) = \{\diamond[=t] : t \in T\}^{\vee}$

This set is exactly the same as what we assumed for the alternative set of (1), except now we have points of time t instead of natural numbers n, and so it's no surprise that the results here will be the same. In particular, the first round of exhaustifying the meaning of (6) ($\diamond[\leq \mathsf{time.of.submission}]$) will exclude all *exactly* alternatives involving times later than the time of submission, since those alternatives are all IE, while the second (recursive) round will derive FC inferences regarding all *exactly* alternatives involving times equal to or earlier than the time of submission.

Similarly, let's assume the exact same set of alternatives for (5) (*The catering premises may open at the earliest at 5:00 AM*). Then, in a completely parallel way, the first round of exhaustifying the meaning of (5) ($\diamond[\geq \mathsf{5{:}00\ AM}]$) will exclude all *exactly* alternatives involving times earlier than 5:00 AM, since those alternatives are all IE, while the second (recursive) round will derive FC inferences regarding all *exactly* alternatives involving times equal to or later than 5:00 AM.

Finally, we turn to (7) (*The Speaker is allowed to appoint between three and seven MPs*), and assume the set of alternatives below. The first round of exhaustifying the meaning of (7) ($\diamond[3,\ldots,7]$) will exclude all *exactly* alternatives involving numbers less than 3 or more than 7, since those alternatives are all IE (this derives both a LB and an UB), while the second (recursive) round will derive FC inferences regarding all *exactly* alternatives involving numbers in the range $[3,7]$.

(33) $\text{alt}(7) = \{\diamondsuit[=n] : n \in \mathbb{N}_0\}^{\vee}$

Examples like (8c) (*The guild may have at least three and at most 100 members*) can be analyzed in the same way.

4.3. Open problem

The reader has by now probably noticed that our account also predicts a LB authoritative reading for sentences like (3) (*You're allowed to draw at least three cards*). The reason why we derive such a reading is essentially the same as the reason why we derive the attested reading of (5) (*The catering premises may open at the earliest at 5:00 AM*): recursive exhaustification excludes *exactly* alternatives with numbers below 3, and derives FC regarding all *exactly* alternatives with numbers greater than or equal to 3.

This prediction is only partially correct. Recall the contrast in (9), repeated below.

(9) The syllabus states that ...
  a. #You're allowed to write at least three pages.
  b. You're allowed to (either) give a presentation or (else) write at least three pages.

Unfortunately, we still have no explanation for this contrast, or more generally why *at least* only sometimes has a LB authoritative reading.

A tempting explanation for why (3) has no LB authoritative reading is that it's somehow 'blocked' by the more natural, straightforward, and seemingly equivalent sentence with *required*, given in (34). One rationale would be that, all else being equal, it's more economical to use *require* than to use *allow* with recursive exhaustification.

(34) You're required to draw at least three cards.

However, we're doubtful whether a coherent and convincing story along these lines can be told, and here's why: (3) (on what would be its LB authoritative reading) and (34) are not completely equivalent. Given the infelicity of using (3) authoritatively, this non-equivalence is hard to detect, but we can observe a clear contrast when we move to *between*. Specifically, (35a), on its authoritative reading, says nothing about whether the speaker is required to appoint any MPs at all; in fact, it seems to imply that not appointing any MPs at all is a possibility (it's just that, if the Speaker decides to appoint some number of MPs, that number must fall in the range $[3,7]$). By contrast, (35b) clearly excludes the possibility of not appointing any MPs.

(35) a. The speaker is allowed to appoint between three and seven MPs.
  b. The speaker is required to appoint between three and seven MPs.

We can highlight the contrast even further with the minimally different pair of examples in (36). In particular, if one of the laws of Absurdistan is (36a), then presumably you won't be breaking the law if you have no children; it's just that, if you have children, then the number of children you have must not exceed three. By contrast, if the law is (36b), then anyone (of child-bearing age, etc.) without children is a criminal.

(36) a. In Absurdistan, you're allowed to have between one and three children.
b. In Absurdistan, you're required to have between one and three children.

The point here is that a sentence of the form *allow* ... $[m,n]$, on its authoritative reading, is not in general equivalent to *required* ... $[m,n]$ (on its strengthened, FC reading): specifically, for $m > 0$, the *allow* sentence says nothing about the 'zero' case, whereas the *required* sentence explicitly excludes it, as the *between* examples above illustrate. Moreover, for $m = 0$, judgments are clear that the resulting *allow* and *require* sentences are not only both felicitous, but also equivalent to one another, which is unexpected under a blocking account.[20]

(37) a. In Absurdistan, you're allowed to have between zero and three children. ✓
$\Leftrightarrow$
b. In Absurdistan, you're required to have between zero and three children. ✓

(38) a. Abstracts may be at most three pages long. ✓
$\Leftrightarrow$
b. Abstracts must be at most three pages long. ✓

Thus, a blocking account of the infelicity of an authoritative use of (3) that relies on semantic equivalence between the blocking sentence, (34), and the blocked sentence, (3), cannot be right, for two reasons: (3) and (34) are not actually equivalent, and cases of actual equivalence, e.g. (37a) and (37b), and (38a) and (38b), do not in fact give rise to blocking effects.[21]

4.4. Concluding remarks

We showed that authoritative readings of sentences like (1), (7), and many others can be accounted for by assuming that their weak, surface-scope meanings are recursively strengthened in the same way that weak meanings of disjunctive permission sentences like (11) are. We conclude by addressing an important question raised by our analysis.

We assumed that FC inferences are total, not partial (cf. fn. 2), and our amended system of recursive exhaustification delivers exactly this. However, there are clear cases where FC is not total. For example, (39) is a perfectly coherent two-sentence discourse, but a blind application of our proposal predicts it to be contradictory: the second sentence should entail, e.g., that you're allowed to draw exactly five cards, which contradicts the first sentence.

20. Parallel observations arise even for FC disjunction. For instance, compare (11) (*You're allowed to have cake or gelato*), which of course does not entail that you must have a dessert, and *You're required to have cake or gelato*, which does. In addition, *You're allowed to have cake or gelato or neither* and *You're required to have cake or gelato or neither* are both intuitively equivalent, but the *require* sentence does not block the *allow* sentence.

21. Because we assume that *exactly n* alternatives include *exactly* 0 alternatives (n ranges over the set $\mathbb{N}_0$, which includes 0), our account *does* exclude the 'zero' cases for sentences like (7). However, this issue can be resolved simply by assuming that a numeral like *three* cannot be replaced by *zero*, an assumption that doesn't seem too far-feteched, given the odd nature of *zero* compared to all other numerals. Note that taking this route would then mean that the authoritative reading of (1), with *at most*, does not actually entail the possibility of drawing zero cards, but is nonetheless compatible with such a state of affairs. We think that this result is adequate and that the inference that drawing zero cards is possible (to the extent that it's available) is either a contextual entailment or an implicature.

(39) You're required to draw an even number of cards. You're allowed to draw at most ten cards.

We think that our account is compatible with such observations, once we acknowledge that relevance considerations can restrict what counts as an alternative (Fox and Katzir, 2011), and/or that alternatives may, under certain conditions, be 'pruned' (Crnič et al., 2015). In this particular case, one effect of the first sentence is to rule out *exactly n* alternatives where *n* is odd.

However, as intuitive as this resolution may appear, it must be constrained in certain ways. For example, FC regarding the numeral mentioned in the sentence seems to always be available, which suggests that the corresponding *exactly* alternative must always be active. Witness the oddity of the two-sentence discourse in (40): the second sentence implies that you're allowed to draw exactly nine cards, which contradicts the first sentence.

(40) # You're required to draw an even number of cards. You're allowed to draw at most nine cards.

An obvious line to pursue is to say that if a numeral *n* is mentioned, then the *exactly n* alternative is relevant and cannot be pruned. Independent support for this approach comes from FC disjunction: each disjunct, because it's explicitly mentioned, must be a possibility, as the oddity of the following discourse illustrates.[22]

(41) # You're required to eat only green vegetables (non-green vegetables are not allowed). You're allowed to eat broccoli, spinach, or red cabbage.

References

Büring, D. (2008). The least *at least* can do. In C. B. Chang and H. J. Haynie (Eds.), *West Coast Conference on Formal Linguistics (WCCFL)*, Volume 26, Somerville, MA, pp. 114–120. Cascadilla Proceedings Project. http://www.lingref.com/cpp/wccfl/26/paper1662.pdf.

Chierchia, G., D. Fox, and B. Spector (2012). Scalar implicature as a grammatical phenomenon. In C. Maienborn, K. von Heusinger, and P. Portner (Eds.), *Semantics: An International Handbook of Natural Language Meaning*, Volume 3, pp. 2297–2331. Berlin, Germany: Mouton de Gruyter.

Coppock, E. and T. Brochhagen (2013). Raising and resolving issues with scalar modifiers. *Semantics and Pragmatics 6*(3), 1–57. https://doi.org/10.3765/sp.6.3.

Crnič, L., E. Chemla, and D. Fox (2015). Scalar implicatures of embedded disjunction. *Natural Language Semantics 23*, 271–305. https://doi.org/10.1007/s11050-015-9116-x.

Fox, D. (2007). Free choice and the theory of scalar implicatures. In U. Sauerland and P. Stateva (Eds.), *Presupposition and Implicature in Compositional Semantics*, Palgrave Studies in

22. An important caveat: the parallelism drawn here is not perfect. In the case of FC disjunction, the non-prunable (obligatorily relevant) alternative is derived by replacing the full disjunction with a single disjunct (one of the ones mentioned); in the case of (40), the non-prunable alternative is derived by replacing *at most* with *exactly*. In some sense, the former case in more natural, insofar as the lexical material required for the alternative is fully present in the utterance (thus, the appeal to 'mentioning' is sensible), whereas in the latter case, it's not (*exactly* was not mentioned). However, both cases involve structural manipulation of the uttered sentence, and so from a formal perspective, there's little distinction between the two.

Pragmatics, Language and Cognition Series, Chapter 4, pp. 71–120. New York, NY: Palgrave Macmillan. https://doi.org/10.1057/9780230210752_4.

Fox, D. and M. Hackl (2006). The universal density of measurement. *Linguistics and Philosophy 29*(5), 537–586. https://doi.org/10.1007/s10988-006-9004-4.

Fox, D. and R. Katzir (2011). On the characterization of alternatives. *Natural Language Semantics 19*(1), 87–107. https://doi.org/10.1007/s11050-010-9065-3.

Geurts, B. and R. Nouwen (2007). *At least* et al.: The semantics of scalar modifiers. *Language 83*(3), 533–559. http://www.jstor.org/stable/40070901.

Kamp, H. (1973). Free choice permission. In *Proceedings of the Aristotelian Society*, Volume 74, pp. 57–74. http://www.jstor.org/stable/4544849.

Katzir, R. (2007). Structurally-defined alternatives. *Linguistics and Philosophy 30*(6), 669–690. https://doi.org/10.1007/s10988-008-9029-y.

Nouwen, R. (2010). Two kinds of modified numerals. *Semantics and Pragmatics 3*(3), 1–41. https://doi.org/10.3765/sp.3.3.

Nouwen, R. (2015). Modified numerals: The epistemic effect. In L. Alonso-Ovalle and P. Menéndez-Benito (Eds.), *Epistemic Indefinites: Exploring Modality Beyond the Verbal Domain*, Oxford Linguistics, Chapter 11, pp. 244–266. Oxford, England: Oxford University Press.

Penka, D. (2014). *At most* at last. In E. Csipak and H. Zeijlstra (Eds.), *Sinn und Bedeutung*, Volume 19, pp. 463–480. http://nbn-resolving.de/urn:nbn:de:bsz:352-0-304168.

Sauerland, U. (2004). Scalar implicatures in complex sentences. *Linguistics and Philosophy 27*(3), 367–391. https://doi.org/10.1023/B:LING.0000023378.71748.db.

Schwarz, B. (2016). Consistency preservation in Quantity implicature: The case of *at least*. *Semantics and Pragmatics 9*(1), 1–47. https://doi.org/10.3765/sp.9.1.

Spector, B. (2017). Comparing exhaustivity operators. *Semantics and Pragmatics 9*. https://doi.org/10.3765/sp.9.11.

von Wright, G. H. (1969). *An Essay on Deontic Logic and the Theory of Action*. Amsterdam, Netherlands.