Alsop, Anna. 2024. The pragmatics of free choice *any*. *Glossa: a journal of general linguistics* 9(1). pp. 1–36. DOI: https://doi.org/10.16995/glossa.9629

The pragmatics of free choice *any*

Anna Alsop, New York University, US, aalsop@nyu.edu

In this paper, I examine previous claims by Menéndez-Benito (2005; 2010) and Dayal (2013) that *You may read any book* entails *Every book may be read on its own*. In light of several key counterexamples, including those raised by Chierchia (2013) and Szabolcsi (2019), I argue that this is not an entailment at all, but a particularly robust implicature that arises in the pragmatics. In order to derive this implicature, I combine Szabolcsi's semantics for universal free choice items (FCIs) with a formal pragmatic model in the Rational Speech Act (RSA) framework. I update Champollion et al.'s (2019) RSA model of free choice with technical innovations from Franke & Bergen (2020) to model how a listener interprets an utterance containing the universal FCI *any* given a range of possible exhaustified parses. This pragmatic model predicts the robustness of the implicature observed by Menéndez-Benito and Dayal, and further predicts that it is more stable across varied contexts than other implicatures arising from the same utterance.

1 Introduction

English *any* is a universal free choice item (FCI), which can contribute universal force to a proposition as in (1a).

(1) a. You may read any book.
 b. $\forall x \in D_e\ [book(x) \rightarrow \Diamond read(you,x)]$

Universal FCIs differ from universal quantifiers (e.g. *every*) in several important respects. They are only felicitous in a limited set of modal environments, and there is strong cross-linguistic and diachronic evidence that the semantics of universal FCIs are closely related to those of existential quantifiers. I will focus on a set of recent implicature-based accounts, beginning with Chierchia (2013), that derive the meaning and distribution of universal FCIs by assuming that they are underlyingly existential quantifiers. This existential quantifier is obligatorily strengthened by a covert exhaustification operator *O* in the grammar to take on its characteristic universal force.

Subsequent implicature-based accounts of universal FCIs by Dayal (2013) and Szabolcsi (2019) take Chierchia's basic idea on board. But while these implicature-based accounts seek to derive the universal meaning indicated in (1b), there remains little consensus on the basic truth conditions of phrases containing universal FCIs. This comes from the intuition that the truth conditions of (1a) should be stronger than the logical expression in (1b). Crucially, Menéndez-Benito (2005; 2010) points out that upon hearing (1a), a listener should conclude that (2) must be true.

(2) You may read just Book A (without reading Book B, Book C…).

Following this intuition, Dayal (2013) formulates the Viability Constraint: a semantic constraint on universal FCIs that requires that every book may be read on its own. The Viability Constraint ensures that (1a) entails (2), capturing the intuitive interpretation of (1a). However, Szabolcsi (2019) points out that this constraint is too strong on the basis of the counterexample in (3).

(3) Any bishop may meet a bishop.

In order for (3) to be felicitous, the Viability Constraint requires that it entail (4).

(4) Bishop A may be the only bishop who meets a bishop.

Assuming that all deontically possible worlds are also metaphysically possible, (4) can never be true—if Bishop A meets Bishop B, Bishop B must also meet Bishop A. The Viability Constraint thus mistakenly predicts that sentences involving symmetric predicates in this way should always be infelicitous. To resolve this issue, Szabolcsi proposes a revision to the Viability Constraint, which does away with the prediction that (1a) entails (2). Instead, the Revised Viability Constraint simply requires that (1a) entail *For every book, you are permitted to but not required to read it*. This Revised Viability Constraint better captures the felicity of universal FCIs in sentences like (3), but now we are left with no explanation for why (1a) so robustly communicates the meaning in (2). In this paper, I argue that this is not a problem for Szabolcsi's semantic analysis, and provide a pragmatic explanation. The reading captured by Dayal's Viability Constraint is robust when it arises, but can be blocked by world knowledge, and as Szabolcsi points out, symmetric predicates. Based on this evidence, I conclude that Dayal's interpretation is the result of a particularly robust implicature, which I call the *exclusiveness* implicature (borrowing terminology from Menéndez-Benito 2005; 2010).

In the system of alternative semantics that Dayal and Szabolcsi both adopt from Chierchia, Dayal's truth conditions and Szabolcsi's truth conditions correspond to two different exhaustified parses (semantic meanings) of utterances containing free choice (FC) *any*. Upon hearing (1a), then, a listener is faced with uncertainty about which parse of the utterance a speaker intends. I model a listener who resolves this uncertainty using principles of Gricean reasoning. Given that Dayal's truth conditions rule out more worlds, a pragmatic listener will often arrive at this stronger interpretation with a high degree of certainty. However, when the relevant worlds picked out by Dayal's semantics are completely ruled out by semantic or pragmatic factors, the listener must fall back on Szabolcsi's weaker interpretation.

Relatedly, I find that other implicatures arising from utterances containing FCI *any* appear to be less robust than the exclusiveness implicature. These readings are easily eroded or boosted by the listener's prior expectations about the world prior to hearing the utterance. I suggest that this is because these readings do not arise from literal parses, so the listener is forced to rely on prior knowledge alone.

To formalize the process of Gricean reasoning for utterances containing FCI *any*, I construct a game-theoretic model in the Rational Speech Act (RSA) framework (Frank & Goodman 2012). This formalization extends key innovations from Champollion et al. (2019) and Franke & Bergen (2020) in order to model a listener entertaining multiple possible exhaustified parses of a single utterance. This pragmatic model predicts the robustness of the exclusiveness implicature in contexts where it is not blocked by world knowledge or the semantics of the predicate. Moreover, it captures the relative robustness of the exclusiveness implicature as compared to other possible implicatures.

2 Free choice *any*

FCI *any* is only felicitous in a limited set of environments. Its distribution, as well as the distribution of universal FCIs in other languages, has been the subject of extensive prior study (e.g. Horn 1972; Carlson 1980; Kadmon & Landman 1993; Haspelmath 1997; Dayal 2004; 2009; Chierchia 2006; 2013). Semantic accounts of FCI *any* are thus tasked with finding a single theory that captures not only the underlying meaning of this item, but also its distribution. Neither task is straightforward. I will begin with a brief overview of the distribution of FCI *any* in Section 2.1. In Section 2.2, I then review cross-linguistic and diachronic evidence that suggest FCI *any* and other universal FCIs are underlyingly existential quantifiers.

2.1 Distribution of free choice *any*

Following previous work, I will assume that *any* functions as both a negative polarity item (NPI) and a universal FCI. NPI *any* is licensed in a limited set of environments, most of which are downward entailing (Ladusaw 1979). For example, *any* is licensed under the scope of negation, as in (5a). The restriction of the quantifier *every*, another downward entailing environment, also licenses NPI *any*, as in (5b).

(5) NPI *any*
  a. Jane didn't buy any shirts.
  b. Tom thanked every guest who brought any drinks.

While NPI *any* is limited in its distribution, it is similar to plain indefinites (e.g. *a*, *some*) in that it contributes existential force. The distribution of NPI *any* is interesting in its own right, but will not be the focus of this paper.

In a second set of environments, *any* functions as a universal FCI, contributing universal force to a proposition. FCI *any* with an unmodified complement[1] is infelicitous in unmodalized episodic statements as in (6a), but is licensed by possibility modals, as in (6b). Not all modal environments license FCI *any*, however—necessity modals do not act as licensers, as can be seen in (6c).

(6) FCI *any*
  a. *Carrie read any book.
  b. Carrie may read any book.
  c. *Carrie must read any book.

The full distribution of FCI *any* is much more intricate, involving interactions with generic operators and covert epistemic modals, among other factors (see Dayal 2009; 2013 for a detailed treatment). To streamline the discussion, I will limit my focus to the basic facts summarized in (6).

2.2 Universal free choice items as existential quantifiers

Since FCI *any* contributes universal force to a proposition, it is tempting to treat it as a universal quantifier subject to distributional constraints. However, several striking cross-linguistic and historical facts militate against this analysis. Haspelmath (1997) finds that roughly half of the

[1] This construction is what Dayal (2013) calls unmodified *any*. This rules out partitive constructions (e.g. *any of the books*), *any*-numeral constructions (e.g. *any two books*), and other modified constructions (e.g. *any book that you like*).

languages in a large cross-linguistic survey of indefinites contain items that function as both a NPI and a FCI (English *any* falls into this category), suggesting semantic relatedness between the two. Additionally, Chierchia (2013) points out that English *any* arose from an Old English plain indefinite, which first became an NPI before further gaining the FCI reading. Typologically and historically, FCIs appear to be closely semantically related to existential items, namely NPIs and plain indefinites.

These facts motivate a treatment of all universal FCIs as underlyingly existential quantifiers (e.g. Horn 1972; Carlson 1980). A set of recent implicature-based accounts, beginning with Chierchia (2013), derive the meaning of FCI *any* by assuming that it is underlyingly an existential quantifier that undergoes obligatory strengthening to a universal meaning. Implicature-based accounts further impose distributional constraints on FCI *any* to capture the limited distribution I describe in Section 2.1.

3 Three implicature-based accounts

Implicature-based accounts of FCI *any*, including the account I will be developing, assume a framework of implicature calculation first introduced by Chierchia (2013). In this framework, sets of alternatives are built into the grammar, and alternative-sensitive operators may act on these alternatives to produce strengthened meanings. All items that lie on a grammaticized scale, such as quantifiers and logical connectives *and* and *or*, carry the feature bundle [σ, D]. The two features in the bundle each correspond to sets of alternatives associated with the scalar term. σ corresponds to the set of scalar alternatives (σA), comprised of the items on the same grammaticized scale. For example, for *some*, σA = {*some*,…*half*,…*most*, *all*}. In subsequent derivations, I will often omit intermediate members of the scale, illustrating the derivation with the minimal set σA = {*some*, *all*}. D corresponds to the set of domain alternatives (DA), comprised of each sub-domain of the quantificational domain. For example, if the restriction of *some* is *classes*, then DA is the set of all sub-domains of classes (*Semantics*, *Phonology*, etc.).

Additionally, lexical items may carry a focal feature F, corresponding to the set of focus alternatives. When the subject *Jill* is interpreted with semantic focus in the utterance *Jill baked bread*, the set of focus alternatives (FA) of *Jill* is comprised of the contextually relevant individuals in the discourse. In this case, FA would look something like {*Bill*, *Sandy*, *Molly*…}.

These syntactic features interact with a covert exhaustification operator O. O probes for alternatives of the proposition it combines with, checking the corresponding syntactic features associated with these alternatives in its c-command domain. When a feature is checked in this way, it receives a + value. Otherwise, when a feature is never checked by O, it receives a – value by default.

O comes in several varieties, differentiated by the types of alternatives it targets. The O operators that will be relevant for universal FCIs are given in (7).

(7) *O* types and their target alternatives (Chierchia 2013)

a. O_{ALT}: All alternatives

b. $O_{\sigma A}$: Scalar alternatives

c. O_{DA}: Domain alternatives

d. $O_{Exh\text{-}DA}$: Pre-exhaustified domain alternatives (with O_{ALT} applied to each of the domain alternatives)

The inclusion of an *O* operator in the syntax has powerful consequences for the semantics of an utterance. *O* takes in the prejacent and the target alternatives, and asserts that the prejacent is true, while any alternatives not entailed by the prejacent are false. For example, $O_{\sigma A}$ applied to a proposition containing *some* derives the scalar implicature in (8c), which when conjoined to the prejacent results in the meaning in (8d).

(8) $O_{\sigma A}$ (Mary took some$_{[+\sigma]}$ of the classes)

a. Prejacent: $\exists x \in D[class(x) \wedge took(mary,x)]$

b. $\sigma A = \{\exists x \in D[class(x) \wedge took(mary,x)], \forall x \in D[class(x) \rightarrow took(mary,x)]\}$

c. Scalar Implicature: $\neg\forall x \in D[class(x) \rightarrow took(mary,x)]$

d. Meaning: $\exists x \in D[class(x) \wedge took(mary,x)] \wedge \neg\forall x \in D[class(x) \rightarrow took(mary,x)]$
(Paraphrase: Mary took some but not all of the classes.)

In this way, *O* has roughly the same contribution to the parse as *only* in the utterance *Mary took only some of the classes*. The only difference Chierchia (2013) assumes is that *O* asserts the prejacent, while *only* presupposes it (Horn 1969).

3.1 Deriving the universal meaning

Chierchia differentiates universal FCIs from plain indefinites by giving them obligatorily active alternatives that require a higher alternative-sensitive operator in the syntax. If the alternatives of a universal FCI go unchecked, the tree is not syntactically well-formed, and the derivation crashes. An obligatory $+D$ feature on universal FCIs is enough to predict their basic universal reading. To illustrate how obligatory implicature derives the desired reading, I will work with the following toy example.

Toy example

Utterance: *You may take any class.*[2]

[2] While I restrict the domain to two classes in the toy example, *any* appears to prefer a domain containing three or more entities.

(i) ??You may take any of these two classes.

(ii) You may take any of these three classes.

I take the infelicitousness of (i) to arise from pragmatic competition with *either*, not from the semantics of *any* itself.

Unexhaustified meaning: $\exists x \in D[class(x) \wedge \Diamond take(you,x)]$

Domain: $D = \{S,P\}$ (S: Semantics, P: Phonology)

Abbreviations: $\phi_S = take(you,S)$, $\phi_P = take(you,P)$

I will take *any* to be unfocused in the utterance, and will set aside the interaction between FCI *any* and focus in subsequent discussion.

In order to check the domain alternatives of the FCI, the utterance in our toy example must be parsed with an exhaustification operator that targets these alternatives. O_{DA} could merge in at the root node to satisfy the syntactic requirement. However, the negation of each domain alternative is not logically consistent with the prejacent. For Chierchia, the domain alternatives of the sentence are $\Diamond\phi_S$ and $\Diamond\phi_P$—negating these and conjoining these with the prejacent would amount to the contradictory expression $\exists x \in D[class(x) \wedge \Diamond take(you, x)] \wedge \neg\Diamond\phi_S \wedge \neg\Diamond\phi_P$ (there is a class you may take and each class is forbidden to take). Instead, the operator $O_{Exh\text{-}DA}$ merges in, avoiding this contradiction by first pre-exhaustifying the domain alternatives. It then operates on these pre-exhaustified domain alternatives to derive the FC implicature. The derivation is shown in (9).

(9) $O_{Exh\text{-}DA}$ $(\exists x \in D[class(x) \wedge \Diamond take(you,x)])$
  a. Exh-DA $= \{\Diamond\phi_S \wedge \neg\Diamond\phi_P, \Diamond\phi_P \wedge \neg\Diamond\phi_S\}$
  b. FC Implicature: $\neg(\Diamond\phi_S \wedge \neg\Diamond\phi_P) \wedge \neg(\Diamond\phi_P \wedge \neg\Diamond\phi_S) = \Diamond\phi_S \leftrightarrow \Diamond\phi_P$
  c. Meaning: $\exists x \in D[class(x) \wedge \Diamond take(you,x)] \wedge (\Diamond\phi_S \leftrightarrow \Diamond\phi_P)$
  $= \forall x \in D[class(x) \wedge \Diamond take(you,x)]$
  (Paraphrase: For every class, you may take it.)

The implicature-based accounts I discuss here, Chierchia (2013), Dayal (2013) and Szabolcsi (2019), assume that the universal free choice reading of *any* results from the inclusion of $O_{Exh\text{-}DA}$ to check this obligatory $+D$ feature in certain environments. When it comes to deriving the distribution of FCI *any*, however, these accounts diverge in their approach and subsequent semantic predictions.

3.2 Chierchia (2013): Modal Containment

In Chierchia's original implicature-based account, FCI *any* additionally carries an obligatory $+\sigma$ feature, which causes $O_{\sigma A}$ to merge into the syntax. Applying $O_{\sigma A}$ to the utterance from the same toy example from the previous section results in the scalar implicature in (10).

(10) $O_{\sigma A}(\exists x \in D[class(x) \wedge \Diamond take(you,x)])$
  a. $\sigma A = \{\forall x \in D[class(x) \rightarrow \Diamond take(you,x)]\}$
  b. Scalar implicature: $\neg\forall\, x \in D[class(x) \rightarrow \Diamond take(you,x)]$

This scalar implicature is not consistent with the universal FC reading. When it is conjoined to the prejacent and FC implicature derived in the previous section, the result is contradiction. The two implicatures are thus incompatible with the prejacent.[3]

This implicature clash can be resolved by weakening the modal base of the scalar implicature in a process called Modal Containment, as defined in (11).

(11) **Modal Containment** (Chierchia 2013)
SC ⊂ FC, where FC is the modal base for the FC implicature and SC is the modal base for the scalar implicature.

The FC implicature conjoined with the prejacent ($\forall x \in D[class(x) \rightarrow \Diamond take(you,x)]$) is evaluated with respect to the full modal base, and is true iff $\Diamond\phi_S$ and $\Diamond\phi_P$ are both true (i.e. both classes may be taken). The scalar implicature ($\neg\forall x \in D[class(x) \rightarrow \Diamond take(you,x)]$) is evaluated with respect to a non-empty subset of the full modal base, and is true iff at least one of the domain alternatives $\Diamond\phi_S$ or $\Diamond\phi_P$ is false (i.e. at least one of the classes may not be taken).

Let us walk through the implications of Modal Containment for the full modal base. The FC implicature is true if every class is permitted, and the scalar implicature is true if there is some class that is not permitted. For there to be a non-empty subset of the full modal base that makes the scalar implicature true, there must be at least one accessible world in which you do not take the full set of classes. For example, we can consider a modal base that contains three worlds in which you take both Semantics and Phonology (w_1, w_2, w_3), and one world in which you take only Phonology (w_4). This modal base makes the FC implicature true, since there are accessible worlds in which you take Semantics (w_1, w_2, w_3), and accessible worlds in which you take Phonology (w_1, w_2, w_3, w_4). For the scalar implicature, we can narrow the modal base to w_4. In this subset, the scalar implicature is true because there is indeed some class that is not permitted: Semantics. In order for the computation to proceed successfully in this way, the full FC implicature modal base must contain at least one accessible world like w_4 in which ϕ_S or ϕ_P is false. In other words, at least one of $\neg\Box\phi_S$ and $\neg\Box\phi_P$ must be true in the full modal base. It follows that Modal Containment can only rescue *You may take any class* if the logical expression in (12) is true in the full FC implicature modal base.

[3] If the existential quantifier takes narrow scope, the implicatures conjoined with the prejacent do not result in contradiction. Chierchia (2013) proposes an additional Wide Scope Constraint on universal FCIs, forcing them to take wide scope over clausemate modal operators whenever possible in order to derive the implicature clash.

In the case of *any*-numeral constructions, Chierchia argues that they may take narrow scope with respect to a modal, which allows them to co-occur with necessity modals where unmodified *any* constructions cannot (e.g. **You must take any book* vs. *You must take any two books*). I will set these constructions aside, limiting my discussion to instances in which *any* is predicted to take wide scope with respect to the modal.

(12) Chierchia (2013) predicted meaning for *You may take any class*:
$\Diamond\phi_S \wedge \Diamond\phi_P \wedge (\neg\Box\phi_S \vee \neg\Box\phi_P)$
(evaluated with respect to the FC modal base)

For Chierchia, Modal Containment provides the explanation for the limited distribution of these FCIs. Episodic statements are only evaluated with respect to the actual world, and so are the scalar and FC implicatures. This means that Modal Containment cannot rescue the implicature clash, since neither implicature quantifies over a modal base. In the case of a necessity modal (*#You must take any class*), Modal Containment also fails to resolve this clash. The meaning of *You must take any class* in the toy model evaluates to $\Box\phi_S \wedge \Box\phi_P \wedge (\neg\Box\phi_S \vee \neg\Box\phi_P)$, which is contradictory.

3.3 Dayal (2013): Viability Constraint

Dayal (2013) takes on board that universal FCIs have an obligatory $+D$ feature, but presents an alternative analysis of the distribution of universal FCIs. Rather than assuming an implicature clash followed by Modal Containment, this analysis assumes that universal FCIs are subject to the Viability Constraint.

(13) Viability Constraint (Adapted from Dayal 2013)
[…FCI…] is felicitous iff each exhaustified domain alternative is true w.r.t. to some subset of the modal base.

In the context of our toy example, the exhaustified domain alternatives are $\Diamond\phi_S \wedge \neg\Diamond\phi_P$ (You may only take Semantics) and $\Diamond\phi_P \wedge \neg\Diamond\phi_S$ (You may only take Phonology). The Viability Constraint is only satisfied with a modal base containing some worlds in which you only take Semantics, and some worlds in which you only take Phonology. This is a stronger basic meaning than the one proposed by Chierchia (2013) in (12). For example, given a modal base containing w_1 in which you only take Semantics, and w_2 in which you take both Semantics and Phonology, Chierchia predicts the utterance should be felicitous, but Dayal does not. Only when the modal base expands to include w_3, in which you only take Phonology, is the Viability Constraint met.

Dayal assumes that the Viability Constraint is evaluated at the point in the derivation where the FCI is introduced. This means universal FCIs must be interpreted as taking wide scope with respect to a clausemate modal—otherwise, there is no modal base with which to satisfy the Viability Constraint.

3.4 Szabolcsi (2019): Revised Viability Constraint

Dayal's Viability Constraint on universal FCIs turns out to be too strong. Szabolcsi (2019) points out example (14) as a crucial problem for the Viability Constraint.

(14) Any bishop may meet a bishop.

In this example, due to the symmetric nature of the predicate, no subsets of the modal base can make the exhaustified domain alternatives true. The set of exhaustified domain alternatives contains propositions in which only one bishop may meet a bishop: {Only Bishop A may meet a bishop, Only Bishop B may meet a bishop…}. But in classical one-agent deontic logic, it is impossible for these exhaustified domain alternatives to be true in any modal base—if one bishop meets another, the latter bishop has also met the first. This utterance should always violate the Viability Constraint, so Dayal's account predicts that this sentence should be infelicitous.

Given the acceptability of the example in (14), Szabolcsi (2019) argues that the original Viability Constraint is too restrictive; it must be relaxed to allow for these types of sentences. To this end, Szabolcsi suggests a weaker semantics more along the lines of Chierchia's, formalized as a revised Viability Constraint.[4]

(15) Revised Viability Constraint (Adapted from Szabolcsi 2019)
[…FCI…] is felicitous iff each domain alternative is true w.r.t. some subset of the modal base, and false w.r.t. some subset of the modal base.

This weaker formulation of the Viability Constraint still captures the distribution of universal FCIs, while avoiding the issue with examples like (14). Szabolcsi notes that this constraint can be thought of as a scalar implicature of sorts, since the condition requires that each domain alternative (in our toy example, $\Diamond\phi_S$ and $\Diamond\phi_P$) be true in *some but not all* non-empty subsets of the modal base. I will come back to this observation in my own account in Section 5.

3.5 Comparing truth conditions

Stepping back, Chierchia (2013), Dayal (2013) and Szabolcsi (2019) all predict the basic universal interpretation of universal FCIs, in other words, they predict that the utterance *You may take any class* always entails *For every class, you have permission to take it*. However, Chierchia (2013) and Szabolcsi (2019) predict different truth conditions from Dayal (2013) when one or more of the exhaustified domain alternatives cannot be true in a subset of the modal base. I will focus on three possible scenarios of the toy example in which the predictions differ. The first, which I will refer to as "S or 2", corresponds to a scenario in which you may take Semantics on its own, but if you take Phonology you must take Semantics with it. The second, "P or 2", corresponds to a scenario in which you may take Phonology on its own, but if you take Semantics you must take Phonology with it. Finally, in the third scenario "Only 2", you are only allowed to take both classes together—neither exhaustified domain alternative is true in any subsets of the modal

[4] I adapt the terminology used in Szabolcsi (2019) to be consistent with Chierchia's (2013). In the toy example I use, Szabolcsi would refer to ϕ_S and ϕ_P as the relevant domain alternatives instead.

base. In all of these scenarios, I assume the modal base includes worlds in which you do not take any classes at all.

Given that the modal base in each scenario contains worlds in which you take no classes at all, there is no requirement to take a class. Szabolcsi (2019) predicts the utterance should be true in all three scenarios, since it is permitted but not required to take each class. Chierchia (2013) similarly predicts the utterance should be true in all three scenarios. Again, this is because each of the modal bases contains worlds in which you take no classes at all—the scalar implicature in his derivation is true in a subset of the modal base containing only these worlds. Finally, Dayal (2013) predicts the utterance should be false in each of the three scenarios, because at least one exhaustified domain alternative cannot be true in any subset of the modal base, violating the Viability Constraint.

Since Szabolcsi and Chierchia's truth conditions are compatible with bishop sentences like (3), they predict that universal FCIs are felicitous even when there are no subsets of the modal base in which any of the exhaustified domain alternatives are true (the "Only 2" scenario). But this leaves the question of why non-bishop sentences like *You may take any class* are so readily interpreted as meaning *You may take any class on its own*—the intuition that underpins Dayal's account.

In the next section, I review the effect of pragmatic factors such as world knowledge and discourse context on the interpretation of utterances containing FCI *any*. Following Chierchia and Szabolcsi, I argue that Dayal's reading is not an entailment, but instead the result of a robust implicature. I will call this implicature the *exclusiveness* implicature, borrowing terminology from Menéndez-Benito (2005; 2010). The exclusiveness implicature is more robust than other inferences, which only appear to arise in certain contexts.

4 Implicatures arising from free choice *any*

4.1 Exclusiveness implicature

Dayal's Viability Constraint follows previous work by Menéndez-Benito (2005; 2010), who presents the Canasta scenario in (16) as a crucial data point.

(16) The Canasta Scenario (Menéndez-Benito 2005; 2010)
One of the rules of the card game Canasta is: When a player has two cards that match the top card of the discard pile, she has two options: (i) take all the cards in the discard pile and (ii) take no card from the discard pile (but take the top card of the regular pile instead). Those are her two only options.

Given the rules of Canasta, Menéndez-Benito judges (17) to be unambiguously false.

(17) In Canasta, you can take any of the cards in the discard pile when you have two cards that match its top card.

She argues on the basis of this judgment that FCI *any* carries an exclusiveness requirement, requiring the utterance in (17) to entail that each card in the discard pile may be taken on its own.[5] This is the prediction built into Dayal's original Viability Constraint. However, the strength of this prediction in this particular scenario is not enough to prove entailment—I argue that it is the result of a strong exclusiveness implicature instead. We can't straightforwardly test for implicature by embedding FCI *any* under negation, since *any* in the scope of negation functions as an NPI. While Chierchia treats these two instances of *any* as the same item, Dayal and Menéndez-Benito restrict their predictions to FCI *any* only, making interpretations of sentences with NPI *any* irrelevant for the exclusiveness implicature. Instead, I offer the following dialogue in (18) as evidence against the entailment Menéndez-Benito proposes in the Canasta scenario.[6]

(18) a. A: Is it true that in Canasta, you can take any of the cards in the discard pile when you have two cards that match its top card?
b. B: Technically yes…but if you take one card from the discard pile, you must take all of them.

Here, A asks B whether the utterance containing the universal FCI is true. B may felicitiously respond *technically yes*, which should not be possible if A's statement is truly false. Following this line of reasoning, A's statement is literally true in the Canasta scenario. The slightly simpler scenario in (19) also gets at the same intuition.

(19) Context: The local animal shelter has a strict policy of not separating animals from their siblings. Annie goes to the shelter to adopt a cat, and walks into one room with three cats inside.
a. Annie: Can I take any of these cats?
b. Shelter employee: Technically yes…but they're siblings, so if you take one, you're required to take them all.

Besides these test cases, there are numerous other cases in which the prediction Menéndez-Benito and Dayal make does not go through. As Szabolcsi (2019) points out, sentences containing symmetric predicates are clear counterexamples in which the exclusiveness implicature does not arise. Outside of the bishop example she raises, all similar sentences involving symmetric predicates have the same issue. This can be seen in additional examples (20a), (20b) and, (20c).

5 Menéndez-Benito's full exclusiveness requirement predicts that (17) entails that any combination of cards can be taken from the discard pile (including any set of two cards or more). Here, I focus on the prediction that any one card may be taken on its own, which is later adopted by Dayal (2013).

6 The novel data presented in this paper represent judgments of the author and 3 additional native speakers of American English. The 'technically yes' examples (18) and (19), while judged to be less natural by all respondents, are still acceptable. Crucially, (18b) and (19b) are judged to be true and non-contradictory. Interestingly, all but one of the respondents found (18b) and (19b) to be more acceptable when uttered in a 'joking' or 'sarcastic' tone.

(20) a. Any bishop may gather with other bishops.
$\not\rightsquigarrow$ Bishop A may be the only bishop who gathers with other bishops.
b. Any bishop may converse with another bishop.
$\not\rightsquigarrow$ Bishop A may be the only bishop who converses with other bishops.
c. Any bishop may pair up with another bishop.
$\not\rightsquigarrow$ Bishop A may be the only bishop who pairs up with other bishops.

Outside of strictly semantic considerations, the exclusiveness implicature is also sensitive to world knowledge and discourse context, which Chierchia (2013) raises as a concern for Menéndez-Benito's account. For example, in the context given in (21), the exclusiveness implicature is consistent with a graduate student's expectations about their courseload.

(21) Context: You are a graduate student only required to take one class per semester.
a. You: What classes are we allowed to take in our first year?
b. Classmate: We may take any class.
$\rightsquigarrow$ We may take Semantics without taking another class.

Compare this to the context given in (22), in which the exclusiveness implicature is inconsistent with an undergraduate student's strong expectations about their courseload.

(22) Context: You are an undergraduate student required to take four classes per semester.
a. You: What classes are we allowed to take in our first year?
b. Classmate: We may take any class.
$\not\rightsquigarrow$ We may take Semantics without taking another class.

Here, the undergraduate student's certainty regarding their courseload is enough to block the exclusiveness implicature. In this way, this implicature is sensitive to world knowledge.

Moreover, despite the robustness of the exclusiveness implicature when it arises, it may be canceled after the fact. In (23), the expectation that each cat can be adopted on its own is felicitously denied in (23b).

(23) Context: Annie goes to the shelter to adopt a cat.
a. Shelter employee: You may adopt any cat here…
b. But if you adopt a cat, you also have to adopt its siblings.

To summarize, the exclusiveness implicature is context-sensitive and cancelable, as we would expect from a pragmatic implicature.

4.2 Not every

In her discussion of the example in (24), Dayal (2013: 93–94) argues that Bill is unambiguously permitted to simultaneously read all of the books in the contextually supplied domain.

(24) Context: Out of the blue.
  a. Bill may read any of these books. (Dayal 2013)

I argue that this is not the case—out of the blue, the sentence in (24) leaves open whether reading the full set of books is permitted. This uncertainty may be resolved given a context, allowing *any* to be interpreted as *any but not every*, or even the much stronger interpretation *any but not more than one*, as in (25).

(25) Context: You are a part-time student who is forbidden from taking more than one class in a semester. Your university offers hundreds of classes each semester.
  a. You: What classes are we allowed to take in our first year?
  b. Classmate: We may take any class.
    $\not\rightsquigarrow$ We may take two classes at the same time.
    $\not\rightsquigarrow$ We may take three classes at the same time.
    …
    $\not\rightsquigarrow$ We may take every class at the same time.

In this context, the utterance in (25b) appears to leave open whether the listener may take more than one class at a time. The student has no choice but to fall back on their previous world knowledge, which rules out worlds in which they may take more than one class. In other words, the listener's prior knowledge appears to be the source of the inference. We can generate any number of variations on the scenario in (25), manipulating the number of courses the student may take in a semester and thereby manipulating the final interpretation. If the student may only take two courses in a semester, for example, the interpretation of *any* shifts to be *any but no more and no less than two*.

Another set of interesting examples involves the conjunction of *any* with a universal quantifier, as illustrated in (26).

(26) a. Anyone and everyone is welcome to come.
  b. We welcome any and all questions.

While these uses are somewhat idiomatic (the same examples are less felicitous when the order of the conjuncts in reversed, and *any and all* in particular is a codified legal expression), they also point to the fact that phrases containing FCI *any* do not necessarily entail their counterparts containing a (presumably narrow-scope) plain universal quantifier. Far from being redundant, the conjunction appears to clarify an otherwise ambiguous expression. In (26a), the first conjunct asserts that every individual in the domain is welcome, with the second conjunct adding that there is no limit to the number of individuals who may come. Similarly, in (26b), the second conjunct serves to clarify that there is no limit on the number or type of questions that may be posed.

Conjoining a plain universal quantifier to FCI *any* affects its acceptability in environments that disprefer a narrow-scope universal reading. For example, a restaurant may display a sign reading (27a). The same sentence with a universal quantifier conjoined to *anyone* is infelicitious, as in (27b) (Lucas Champollion p.c.). It is highly unlikely based on world knowledge that the restaurant would simply refuse to serve everyone—this scenario may be felicitously ruled out of the interpretation of (27a). In the case of (27b), the inclusion of the universal quantifier explicitly rules in this scenario, resulting in an infelicitous utterance.

(27) a. We reserve the right to refuse to serve anyone.
b. #We reserve the right to refuse to serve anyone and everyone.

Finally, the implicatures *not two*, *not three*,…*not every*, like the exclusiveness implicature, appear to be cancelable inferences. For example, take the scenario in (28). In this context, the mother has returned with just enough souvenirs for every one of her children to have one. Given this, the listener is most likely to interpret (28a) as only permitting them to take one souvenir.

(28) Context: You are one of three siblings. Your mother returns from a business trip with three souvenirs and presents them to just you first.
a. Your mother: You may have any of these souvenirs…
b. In fact, you may even have all of them.

Assuming the listener is reasonable, the *not two* and *not every* implicatures are robust due to strong prior expectations. However, they can be felicitously denied, as shown in (28b). The listener may be surprised by the sentence in (28b), but it is certainly not a contradiction.

5 A pragmatic account of free choice *any*

In the previous section, I demonstrated that utterances containing FCI *any* can but need not receive stronger interpretations than those predicted by Chierchia and Szabolcsi. The exclusiveness implicature is generally robust, but may be blocked by world knowledge and symmetric predicates of the sort raised by Szabolcsi. On the other hand, out of the blue an utterance like *Bill may read any of these books* appears to leave open whether Bill may take two books, three books, or even all of the books. The number of books that may be read must instead be inferred by the context, which may give rise to implicatures *not two*, *not three*, and so on until *not every*. In this section, I show that the semantics proposed by Szabolcsi and a formal model of Gricean pragmatics is all we need to capture these facts.

Following Chierchia (2013), I will assume that a single utterance may be mapped to more than one semantic meaning due to covert exhaustification operators *O*. An exhaustified parse may be considered a possible semantic meaning so long as it ensures that the syntactic requirements of each lexical item are met, and the resulting semantic derivation does not result in contradiction.

I will assume Szabolcsi's semantics for FCI *any*, which I will show is straightforward to formalize using exhaustification operators. Given her assumptions, one constraint on the parse is the obligatory universal interpretation of FCI *any*, which requires the parse to include $O_{Exh\text{-}DA}$ to act on the active domain alternatives of the FCI. A second constraint is the Revised Viability Constraint, which requires each domain alternative to be true in some but not all subsets of the modal base. As discussed in Section 3.4, Szabolcsi notes that this constraint can be thought of as a scalar implicature of sorts, since the condition requires that each domain alternative be true in *some but not all* worlds. This can be satisfied by the inclusion of $O_{\sigma A}$ to check the scalar alternatives of the possibility modal.

With these constraints in hand, we can derive two relevant licit parses for the toy example discussed in Section 3, repeated below.

Toy example

Utterance: *You may take any class.*

Unexhaustified meaning: $\exists x \in D[class(x) \wedge \Diamond take(you,x)]$

Domain: $D = \{S,P\}$ (S: Semantics, P: Phonology)

Abbreviations: $\phi_S = take(you,S)$, $\phi_P = take(you,P)$

The weakest possible meaning of *You may take any class* in the toy example is parsed with the minimum number of exhaustification operators possible to check the obligatory $+D$ feature and satisfy the Revised Viability Constraint. Utterances can meet Szabolcsi's Revised Viability Constraint with a scalar implicature resulting from exhaustification of the possibility modal. The derivation is illustrated in (29) below.

(29) Parse 1: $O_{Exh\text{-}DA}$ $(\exists x \in D[class(x) \wedge O_{\sigma A}(\Diamond take(you,x))])$

a. $\sigma A = \{\Diamond take(you,x), \Box take(you,x)\}$

b. $Exh\text{-}DA = \{O_{\sigma A}(\Diamond\phi_S) \wedge \neg O_{\sigma A}(\Diamond\phi_P), O_{\sigma A}(\Diamond\phi_P) \wedge \neg O_{\sigma A}(\Diamond\phi_S)\}$

c. FC Implicature: $\neg(O_{\sigma A}(\Diamond\phi_S) \wedge \neg O_{\sigma A}(\Diamond\phi_P)) \wedge \neg(O_{\sigma A}(\Diamond\phi_P) \wedge \neg O_{\sigma A}(\Diamond\phi_S))$
$= O_{\sigma A}(\Diamond\phi_S) \leftrightarrow O_{\sigma A}(\Diamond\phi_P)$

d. Meaning: $\exists x \in D[class(x) \wedge O_{\sigma A}(\Diamond take(you,x))] \wedge [O_{\sigma A}(\Diamond\phi_S) \leftrightarrow O_{\sigma A}(\Diamond\phi_P)]$
$= O_{\sigma A}(\Diamond\phi_S) \wedge O_{\sigma A}(\Diamond\phi_P)$
$= \Diamond\phi_S \wedge \neg\Box\phi_S \wedge \Diamond\phi_P \wedge \neg\Box\phi_P$
(Paraphrase: For every class, you can but don't have to take it.)

The listener may choose to parse the utterance with an additional exhaustification operator that checks a *F* feature on the variable contained in the VP. The parse in (30), adapted from Xiang (2020), asserts the stronger reading that each class can be taken on its own. This satisfies both the original and revised Viability Constraints.

(30) Parse 2: $O_{Exh\text{-}DA}(\exists x \in D[class(x) \wedge O_{\sigma A}(\Diamond O_{ALT}(take(you,x)))])$

a. $ALT = \{\phi_S, \phi_P\}$

b. $\sigma A = \{\Diamond O_{ALT}(take(you,x)), \Box O_{ALT}(take(you,x))\}$

c. $Exh\text{-}DA = \{O_{\sigma A} \Diamond O_{ALT}(\phi_S) \wedge \neg O_{\sigma A} \Diamond O_{ALT}(\phi_P), O_{\sigma A} \Diamond O_{ALT}(\phi_P) \wedge \neg O_{\sigma A} \Diamond O_{ALT}(\phi_S)\}$

d. FC Implicature: $\neg[O_{\sigma A} \Diamond O_{ALT}(\phi_S) \wedge \neg O_{\sigma A} \Diamond O_{ALT}(\phi_P)] \wedge \neg[O_{\sigma A} \Diamond O_{ALT}(\phi_P) \wedge \neg O_{\sigma A} \Diamond O_{ALT}(\phi_S)]$
$= O_{\sigma A} \Diamond O_{ALT}(\phi_S) \leftrightarrow O_{\sigma A} \Diamond O_{ALT}(\phi_P)$

e. Meaning: $\exists x \in D[class(x) \wedge O_{\sigma A}(\Diamond O_{ALT}(take(you,x)))] \wedge$
$[O_{\sigma A}(\Diamond O_{ALT}(\phi_S)) \leftrightarrow O_{\sigma A}(\Diamond O_{ALT}(\phi_P))]$
$= O_{\sigma A}(\Diamond O_{ALT}(\phi_S)) \wedge O_{\sigma A}(\Diamond O_{ALT}(\phi_P))$
$= \Diamond O_{ALT}(\phi_S) \wedge \neg\Box O_{ALT}(\phi_S) \wedge \Diamond O_{ALT}(\phi_P) \wedge \neg\Box O_{ALT}(\phi_P)$
$= \Diamond(\phi_S \wedge \neg\phi_P) \wedge \neg\Box(\phi_S \wedge \neg\phi_P) \wedge \Diamond(\phi_P \wedge \neg\phi_S) \wedge \neg\Box(\phi_P \wedge \neg\phi_S)$
(Paraphrase: For every class, you have the option to take it on its own but you don't have to.)

I rule out narrow-scope parses of the FCI, following Dayal (2013)'s assumption that the Viability Constraint applies at the level in the semantic derivation at which the FCI is introduced. Without a modal within the scope of *any*, the Revised Viability Constraint cannot be satisfied. Further, I will set aside licit parses that focus the content words *you*, *class*, and *take*, since the resulting implicatures only bear on other permissions outside of which classes the listener may take (e.g. You may *take* a class, but you may not *teach* a class). I assume the Question Under Discussion (QUD, Roberts 2012) in the toy example is *Which classes may the listener take?*, so these other implicatures are not immediately relevant for the discourse. I will also set aside parses in which *any* is focused for domain-widening effect (Chierchia 2013), as the toy example does not build in wider domains of classes the listener might consider. This means in the model at hand, there are only two possible parses of *any* that differ in their answers to the QUD: one corresponding to Szabolcsi's reading, and one corresponding to Dayal's.

Upon hearing the utterance *You may take any class*, then, a listener can only be certain that they are permitted but not required to take Semantics (and similarly for Phonology). The listener must rely on factors such as the meaning of the predicate and world knowledge to determine the most likely meaning the speaker intended to communicate. In a context in which Parse 2 is not a priori very unlikely or impossible (e.g. you are a graduate student permitted to take one course per semester), the listener expects a cooperative speaker to intend this parse—it is more specific, and allows the listener to rule out more possible worlds. As for implicatures of the type *not every*, no literal parses of the utterance settle the issue of whether the listener may take both classes at once. The listener instead relies on their prior expectations in determining whether or not the speaker intended to communicate *any but not every*.

We can formalize the behavior of a neo-Gricean listener facing this problem using a probabilistic pragmatic model in the Rational Speech Act (RSA) framework (Frank & Goodman 2012). RSA models assume a conversation to be a cooperative game between speakers and listeners. The two conversational agents share the mutual goal of allowing the listener to correctly determine the state of the world they are in. The speaker and listener reason about each other's behavior, coordinating their actions to select moves that maximize the probability of realizing this goal. Because it models speakers and listeners as cooperative and rational agents, RSA captures many aspects of formal Gricean reasoning. When a pragmatic listener in RSA is provided with an utterance, it considers alternative utterances the speaker could have chosen, and chooses an interpretation accordingly.

In the next section, I build on previous work by Champollion et al. (2019), who derive the reading of free choice disjunction (e.g. *You may take Semantics or Phonology*) by assuming that utterances are ambiguous due to optional insertion of covert exhaustification operators. They construct an RSA model that contains two literal interpretation functions: one that maps each utterance to its unexhaustified parse, and one that maps each utterance to one possible exhaustified parse. The pragmatic listener in this model is aware of these two interpretation functions, and reasons about which interpretation function the speaker most likely intended.

While my approach to free choice is similar in spirit, I will adopt a more intuitive implementation recently proposed by Franke & Bergen (2020), which was unavailable to Champollion et al. at the time of their writing. Rather than constructing fixed interpretation functions, Franke & Bergen suggest a system in which each utterance is independently associated with a number of parses, some of which are strengthened with covert exhaustification operators. In their *Global Intentions* model, the listener reasons about which parse of an individual utterance the speaker most likely intended. Unlike the Champollion et al. model, the listener simply picks out the most likely parse for the utterance, rather than an interpretation function for the entire set of utterances in the model.

In Section 5.1, I present a Global Intentions model of FCI *any*, reviewing the utterances, utterance parses, state space, and model architecture I assume. I then review the predictions of this model in Section 5.2.

5.1 A Rational Speech Act model of free choice *any*

5.1.1 State space

For the toy example I have been considering, I have assumed that the QUD is *What classes may the listener take?* The mutual goal of the speaker and listener in this model will be to settle this QUD—in other words, to have the listener correctly determine their class-related permissions. To this end, we can impose a partition on the set of possible worlds, with each cell corresponding to a different set of class-related permissions. I will refer to each cell of the partition as a *state*.

For presentational purposes, let us begin by defining a minimal set of accessible worlds in our modal base—one in which the listener takes only Semantics (w_S), only Phonology (w_P), both classes together (w_{SP}), and no classes at all (w_0). The worlds accessible from each state in the model are shown in **Table 1**. To begin, the model contains two states in which the basic universal free choice reading of *You may take any class* is false. In the "Only S" state, the listener is only permitted to take Semantics, and in the "Only P" state, the listener is only permitted to take Phonology. In both states, it is also permitted to take no class at all.

	w_0	w_S	w_P	w_{SP}
Only S	1	1	0	0
Only P	1	0	1	0
Only 1	1	1	1	0
Any #	1	1	1	1
S or 2	1	1	0	1
P or 2	1	0	1	1
Only 2	1	0	0	1

Table 1: Accessible worlds for each state.

Outside of these two states, there are five possible states in which the basic universal free choice reading of *You may take any class* is true, which leaves significant room for pragmatic strengthening. In two of these states, the exclusiveness implicature holds, with the listener permitted to take each class on its own. In the "Only 1" state, the listener is allowed to take Semantics and allowed to take Phonology, but not both classes together. The "Only 1" state corresponds to the *not every* reading of *You may take any class*. On the other hand, in the "Any #" state, the listener is allowed to take Semantics and allowed to take Phonology, or both together. Again, in both states, the listener is permitted to take no class at all.

Finally, I include the three states introduced in Section 3.5 as the states of affairs for which previous implicature-based theories of universal FCIs diverge in their predictions: "S or 2", "P or 2", and "Only 2". While they are compatible with the basic universal free choice reading, in these states one or both of the exhaustified subdomain alternatives is false. These are the states ruled out by the exclusiveness implicature. As with the other states, the listener is permitted to take no class at all.

I will set aside all states that no utterance can truthfully refer to, including a state in which the listener is not permitted to take any classes at all. Additionally, assuming the Revised Viability Constraint, the crucial utterance *You may take any class* can only refer to states in which each class is permitted but not required to take. This is a hard semantic constraint that is necessary to explain why a necessity modal cannot on its own license FCI *any*. States in which the listener is required to take some set of classes thus cannot be truthfully described by the crucial utterance. Including them in the model provides no additional information about the pragmatics of this utterance, so I will also set them aside.

5.1.2 Utterance set

In order to communicate a given state to a listener, a speaker in the RSA model may choose from a set of possible messages: the utterance set. Here, I define a minimal set of plausible Gricean pragmatic alternatives for the crucial utterance *You may take any class* in the context of our toy example. I assume the set of pragmatic alternatives are the domain and scalar alternatives of the utterance assumed by Chierchia (2013). This gives us the four utterances given in (31).

(31) Utterance set
May S: *You may take Semantics.*
May P: *You may take Phonology.*
May Any: *You may take any class.*
May Every: *You may take every class.*

In line with my overarching assumptions and a desire for simplicity, this set of utterances corresponds to a neo-Gricean set. This is not a trivial choice (e.g. Cremers et al. 2023). As a reviewer points out, some further candidate utterances that could be considered include *You may take both classes* and *You may only take Semantics*. It is possible to include these utterances within RSA—they are longer and therefore are typically regarded as harder to produce, but this factor can be captured by including a production cost parameter to the model. It is not straightforward, however, to justify the inclusion of certain non-Gricean alternatives over others, as the speaker in principle has an infinite set of possible utterances to choose from. Therefore I take the set of neo-Gricean alternatives to be a starting point, as it provides a minimal set of utterances with which to present the dynamics of the model.

Likewise, it is standard to include a 'null' message that truthfully refers to all states in the model, but as it does not affect the predictions, I will omit it for simplicity. The utterance set could also include utterance alternatives containing *must*, the scalar alternative of *may*: *You must take Semantics*, *You must take Phonology*, *You must take every class* (the unacceptable sentence **You must take any class* would be ruled out). Since the implicatures of interest do not deal with whether taking certain classes is required, I will choose to keep the utterance set minimal. However, *You may take Phonology* also gives rise to the implicature that you are not required to take phonology. This and related implicatures that involve the interplay of *may* and *must* could in principle be captured by a model that includes utterances with *must*, as well as an expanded state space. I will leave an investigation of such an extension for future work.

RSA models associate each utterance with a production cost, or estimated effort required for a speaker to produce that utterance. As the utterances considered here are roughly equivalent in length and complexity, I will assume equal costs for each utterance.

5.1.3 Utterance parses

In the Global Intentions model, each utterance is associated with a set of literal parses. Each parse is represented by the set of states in which it is true, much in the same way that a proposition may be represented by a set of possible worlds in which it is true. To maintain theoretical consistency with Szabolcsi's account, the set of utterance parses is derived using the system of covert exhaustification she adopts from Chierchia. Each unique utterance parse in the toy example is listed in (32)–(35), along with one possible LF it may correspond to. For the scalar alternative *You may take every class*, I assume that there is scopal ambiguity, allowing both wide-scope and narrow-scope parses of the universal quantifier.

(32) Parses of *You may take Semantics*

a. $\Diamond take(you,S)$: {Only S, Only 1, Any #, Only 2, S or 2, P or 2}
Paraphrase: You may take Semantics.

b. $\Diamond O_{ALT}(take(you,S))$: {Only S, Only 1, Any #, S or 2}
Paraphrase: You may only take Semantics.

c. $O_{\sigma A}(\Diamond O_{ALT}(take(you,S)))$: {Only S}
Paraphrase: You only may take Semantics.

(33) Parses of *You may take Phonology*

a. $\Diamond take(you,P)$: {Only P, Only 1, Any #, Only 2, S or 2, P or 2}
Paraphrase: You may take Phonology.

b. $\Diamond O_{ALT}(take(you,P))$: {Only P, Only 1, Any #, P or 2}
Paraphrase: You may only take Phonology.

c. $O_{\sigma A}(\Diamond O_{ALT}(take(you,P)))$: {Only P}
Paraphrase: You only may take Phonology.

(34) Parses of *You may take any class*

a. $O_{Exh\text{-}DA}O_{\sigma A}(\exists x \in D[class(x) \wedge \Diamond take(you,x)])$:
{Only 1, Any #, Only 2, S or 2, P or 2}
Paraphrase: For every class, you may but don't have to take it.

b. $O_{Exh\text{-}DA}(\exists x \in D[class(x) \wedge O_{\sigma A}(\Diamond_{[+\sigma]}\, O_{ALT}(take(you,x))]))$:
{Only 1, Any #}
Paraphrase: For every class, you may but don't have to take it on its own.

(35) Parses of *You may take every class*

a. $\forall x \in D[class(x) \wedge O_{\sigma A}(\Diamond take(you,x))]$:
{Only 1, Any #, Only 2, S or 2, P or 2}
Paraphrase: For every class, you may but don't have to take it.

b. $O_{\sigma A}(\Diamond \forall x \in D[class(x) \wedge take(you,x)])$:
{Any #, Only 2, S or 2, P or 2}
Paraphrase: You may but don't have to take every class at once.

c. $\forall x \in D[class(x) \wedge O_{\sigma A}(\Diamond O_{ALT}(take(you,x)))]$:
{Only 1, Any #}
Paraphrase: For every class, you may but don't have to take it on its own.

d. $O_{\sigma A}(\Diamond O_{\sigma A}(\forall x \in D[class(x) \wedge take(you,x)]))$:
{Only 2}
Paraphrase: You only may take every class at once.

The parses for each utterance in the model are summarized in **Table 2**. The states truthfully described by a given utterance parse are assigned a 1, otherwise they are assigned a 0.

	Only S	Only P	Only 1	Any #	Only 2	S or 2	P or 2
May S: (32a)	1	0	1	1	1	1	1
May S: (32b)	1	0	1	1	0	1	0
May S: (32c)	1	0	0	0	0	0	0
May P: (33a)	0	1	1	1	1	1	1
May P: (33b)	0	1	1	1	0	0	1
May P: (33c)	0	1	0	0	0	0	0
May Any: (34a)	0	0	1	1	1	1	1
May Any: (34b)	0	0	1	1	0	0	0
May Every: (35a)	0	0	1	1	1	1	1
May Every: (35b)	0	0	0	1	1	1	1
May Every: (35c)	0	0	1	1	0	0	0
May Every: (35d)	0	0	0	0	1	0	0

Table 2: Truth conditions for each utterance and parse.

5.1.4 The Global Intentions model

As in all RSA models, the Global Intentions model defines pragmatic reasoning recursively, constructing multiple layers of conversational agents. The first layer, level-0, contains a listener who is provided with an utterance, a single literal parse of the utterance, and a set of possible world states. This level-0 listener (L_0) rules out states in which the utterance parse is false, but is left with unresolved uncertainty among the remaining states. When L_0 a priori considers all states equally likely, it assigns each remaining state equal probability. When L_0 a priori considers one remaining state to be more likely than the rest, it assigns a higher probability to that state. In this way, L_0 is rational but not pragmatic—it only considers a single semantic meaning of an utterance and its own prior expectations in selecting a state.

The probability that L_0 selects state s given an utterance u, utterance parse p (represented by the set of states the utterance can truthfully describe), and interpretation function $[\![\cdot]\!]^p$ is given in (36). For any utterance u, $[\![u]\!]^p$ denotes a function from states to truth values, which returns 1 if $s \in p$, and 0 otherwise.

(36) $$P_{L_0}(s|u,p) \propto P(s) \times [\![u]\!]^p(s)$$

This probability distribution can be conceptualized as the listener's interpretation of an utterance. Given an utterance and utterance parse, a listener narrows the set of possible states according to its literal semantics. This is analogous to the set of possible worlds a proposition communicates, but with each world assigned a probability by the listener.

With level-0 defined, we can now begin adding layers of pragmatic agents that additionally reason about an interlocutor in choosing their actions. A level-1 speaker (S_1) takes themselves to be facing L_0, and chooses both an utterance and an utterance parse to pass to the listener. S_1 behaves according to a utility function U_1, which assigns a utility to each utterance-parse pair given a state. This utility can be conceptualized as a trade-off between two things: the probability that L_0 successfully chooses the correct state given the utterance and parse, and the associated cost of the utterance $C(u)$. The utility of an utterance u and parse p given state s and utterance cost $C(u)$ is formalized by the equation in (37). The model presented here assumes that each utterance has equal cost, so this term will not be crucial.

(37) $$U_1(u,p|s) = log(P_{L_0}(s|u,p)) - C(u)$$

S_1 is more likely to choose utterance-parse pairs with higher utilities, but will sometimes still make sub-optimal choices. The speaker's tendency to choose optimal utterance-parse pairs can be modulated by the speaker optimality parameter $\alpha > 0$, where increasing the value of α increases the probability that the speaker will make one of the most optimal choices. The joint probability that the speaker will choose utterance u and utterance parse p given state s and speaker optimality parameter α is described by the equation in (38). We can focus on the predictions for S_1 utterance choice by marginalizing over parses, as in (39).

(38) $$P_{S_1}(u,p|s) \propto exp(\alpha U_1(u,p|w))$$

(39) $P_{S_1}(u|s) \propto \sum_p P_{S_1}(u,p|s)$

Finally, we introduce a listener that takes themselves to be facing a level-1 speaker. This level-1 listener (L_1) is only provided with an utterance, and must reason about the state the speaker meant to communicate, as well as about the utterance parse the speaker intended. The joint probability that this listener will choose state s and utterance parse p given utterance u is described by the equation in (40). As we did with S_1, we can also marginalize over parses as in (41) to obtain just L_1's inferences about the state of the world.

(40) $P_{L_1}(s,p|u) \propto P(s) \times P_{S_1}(u,p|s)$

(41) $P_{L_1}(s|u) \propto \sum_p P_{L_1}(s,p|u)$

This formalization closely approximates the listener we set out to model. When the level-1 listener hears an utterance, they are unsure of which parse the speaker intended. The listener has no choice but to rely on their prior biases about the most likely state of affairs ($P(s)$), along with their best approximation of the speaker's behavior ($P_{S_1}(u,p|s)$).

5.2 Model results

Given the utterance set, state space, and utterance parses assumed in the previous section, the Global Intentions model captures the behavior of the exclusiveness implicature, as well as the weaker *not every* implicature. In this section, I will discuss model results in the context of each of these implicatures in turn.

5.2.1 Exclusiveness implicature

In our toy example, the crucial utterance (42) is associated with the robust exclusiveness implicature (42a).

(42) You may take any class.
  a. ⇝ You may take any class on its own. (exclusiveness implicature)

In the context of the model, when pragmatic listener L_1 derives the exclusiveness implicature (42a), we can expect higher posterior probability on states "Only 1" and "Any #" over states "Only 2", "S or 2" and "P or 2" that do not guarantee permission to take each class on its own. This result on its own is insufficient to diagnose implicature in the context of our model, however—further examination of the effect of the state prior and S_1 behavior will be needed to ascertain which components of the model are driving the results.

As discussed in Section 4.1, the exclusiveness implicature is highly robust, arising in most contexts except those in which world knowledge rules it out. To test whether the model derives this implicature in a neutral context, we will begin with a look at results given uniform state prior probabilities. This prior setting corresponds to a scenario in which the listener a priori assumes all states are equally likely before hearing the speaker's utterance. As discussed in Section 5.1.2, utterance costs are set equal. Additionally, the speaker optimality parameter α is set to 100 to approximate a listener who believes the speaker almost always makes the most rational moves in the conversation. Given these model settings, posterior distributions for L_1 and S_1 are given in **Tables 3** and **4**. **Table 3** is best read row-by-row to see how L_1's posterior changes according to the observed utterance. For example, the first row corresponds to the posterior distribution over states after hearing May S, or *You may take Semantics*. By contrast, **Table 4** is best read column-by-column, as each column presents S_1's posterior distribution over utterances in a given state. In this table, the first row corresponds to S_1's posterior in the "Only S" state.

	Only S	**Only P**	**Only 1**	**Any #**	**Only 2**	**S or 2**	**P or 2**
May S	0.67	0	~0	~0	~0	0.33	~0
May P	0	0.67	~0	~0	~0	~0	0.33
May Any	0	0	0.50	0.50	~0	~0	~0
May Every	0	0	0.17	0.17	0.33	0.17	0.17

Table 3: L_1 with uniform prior, $\alpha = 100$.

	Only S	**Only P**	**Only 1**	**Any #**	**Only 2**	**S or 2**	**P or 2**
May S	1	0	~0	~0	~0	0.50	~0
May P	0	1	~0	~0	~0	~0	0.50
May Any	0	0	0.50	0.50	~0	~0	~0
May Every	0	0	0.50	0.50	~1	0.50	0.50

Table 4: S_1 with uniform prior, $\alpha = 100$.

Looking first at **Table 3**, we see that an L_1 who has heard May S assigns 67% probability to the "Only S" state and 33% probability to the "S or 2" state. Concretely, this posterior distribution describes a listener who, upon hearing *You may take Semantics*, concludes they may only take Semantics 67% of the time. The other 33% of the time, they conclude they may either take Semantics on its own, or take both classes together. This appears to be an intuitive result—why would the speaker single out Semantics if they meant to communicate a state like "Only 1" in which the permissions are identical for each class? Or a state like "P or 2", in which you are permitted to take Phonology on its own, but not Semantics? We see a parallel outcome for May P as well. We lack the empirical grounds to assess the predictions for these utterances further, but they are a plausible result.

The posterior for May Every is more distributed, with 17% probability on 4 states ("Only 1", "Any #", "S or 2", "P or 2") and 33% on "Only 2". An L_1 in our model that hears May Every, then, rules out the two states in which only one class is permitted. They are left with a good deal of uncertainty with regard to which is the actual state of the world, with the most probable being one in which the only permission is to take both classes together. Again, we lack the evidence with which to assess these results, but they are not obviously unexpected or unintuitive.

Finally, given our crucial utterance May Any, the L_1 posterior places 50% probability on the "Only 1" state and 50% probability on the "Any #" state, with approximately 0% on the remaining states. This means that upon hearing the crucial utterance *You may take any class*, L_1 in this model concludes with near-certainty that they are either in the "Only 1" state or the "Any #" state. In both of these states and no others, the listener is permitted to take each class on its own—this posterior distribution is exactly what we would expect if the listener had derived the exclusiveness implicature.

Tracing these results back to S_1, we can see in **Table 4** that the speaker chooses to utter May Any 50% of the time to communicate the "Only 1" and "Any #" states, and almost never to communicate the other states. And while these tables show the L_1 and S_1 posteriors marginalized over parses, we can further probe to see which utterance parses S_1 chooses. The model predicts that when uttering May Any, S_1 intends its parse (34b), which entails the exclusiveness implicature, almost 100% of the time.

Given that S_1 almost always utters May Any to communicate "Only 1" and "Any #", and L_1 almost always concludes they are in these two states after hearing May Any, we conclude the model does indeed derive the desired exclusiveness implicature given a uniform prior.

The explanation for this result is that the weaker parse (34a) of May Any is never the optimal choice for S_1 to communicate a given state to L_0. The utterance with this weaker parse is true in 5 states: "Only 1", "Any #", "Only 2", "S or 2", and "P or 2". When S_1 is in state "Only 1" or "Any #", the stronger parse (34b) of May Any is more optimal—L_0 has only a 1-in-5 chance of arriving at the correct interpretation with (34a) but a 1-in-2 chance with (34b). As for state "Only 2", parse (35d) of May Every is the only parse that singles out this state, guaranteeing L_0 will arrive at the correct state. Finally, in states "S or 2" and "P or 2", parses of utterances "May S", "May P" and May Every outcompete parse (34a) of May Any yet again. In the case of "S or 2", parse (32b) of May S and parse (35b) of May Every both give L_0 a 1-in-4 chance of selecting the desired state, making these the two optimal parses for S_1 (an analogous explanation applies for "P or 2"). S_1 thus almost never selects the weaker parse (34a) for May Any. As the stronger parse (34b) is an optimal choice for communicating the "Only 1" and "Any #" states, S_1 almost always utters May Any with this parse in mind.

The model also captures the robustness of the exclusiveness implicature across contexts. Champollion et al. (2019) demonstrate that the robustness of an implicature is correlated with its stability with respect to changes in the listener's prior probabilities over states. I extend this insight to the case of implicatures arising from utterances containing FCI *any*.

It takes a substantial increase in prior probability on the "Only 2" state to have any effect on L_1's interpretation of May Any. Setting the prior to 80% with equal probability on the 6 other states, the posterior on "Only 2" is still negligible, and even at 90% prior the posterior is 4%. Increasing the prior probability assigned to the "S or 2" and "P or 2" states has more of an effect. A level-1 listener who assigns 70% prior probability to the "S or 2" state (with 5% prior probability on the 6 remaining states) concludes there is a 2% probability of being in the "S or 2" state upon hearing May Any, as is shown in **Table 5**. This is still rather unlikely, but greater than the near-zero probability predicted in the uniform prior model. These results parallel the empirical generalizations for the exclusiveness implicature—it is strong, but it is defeasible with enough contextual support.

	Only S	Only P	Only 1	Any #	Only 2	S or 2	P or 2
May S	0.13	0	~0	~0	~0	0.87	~0
May P	0	0.50	~0	~0	~0	~0	0.50
May Any	0	0	0.49	0.49	~0	0.02	~0
May Every	0	0	0.06	0.06	0.11	0.78	~0

Table 5: L_1 with 70% listener prior probability on S or 2 state, $\alpha = 100$.

The weaker parse of May Any becomes crucial when the probability of being in a state where the exclusiveness implicature is true approaches or drops to exactly zero. In some cases, the semantics of the predicate itself lead the listener to completely rule out such states. In Szabolcsi's *Any bishop may meet a bishop* example, the symmetry of the verb *meet* ensures that there is no metaphysically possible state of affairs in which one bishop may meet a bishop, without the latter also meeting the former. Classical deontic logic then also ensures that there is also no state of affairs in which one bishop has permission to meet another bishop without another bishop also having permission to meet the former. In other cases, strongly held beliefs about the state of the world lead the listener to completely rule out accessible worlds that make any exhaustified domain alternative true. Undergraduates required to take four classes a semester are not likely to conclude on the basis of the utterance *You may take any class* that they may take a single class on its own.

In these cases, there is only one state that remains: the "Only 2" state. At this point, the pragmatic explanation is simple—the "Only 2" state is the only one possible state of affairs that May Any can refer to, so the listener must conclude they are in that state.

5.2.2 Not every implicature

We now turn to the *not every* implicature, which for our crucial utterance (43) is the implicature (43a).

(43) You may take any class.
  a. ⇝ You may not take every class. (*not every* implicature)

To translate this into the language of our model, when L_1 derives this implicature, there should be decreased posterior probability for states in which it is permitted to take both classes together: "Any #", "Only 2", "S or 2", and "P or 2". Recall that in cases where the priors result in an exclusiveness implicature (as in the uniform prior setting, presented in **Tables 3** and **4**), L_1 almost always concludes that May Any communicates "Only 1" or "Any #". In these cases, an additional *not every* implicature will decrease the probability mass on "Any #" in favor of "Only 1".

In Section 4.2 we argued that, out of the blue, an utterance like (43) leaves open the issue of whether it is permitted to take every class simultaneously. The *not every* implicature is highly context-dependent, and only seems to arise in situations in which the permission to take every class is already unlikely. In a context where all states are equally likely, then, we expect the listener to be uncertain regarding whether the listener also has permission to take both classes at once, as in example (24) in Section 4.2. As we saw in **Table 3**, L_1 concludes that they are just as likely to be in the "Only 1" state as in the "Any #" state. In other words, the listener believes there is a 50–50 chance that taking both classes is not permitted. Thus the model correctly predicts the lack of *not every* implicature in this case.

As in the case of the exclusiveness implicature, we can manipulate the model's state priors to test predictions for the *not every* implicature across contexts. For example, when the listener expects the most likely state of the world is one in which they may take just Semantics or Phonology but not both simultaneously (the "Only 1" state), their prior expectations should facilitate the *not every* reading. I find that the model indeed captures this fact. Given 70% prior probability on the "Only 1" state, L_1 assigns almost 100% posterior probability to the "Only 1" state upon hearing the crucial utterance May Any (see **Table 6**). In this case, the listener is almost certain that *any* is intended to communicate *any but not every*.

	Only S	**Only P**	**Only 1**	**Any #**	**Only 2**	**S or 2**	**P or 2**
May S	~ 1	0	~ 0	~ 0	~ 0	~ 0	~ 0
May P	0	~ 1	~ 0	~ 0	~ 0	~ 0	~ 0
May Any	0	0	~ 1	~ 0	~ 0	~ 0	~ 0
May Every	0	0	0.64	0.09	0.09	0.09	0.09

Table 6: L_1 with 70% listener prior probability on Only 1 state, $\alpha = 100$.

This is a direct consequence of reasoning about S_1's behavior, as we see that the speaker chooses May Any only to communicate the "Only 1" state (see **Table 7**). Given that speaker intent and listener interpretation both align with the *not every* reading, we can conclude that the model derives the desired implicature in this context. More broadly, our results capture the comparative weakness of the *not every* implicature when compared to the exclusiveness implicature. In a neutral context with uniform state priors, the *not every* implicature does not arise at all, with L_1 assigning 50–50 odds to whether taking both classes together is permitted. Only when the context makes this permission unlikely does L_1 prefer the *any but not every* interpretation. This differs greatly from the results for the exclusiveness implicature, which is derived not only with a uniform state prior but even in prior settings where "Only 1" and "Any #" are highly unlikely.

	Only S	**Only P**	**Only 1**	**Any #**	**Only 2**	**S or 2**	**P or 2**
May S	1	0	~0	~0	~0	~0	~0
May P	0	1	~0	~0	~0	~0	~0
May Any	0	0	0.50	~0	~0	~0	~0
May Every	0	0	0.50	~1	~1	~1	~1

Table 7: S_1 with 70% listener prior probability on Only 1 state, $\alpha = 100$.

Given the context-dependent nature of the *not every* implicature, it makes sense to ask whether manipulation of the state prior derives the opposite implicature as well. In a context where the most likely state of the world is one in which it is permitted to take Semantics, Phonology, or both (the "Any #" state), we might expect a listener who hears May Any to conclude that taking both classes is permitted (the reading originally proposed for example (24) in Dayal 2013). As in the previous cases, we can examine the outcome of placing 70% of the prior on the "Any #" state, with 5% equal probability on the 6 remaining states. The L_1 and S_1 posterior distributions given this prior setting are shown in **Tables 8** and **9**. The L_1 posterior shifts significantly, with 93% probability on the "Any #" compared to 50% in the uniform prior model. In other words, with these prior expectations, the listener infers that they are most likely permitted to take both classes simultaneously. Unlike in the case of the *not every* implicature however, the increased posterior probability on "Any #" does not arise from a corresponding increase in the speaker's use of May Any to communicate that state. We can see in **Table 9** that the speaker utters May Any to communicate both "Only 1" and "Any #" in equal proportion, just as in the uniform prior setting. L_1's increased posterior on "Any #" can instead be explained by a combination of the state prior and the speaker's almost categorical preference for the stronger parse (34b). A simple Bayesian conditioning on the truth of (34b) given the state prior results in the L_1 posterior, as shown in Equation (1).

$$
\begin{aligned}
P(\text{Any \#} \mid [\![34b]\!]=1) &= \frac{P([\![34b]\!]=1 \mid \text{Any \#}) \times P(\text{Any\#})}{P([\![34b]\!]=1)} \\
&= \frac{P([\![34b]\!]=1 \mid \text{Any \#}) \times P(\text{Any \#})}{P(\text{Any \#}) + P(\text{Only 1})} \\
&= \frac{1 \times 0.70}{0.70 + 0.05} \\
&\approx 0.93
\end{aligned} \tag{1}
$$

	Only S	Only P	Only 1	Any #	Only 2	S or 2	P or 2
May S	0.67	0	~0	~0	~0	0.33	~0
May P	0	0.67	~0	~0	~0	~0	0.33
May Any	0	0	0.07	0.93	~0	~0	~0
May Every	0	0	~0	0.74	0.11	0.05	0.05

Table 8: L_1 with 70% listener prior probability on Any # state, $\alpha = 100$.

	Only S	Only P	Only 1	Any #	Only 2	S or 2	P or 2
May S	1	0	~0	~0	~0	0.50	~0
May P	0	1	~0	~0	~0	~0	0.50
May Any	0	0	0.50	0.50	~0	~0	~0
May Every	0	0	0.50	0.50	~1	0.50	0.50

Table 9: S_1 with 70% listener prior probability on Any # state, $\alpha = 100$.

Given that this L_1 behavior does not follow the pragmatic speaker S_1's intent, we cannot conclude that the model derives the opposite of the *not every* implicature in the traditional sense of the term. Rather, this is an example in which a listener's guess about the state of the world is skewed heavily by their prior expectations.

Why does increasing the prior probability on the "Only 1" state result in implicature, but not when we do the same for "Any #"? The source of this asymmetry is the May Every utterance, and specifically its parse (35b). When the prior probability on "Only 1" is high, it becomes difficult for S_1 to successfully communicate about the "Any #" state, as almost all parses that are true in the latter are also true in the former. Given these parses, the literal listener L_0 assigns the majority of the weight to "Only 1" due to the state prior, which results in a high rate of communicative failure when in the "Any #" state. Luckily, parse (35b) of May Every is the exception, as it truthfully describes "Any #" but not "Only 1". This makes it the most optimal choice for S_1 in the "Any #" state. The natural effect of this is that S_1 now only ever has reason to utter May Any in the "Only 1" state, deriving the *not every* implicature.

There is no utterance parse for which the opposite situation—true in "Only 1" but not in "Any #"—holds. Increased prior probability on "Any #" thus does not change the set of most optimal utterances, and S_1 is left with parses of May Any and May Every that result in a higher margin of communicative failure than in the uniform prior setting.

While we have successfully derived the *not every* implicature, it is worth pointing out that this derivation is rather brittle. The story breaks down without the crucial parse (35b), and other changes such as the inclusion of additional utterances and parses would likely have a large impact as well. This is in contrast to the exclusiveness implicature, which relies on the stronger parse (34b) of May Any, but is unlikely to be as influenced by other changes to the model.

To sum up, our RSA model of FCI *any* not only derives both the exclusiveness implicature and *not every* implicature, but captures the difference in robustness between these two implicatures as well. We saw that the exclusiveness implicature arises strongly with a uniform state prior, and is very difficult to perturb this result even with a substantial increase of prior probability on states incompatible with the implicature. In contrast, the *not every* implicature does not arise at all given a uniform prior, only coming into play when privileging the "Only 1" state. Finally, we considered model outputs given increased prior probability on the "Any #" state, and found an interesting discrepancy between L_1 and S_1 behavior. While the L_1 results alone suggested an implicature to the effect of an *any and every* reading of May Any, a closer examination revealed that S_1 only intended to communicate this reading half of the time. The L_1 result is best explained as a combination of state priors and the expectation that S_1 almost always selects the strongest parse of May Any.

6 Conclusion

In this paper, I set out to reconcile two apparently conflicting sets of judgments from previous work on universal FCIs. The former set of judgments, first raised by Menéndez-Benito (2005; 2010) and subsequently adopted by Dayal (2013), suggests that *You may read any book* should entail *You may read any one book on its own*—in other words, that there must be subsets of the modal base that make each exhaustified domain alternative true. The latter set of judgments, raised by Szabolcsi (2019), suggests that this prediction is too strong. Examples such as *Any bishop may meet a bishop* rule out the accessible worlds required by Dayal, and yet are perfectly felicitous. Szabolcsi proposes a revision to Dayal's semantics that captures this second set of judgments.

On the face of it, Szabolcsi's account cannot predict the stronger reading of universal FCIs that arises in the first set of examples. However, I demonstrated that a formal pragmatic model provides an explanation, deriving the interpretation of the first set of examples as a robust implicature, which I call the exclusiveness implicature. I assumed that FCI *any* requires a minimum degree of exhaustification corresponding to Szabolcsi's semantics, but also admits further optional exhaustification to Dayal's semantics. The set of licit possible parses are fed into a pragmatic mechanism, through which the listener uses Gricean reasoning and prior knowledge about the most likely state of affairs. Based on these factors, the listener selects the most likely parse and intended interpretation.

I formalized the dynamics of such a system using a game-theoretic model of language understanding in the Rational Speech Act (RSA) framework. I built on a previous account of free choice from Champollion et al. (2019), updating it with technical innovations from Franke & Bergen (2020) to model how a listener may arrive at a range of possible interpretations of FCI *any*. The RSA model predicts the robustness of the exclusiveness implicature observed by Menéndez-Benito and Dayal, and further predicts that it is more stable to changes in discourse context than other possible implicatures arising from utterances containing FCI *any*.

This suggests that FCI *any*, and likely other universal FCIs cross-linguistically, are considerably more pragmatically felicitious when the exclusiveness implicature arises. The robustness of the implicature when it arises is similar to the robustness traditionally associated with entailments. It is only through careful consideration of counterexamples, such as those raised by Chierchia (2013) and Szabolcsi, that we can show that the exclusiveness implicature is indeed defeasible.

This analysis of FCI *any* provides evidence for a more nuanced view of the semantics-pragmatics interface, moving beyond a simple dichotomy between defeasible implicatures and nondefeasible entailment. Given classic semantic ambiguities such as scopal ambiguity among quantifiers and other operators, as well as more recent theories that incorporate optional silent exhaustification operators, it is clear that there can be more than a single literal parse for each utterance. In theories of pragmatics involving exhaustification operators, the resolution of this ambiguity is often cast as a set of economy conditions on operator insertion (e.g. Fox 2007; Chierchia 2013). However, following Champollion et al. (2019) and Franke & Bergen (2020), my account shows that we can use independently motivated principles of Gricean reasoning to resolve this ambiguity. Importantly, the semantic and pragmatic parts of the account are modular. Given a different set of assumptions about constraints on FCI *any* or exhaustification, the RSA machinery can remain as is, and vice versa. This provides a rich framework for testing how different theoretical assumptions in both the semantics and pragmatics may affect the interpretation of sentences containing FCI *any*.

6.1 Open issues

This minimal RSA model of universal FCIs serves as a proof of concept that the Revised Viability Constraint proposed by Szabolcsi (2019) is compatible with a robust exclusiveness implicature arising in the pragmatics. Due to the availability of a stronger literal parse, *You may take any class* is uniquely well-suited for communicating states in which you may take each class on its own. This strong pragmatic advantage leads the listener to conclude that the speaker intends this stronger parse, even in contexts where it is a priori unlikely. Other implicatures, such as *not every*, arise instead from the listener's prior expectations. Following Champollion et al. (2019), I assumed that these readings are thus expected to be easily modulated by the listener's prior probabilities over states.

There remain a number of open questions for further study. First, the acceptability of FCI *any* appears to be further modulated by some measure of the degree of freedom communicated to the listener. In Menéndez-Benito's Canasta scenario, every card must be taken at once. The only possible worlds are i) a world in which no cards are taken from the discard pile, and ii) a world in which every card is taken from the discard pile. In this scenario, the exclusiveness implicature is ruled out given the context. We saw in example (18) that it is possible to fall back on the weaker literal meaning in certain cases, but there seems to be something highly uncooperative about using FCI *any* to communicate an all-or-nothing situation (the "Only 2" state in the RSA model). The model predicts that a listener rules out "Only 2" upon hearing a May Any utterance even when the prior probability on "Only 2" is very high, which tracks with Menéndez-Benito's judgments. On the other hand, given the sentence *Any bishop may meet another bishop*, the exclusiveness implicature should also never be true—and yet, the sentence is perfectly acceptable. The difference between the Canasta and bishop examples likely lies in the fact that bishop sentences still appear to communicate that any pair of bishops may meet on their own, a slightly weaker version of the exclusiveness implicature. Similarly, in the example where an undergraduate student is required to take four courses every semester, *You may take any class* appears to communicate that any set of four classes may be taken on their own. Given that the model only has two classes in the domain, it does not distinguish between a) Canasta examples and b) bishop and class examples.

It would be worthwhile to investigate the model predictions when allowing the domain size to include more than two classes for students to choose from. Besides potentially distinguishing between the Canasta and bishop examples, larger domain sizes may pose an issue for my RSA account of FCI *any*. In such domains, the set of states grows exponentially with domain size, and results obtained for a domain of size 2 may not generalize (See Schreiber & Onea 2021). In particular, as a reviewer points out, with domain sizes of 3 or more, the dynamics between May Every and May Any begin to shift—parse (35d) of May Every no longer distinguishes "Only 1" from all other states compatible with parse (34b) of May Any. As a result, May Any would not be able to single out the "Only 1" state anymore, potentially resulting in different predictions for L_1's interpretation of May Any.

Another issue involves the predictions of the Revised Viability Constraint. This constraint is intended to be a semantic requirement determining the distribution of universal FCIs, requiring each domain alternative to be true in some but not all accessible worlds. However, the meaning required by the Revised Viability Constraint appears to be cancelable. Take this example in (44).[7]

[7] A reviewer points out it is not trivial to assume that the setup of this dialogue (as well as in similar examples (18) and (19)) is equivalent to asking if at least one of the parses for its declarative counterpart is true. While I do not commit to a theory of polarity particles here, the point still stands when replacing *yes* with a declarative statement of the antecedent. For example, the grandmother in (44) could say, *It's technically true that you can do any of these chores—in fact, you must do all of them*. While the wordiness makes this utterance less natural, it still does not result in contradiction.

The grandmother's response not only denies the exclusiveness implicature, but also the *may but not must* scalar implicature required by the Revised Viability Constraint.

(44) Context: This year, Chris decides to spend the summer at his grandmother's house. His first day there, his grandmother pulls him aside to discuss his responsibilities for the summer. She shows him a list of chores she has written down.
   a. Chris: Can I choose any of those chores?
   b. Grandmother: Technically yes—in fact, you must do all of them.

Perhaps this means that only the basic universal free choice inference is a truly obligatory meaning, while the Revised Viability Constraint is a more local constraint on the distribution of universal FCIs. The pragmatic sketch I provide in this paper does not account for the dialogue above, since it assumes that the semantic meaning of utterances containing universal FCIs are constrained by the Revised Viability Constraint.

A reviewer points out that on the present account, a sentence like *You may take any class* represents an interesting counterexample to Magri's (2009) Mismatch Hypothesis, which states that a sentence sounds odd if its strengthened meaning contradicts common knowledge. This hypothesis was motivated by sentences like *#Some Italians come from a warm country*. As the reviewer points out, *You may take any class* remains acceptable even if its strengthened meaning *For each class, you are allowed to take that class by itself* contradicts common knowledge, as for example when it is uttered in the context of a program that requires taking at least two classes. For a thoughtful alternative perspective on the Magri data, see Anvari (2018).

Finally, the presence of utterances with multiple parses raises interesting degrees of freedom for RSA-based models. In my model, the speaker considers each utterance-parse pair as equally likely. Alternatively, a speaker could instead consider each utterance as equally likely, sampling an utterance first and then selecting one from a set of equally weighted corresponding utterance parses. This alternative implementation of the utterance-parse prior privileges utterance-parse pairs from utterances with fewer corresponding parses. May Any receives receive the largest boost given its 2 parses, compared to 3 parses for May S/May P and 4 for May Every. While this implementation produces different numerical results, the main conclusions of this paper are not affected by this choice. However, this implementation is rather intuitive in that the speaker considers utterances as primary and parses as a secondary consideration once the utterance has already been selected.

Acknowledgements

I would like to thank Lucas Champollion, Philippe Schlenker and Yimei Xiang for their helpful feedback on this project. Thank you also to the editor and three anonymous reviewers for their insightful comments and suggestions.

Competing interests

The author has no competing interests to declare.

References

Anvari, Amir. 2018. Logical Integrity. In Maspong, Sireemas & Stefánsdóttir, Brynhildur & Blake, Katherine & Davis, Forrest (eds.), *Proceedings of SALT 28*, 711–726. Cambridge, MA. DOI: https://doi.org/10.3765/salt.v28i0.4419

Carlson, Greg N. 1980. Polarity *any* is existential. *Linguistic Inquiry* 11(4). 799–804.

Champollion, Lucas & Alsop, Anna & Grosu, Ioana. 2019. Free choice disjunction as a rational speech act. In Blake, Katherine & Davis, Forrest & Lamp, Kaelyn & Rhyne, Joseph (eds.), *Proceedings of SALT 29*, 238–257. Los Angeles, CA. DOI: https://doi.org/10.3765/salt.v29i0.4608

Chierchia, Gennaro. 2006. Broaden your views: Implications of domain widening and the "logicality" of language. *Linguistic Inquiry* 37(4). 535–590. DOI: https://doi.org/10.1162/ling.2006.37.4.535

Chierchia, Gennaro. 2013. *Logic in grammar: polarity, free choice, and intervention*. Oxford: Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780199697977.001.0001

Cremers, Alexandre & Wilcox, Ethan G. & Spector, Benjamin. 2023. Exhaustivity and antiexhaustivity in the RSA framework: Testing the effect of prior beliefs. *Cognitive Science* 47(5). e13286. DOI: https://doi.org/10.1111/cogs.13286

Dayal, Veneeta. 2004. The universal force of free choice any. *Linguistic Variation Yearbook* 4. 5–40. DOI: https://doi.org/10.1075/livy.4.02day

Dayal, Veneeta. 2009. Variation in English free choice items. *Universals and Variation: Proceedings of GLOW in Asia VII*, 237–256.

Dayal, Veneeta. 2013. A viability constraint on alternatives for free choice. In Fălăuş, Anamaria (ed.), *Alternatives in semantics* (Palgrave Studies in Pragmatics, Language and Cognition), 88–122. London: Palgrave Macmillan. DOI: https://doi.org/10.1057/9781137317247

Fox, Danny. 2007. Free choice and the theory of scalar implicatures. In Sauerland, Uli & Stateva, Penka (eds.), *Presupposition and implicature in compositional semantics*, 71–120. London: Palgrave Macmillan. DOI: https://doi.org/10.1057/9780230210752_4

Frank, Michael C. & Goodman, Noah D. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084). 998. DOI: https://doi.org/10.1126/science.1218633

Franke, Michael & Bergen, Leon. 2020. Theory-driven statistical modeling for semantics and pragmatics: A case study on grammatically generated implicature readings. *Language* 96(2). e77–e96. DOI: https://doi.org/10.1353/lan.2020.0034

Haspelmath, Martin. 1997. *Indefinite pronouns* (Oxford Studies in Typology and Linguistic Theory). Oxford: Oxford University Press. DOI: https://doi.org/10.1093/oso/9780198235606.001.0001

Horn, Laurence. 1969. A presuppositional analysis of *only* and *even*. In *Proceedings from the annual meeting of the Chicago Linguistic Society 5*, 98–107. Chicago: Chicago Linguistic Society.

Horn, Laurence. 1972. *On the semantics and pragmatics of logical operators in English*. Los Angeles, CA: University of California dissertation.

Kadmon, Nirit & Landman, Fred. 1993. Any. *Linguistics and Philosophy* 16(4). 353–422. DOI: https://doi.org/10.1007/BF00985272

Ladusaw, William A. 1979. *Polarity sensitivity as inherent scope relations*. Austin, TX: University of Texas dissertation.

Magri, Giorgio. 2009. A theory of individual-level predicates based on blind mandatory scalar implicatures. *Natural Language Semantics* 17(3). 245–297. DOI: https://doi.org/10.1007/s11050-009-9042-x

Menéndez-Benito, Paula. 2005. *The grammar of choice*. Amherst, MA: University of Massachusetts dissertation.

Menéndez-Benito, Paula. 2010. On universal Free Choice items. *Natural Language Semantics* 18(1). 33–64. DOI: https://doi.org/10.1007/s11050-009-9050-x

Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5(6). 1–69. DOI: https://doi.org/10.3765/sp.5.6

Schreiber, Alexander & Onea, Edgar. 2021. Are narrow focus exhaustivity inferences Bayesian inferences? *Frontiers in Psychology* 12. 1–19. DOI: https://doi.org/10.3389/fpsyg.2021.677223

Szabolcsi, Anna. 2019. Unconditionals and free choice unified. In Blake, Katherine & Davis, Forrest & Lamp, Kaelyn & Rhyne, Joseph (eds.), *Proceedings of SALT 29*, 320–340. Los Angeles, CA. DOI: https://doi.org/10.3765/salt.v29i0.4616

Xiang, Yimei. 2020. Function alternations of the Mandarin particle *dou*: Distributor, free choice licensor, and 'even'. *Journal of Semantics* 37(2). 171–217. DOI: https://doi.org/10.1093/jos/ffz018