Semantic triviality leads to ungrammaticality through iterated learning

Ciyang Qing & Wataru Uegaki

School of Philosophy, Psychology, and Language Sciences,
University of Edinburgh

c.qing@ed.ac.uk
w.uegaki@ed.ac.uk

20 December, 2024

This research is supported by UKRI Future Leaders Fellowship (MR/V023438/1).

1 *Syntax and Natural Logic*

- Natural language grammar contains **syntax**, a structure-building module that produces an input for semantic interpretation.
- A controversial yet prominent hypothesis: Grammar also contains **'natural logic'**:[1]
  - a device that automatically computes logical properties of structures built by syntax; and
  - filters out those that are logically trivial.
- According to the hypothesis, syntax + natural logic determines the set of grammatical sentences. Call this the Logicality hypothesis.[2]
- Today: outline the empirical motivations for the Logicality hypothesis and suggest an alternative that does away with natural logic.

[1] Fox (2000); Fox and Hackl (2006); Chierchia (2006, 2013); Abrusán (2014)

[2] This term follows Del Pinal (2019).

2 *Meaning-driven ungrammaticality*

*Definiteness Effect in the existential-*there *construction*

- The English existential-*there* construction is compatible with some determiners, but not others.

(1) There is { a / no / *every } smiling cat.

- The existential *there* combined with *every* is always semantically trivial (more specifically tautological) (Barwise and Cooper, 1981).[3]
  - $[\![\text{there be det NP}]\!] \Leftrightarrow [\![\text{det}]\!]([\![\text{NP}]\!])(D)$
    (D: the domain of individuals)
  - $[\![\text{there is no smiling cat}]\!] = 1$ iff $[\![\text{no}]\!]([\![\text{smiling cat}]\!])(D)$
    iff $\{\, x \mid x \text{ is a smiling cat} \,\} \cap D = \emptyset$
  - $[\![\text{there is every smiling cat}]\!] = 1$ iff $[\![\text{every}]\!]([\![\text{smiling cat}]\!])(D)$
    iff $\{\, x \mid x \text{ is a smiling cat} \,\} \subseteq D$

[3] More generally, the distinction between the two classes of determiners correlates with a semantic distinction between **strong** and **weak determiners**:
Determiner D is **positive strong** iff for every model $M = \langle [\![\ ]\!], D_e \rangle$ and every $A \subseteq D_e$, if $[\![D]\!](A)$ is defined, then $[\![D]\!](A)(A) = 1$
D is **negative strong** iff for every model $M = \langle [\![\ ]\!], D_e \rangle$ and every $A \subseteq De$, if $[\![D]\!](A)$ is defined, then $[\![D]\!](A)(A) = 0$
D is **weak** if D is not strong.

Other phenomena hypothesised to involve meaning-driven ungrammaticality

(2) *Exceptives*[4]
  a. **All** students but Sue passed the exam.
  b. ***Few** students but Sue passed the exam.

[4] von Fintel (1993)

(3) *NPI licensing*[5]

a. Sue does**n't** have any eggs.

b. *Sue has any eggs.

(4) *Negative islands in comparatives*[6]

a. Mary is taller than **every** other student is.

b. *Mary is taller than **no** other student is.

(5) *Weak islands in manner questions*[7]

a. How does John **hope** that Peter fixed the car?

b. *How does John **regret** that Peter fixed the car?

(6) *Selectional restrictions of clause-embedding predicates*[8]

a. Xander knows **which** horse won the race.

b. *Xander **thinks** which horse won the race.

[5] Chierchia (2006, 2013)

[6] Gajewski (2002)

[7] Abrusán (2014)

[8] Theiler et al. (2019); Uegaki (2015, 2023)

3 *The Logicality hypothesis*

Architectural assumptions

- Grammar contains syntax as well as natural logic.
- Meaning-driven ungrammaticality arises because natural logic filters out logically trivial sentences.

Issue: Grammatical vs. ungrammatical triviality

- Not all semantically trivial sentences are ungrammatical.

(7) a. Alice is pregnant and is not pregnant.

b. It is raining or it is not raining.

- Natural logic should *rule out* meaning-driven ungrammaticality of the kind mentioned in §2 but *rule in* the acceptable triviality of the kind in (7).
- Intuition:
  - meaning-driven ungrammaticality involves logical triviality that arises purely from the **logical words** in the sentence while
  - the triviality in the acceptable cases like (7) arises due to the meaning of **content words**.

Solution 1: Logical Skeletons (Gajewski)

- The **Logical Skeleton** of a sentence is a representation where each token of a content word is replaced with a new variable.[9]
- A sentence is logically trivial if its logical skeleton is necessarily true or necessarily false under all variable assignments.

[9] Gajewski (2002)

(8) there is every X_1.

(9) x_1 is X_2 and is not X_3.

Solution 2: Rescaling (Del Pinal)

- Optional insertion of the RESCALE operator.[10]
- RESCALE is only applicable to open class items.

(10) (Del Pinal, 2019: 11)
For any open class term P, argument of suitable type x and context c, $\{x|\text{RESCALE}_c(P)(x)\} \subseteq \{x|P(x)\}$.
That is, the meaning modulation is constrained to specialize meanings, where the precise refinement depends on the context parameter c.

(11) a. Alice is pregnant and is not pregnant.
b. **Non-trivial LF**
Alice is pregnant and is not RESCALE$_c$(pregnant)

(12) a. There is every student.
b. **Trivial LF even after applying RESCALE**:
There is every RESCALE$_c$(student).

Upshot

- The logicality hypothesis involves a non-trivial architectural assumption about the interaction between syntax and semantics.
  - The grammar contains a logic with its own deductive system, that acts as a filter for grammatical structures.
- There is no consensus on the exact specification of the natural deductive system (see e.g. Abrusán, 2019; Del Pinal, 2019; Pistoia-Reda and Sauerland, 2021)

4 An alternative: Meaning-driven ungrammaticality through iterated learning

- Instead of assuming that natural logic is a part of the grammar, we assume that the process of **iterated learning** filters out logically trivial sentences in the diachronic process.[11]
- Architectural assumptions of our iterated learning model
  - The language faculty defines the space of possible grammars learners can in principle entertain.
  - **Pragmatic speakers**: Each generation of speakers are pragmatic speakers, who produce utterances that maximize the chance of a listener choosing the intended meaning.
  - **Bayesian syntax learning**: Learners induce (through a Bayesian induction process) a syntax that fits the linguistic input the previous generation produces.
    * We crucially assume that the learners are at first pure syntax learners, who do not take into account the interpretations. (They learn the semantics afterwards.)

[10] Del Pinal (2019)

[11] Iterated learning is developed in the theory of language evolution and has been used to model a variety of phenomena, such as emergence of compositionality (Kirby and Hurford, 2002; Kirby et al., 2008, 2015), emergence of morphological regularity (Smith and Wonnacott, 2010), and word-order harmony (Culbertson and Kirby, 2016).

- The effect of the grammar induction can amplify as the leaning process is iterated through generations.

A toy iterated-learning model of the definiteness effect

- Learners' hypothesis space contains various possible grammars with different combinatorial restrictions:

  G_+: existential *there* is compatible with *a*, *no*, and *every*

  G_{-every}: existential *there* is compatible with *a* and *no* but not with *every*.

  ...

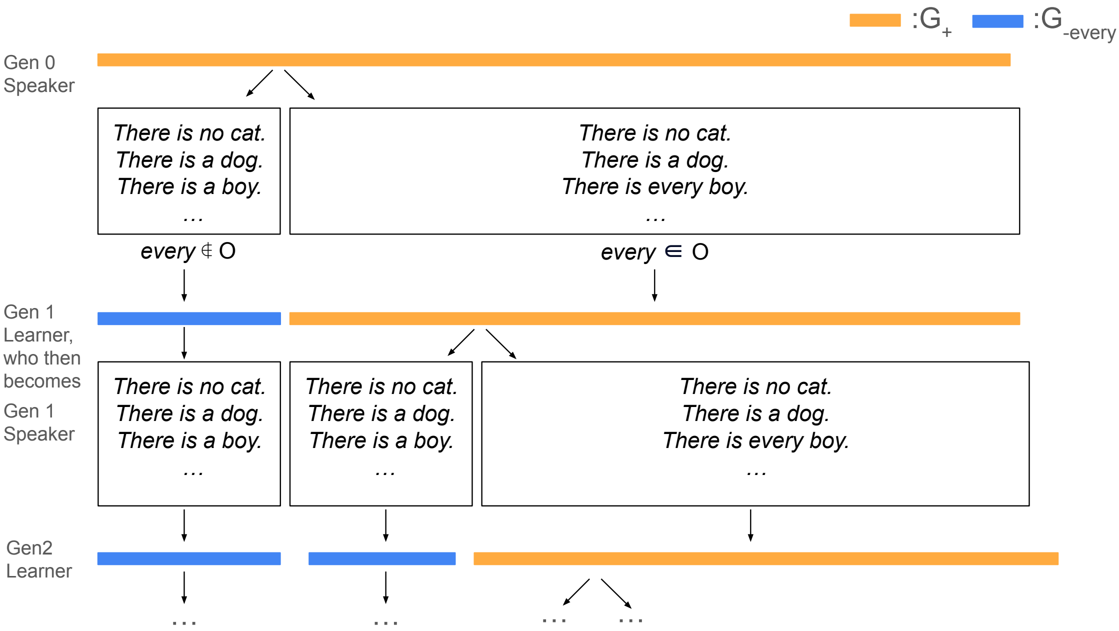

- Speakers of generation 0 possess G_+.
- However, under our assumption that they are pragmatic speakers, they almost never produce *every* in existential *there*, as it would be tautogical and thus informationally useless.[12]
- Thus, some learners in generation 1 never encounter *every* in existential *there* as part of their input data.
- These learners are likely to induce G_{-every} instead of G_+ since the likelihood of observing an input without any occurrence of *every* in existential *there* is significantly greater in G_{-every} than in G_+.[13]
- A motivating analogy: Suppose you observe some natural numbers randomly drawn (with replacement) from a certain set. After observing $4, 8, 2, 2, 6, 12, 6, 18, 4, 16$ (i.e., 10 even numbers and no odd numbers), you will probably conclude that hypothesis

[12] This can be implemented as a pragmatic speaker in the Rational Speech Act (RSA) model (Goodman and Frank, 2016). A pragmatic speaker $S1$ tries (to an extent controlled by the rationality parameter α) to maximize the chance of a literal listener $L0$ choosing the intended world after hearing an existential *there* construction.

[13] A learner uses Bayesian reasoning to infer the most likely grammar G that generates O, the observed uses of existential-*there* from the previous generation:

(i) $L(G \mid O) \propto \Pr(G) \cdot P(O \mid G)$

H_{even} (the set only includes even numbers) is much more likely than H_{all} (the set includes all natural numbers), even though both are technically compatible with your observation.

- Intuitively, this is because the observation is highly unlikely under H_{all} (i.e., $\Pr(O \mid H_{\text{all}}) = (1/2)^{10} = 1/1024$), whereas it is totally expected under H_{even} (i.e., $\Pr(O \mid H_{\text{even}}) = 1$).
- In general, Bayesian learners may take the lack of certain instances in the observation to be evidence that they are not possible.[14]
- The effect of favoring G_{-every} over G_{+} amplifies as the learning is iterated through generations.

[14] Bayesian learning along this line has also been applied to learning word meanings (Xu and Tenenbaum, 2007).

4.1 *Grammatical triviality*

- We still need to account for the contrast between grammatical triviality, as in (7) repeated below, and ungrammatical triviality.

(7) a. Alice is pregnant and is not pregnant.
    b. It is raining or it is not raining.

- Our assumption: the grammars in the learners' hypothesis space do not distinguish between members of an open class such as *pregnant* and *ill* although it does distinguish between different members of the closed class items.[15]
- Consequently, any grammar that can generate non-trivial sentences in the observation such as (13) will also generate (12).

[15] Gajewski's 2009 also discusses an account based on the closed/open class distinction, crucially within the Logicality picture.

(13) a. Alice is pregnant and is not ill.
     b. It is raining or it is not snowing.

4.2 *Nuances*

Diachronic and counterfactual interpretations of meaning-driven ungrammaticality

- We can interpret our account to be making a diachronic claim.
  - There was a point in English grammar when *there is every cat* was grammatical.
  - Through language change, this has become ungrammatical.
- But our account is also compatible with a counterfactual interpretation of the meaning-driven account of relevant ungrammaticality.
  - The English grammar happens to be one that has always ruled out *there is every cat*. (no explanation for why this is the case)
  - If it were to be one that rules in *there is every cat*, the grammar would become one that rules it out after several generations.

The account may be phenomenon-specific

- The account crucially assumes that the target combinatorial restrictions can be accounted for in terms of syntax in the synchronic grammar.
- However, this might not be a reasonable assumption for some of the phenomena discussed in Sect. 2.
- For example, it is broadly agreed in the linguistic community that the distribution of NPIs is semantic/pragmatic in nature.[16]
- This may suggest that our account is applicable to some phenomena (e.g., selectional restrictions of clause-embedding predicates) but not all cases of meaning-driven ungrammaticality.

[16] See for example Kadmon and Landman (1993); Krifka (1994); Chierchia (2013).

5 *Comparison*

- Similarly to the logicality-based accounts which appeal to the difference between logical and content words, our analysis appeals to a difference between open- and closed-class items.
- However, our analysis does not need to postulate an additional natural deductive system sensitive to this distinction to filter out ungrammatical trivial sentences.
- Such sentences end up being *syntactically* ill-formed according to the synchronic grammar.
- This straightfowardly explains why their trivial meanings are non-transparent to non-specialists without semantic analysis and why they intuitively feel ungrammatical.
- In contrast, Logicality-based accounts need to stipulate that their trivial meanings are only accessed subconsciously by the natural deductive system to account for their non-transparency (Chierchia 2013, 53).[17]

[17] Relatedly, the Rescale-based account needs to resort to the possibility of a RESCALE-based interpretation to account for some cases of grammatical triviality, even if it is not the intuitive interpretation of the sentence. Our account does not require computation of such an non-actualized alternative interpretation to account for grammatical triviality. They are simply grammatical according to the synchronic grammar.

6 *Beyond strict triviality*

An additional feature of our analysis: a construction that is not strictly speaking trivial can undergo a similar process and end up ungrammatical.

- As long as the construction is sufficiently uninformative, it will be rarely used and (with other conditions met) learners can induce a grammar that rules it out.

Case study: antirogativity of neg-raising predicates

(14) *Neg-raising predicates cannot take questions*[18]

a. Xander is certain (about) **whether** it is raining.

b. *Xander **believes** whether it is raining.

[18] Theiler et al. (2019); Uegaki (2015, 2023)

- Theiler et al. (2019) propose that neg-raising predicates are antirogative because when they take questions, the assertion is trivial given their Excluded Middle (EM) presupposition (which is assumed to be the source of the neg-raising inference).[19]

[19] See also Zuber's (1982) and Mayr's (2019) analyses along similar (though not identical) lines

(15) Bill doesn't believe that it is raining.
$\rightsquigarrow$ Bill believes that it is not raining.
a. Assertion: $\neg B_b(\textbf{raining})$
b. EM: $B_b(\textbf{raining}) \vee B_b(\neg\textbf{raining})$

(16) Bill believes whether it is raining.
a. Assertion: $B_b(\textbf{raining}) \vee B_b(\neg\textbf{raining})$
(cf. *Bill is certain (about) whether it is raining*.)
b. EM: $B_b(\textbf{raining}) \vee B_b(\neg\textbf{raining})$

- Their analysis relies on the non-trivial assumption that EM is a presupposition and is introduced by a null logical item. But the presuppositional status of EM is questionable.[20]

[20] See e.g. Abusch (2010), Romoli (2013)

- In addition, given that neg-raising is a defeasible inference, one might expect that embedded questions would be fine in contexts where the inference does not arise, but Theiler et al. (2019) observe that this is not the case:

(17) Context: Bill doesn't know who killed Caesar. He isn't even sure whether or not Brutus and Caesar lived at the same time. So, naturally ...
a. Bill doesn't believe that Brutus killed Caesar. $\not\rightsquigarrow$ Bill believes that Brutus didn't kill Caesar.
b. *Bill doesn't believe whether Brutus killed Caesar.

- To account for this, Theiler et al. (2019) crucially assume that the EM presupposition is always present but locally accommodated in such contexts (as opposed to not being triggered at all).

- From our alternative perspective, we can maintain the core insight of their analysis without having to make such strong commitments.
  - We don't need to assume that EM is hard-coded to neg-raising predicates as a (hard) presupposition. All we need is that EM generally holds and therefore it is highly uninformative to assert it.
  - This is enough to cause lack of neg-raisers taking questions in the learners' input.
  - Once neg-raisers become antirogative in the synchronic grammar, they will remain incompatible with questions even in contexts where the neg-raising inference does not arise.

7 Conclusions

- We provide a new analysis of the link between semantic triviality and ungrammaticality, based on two assumptions: Pragmatic speakers and Iterated learning model, based on Bayesian syntax leaning..
  - Pragmatic speakers only rarely utter semantically trivial sentences.
  - Because of this, some learners in the next generation will induce a grammar that rules out such sentences.
  - The effect amplifies through generations.
- The account does away with an additional grammatical module, as is assumed in the Logicality hypothesis.
- It further extends to grammatical restrictions that seem to result from sentences that are systematically highly uninformative rather than strictly trivial.

A *Formal details of the iterated learning model for existential* there

An RSA speaker model We assume that speakers need to use *there is {a / no / every} NP* to communicate one of the 4 possible situations w_i ($i = 0, 1, 2, 3$ representing the cardinality of $[\![\mathrm{NP}]\!]$). A pragmatic speaker $S1$ tries (to an extent controlled by the rationality parameter α) to maximize the chance of a literal listener $L0$ choosing the intended w after hearing an existential-*there* construction (ETC) u:

(18) a. Pragmatic speaker: $S1(u \mid w) \propto [L0(w \mid u)]^{\alpha}$

b. Literal listener: $L0(w \mid u) \propto \Pr(w) \cdot [\![\mathrm{u}]\!](w)$

Assuming the prior $\Pr(w)$ is uniform and $\alpha = 10$, we have the following production probabilities of different ETCs under different situations:

(19) $S1(\textit{no} \mid w_0) = 0.999999$, $S1(a \mid w_0) = 0$, $S1(\textit{every} \mid w_0) = 0.000001$,
$S1(\textit{no} \mid w_i) = 0$, $S1(a \mid w_i) = 0.947$, $S1(\textit{every} \mid w_i) = 0.053$, where $i = 1, 2, 3$

Bayesian grammar learning: Learners in the next generation observe N uses of ETCs generated by the RSA speaker $S1$ for each situation. For concreteness, we take $N = 10$ (the choices of N and α only affect the rate of grammar change, but the qualitative pattern stays the same in the long run). The probability that all 10 instances of ETCs in w_0 use *no* is $0.999999^{10} = 0.99999$, and the probability that all 10 instances of ETCs in w_i use *a* is $0.947^{10} = 0.58$ (for $i \neq 0$). Therefore, the probability that the learner never observes *every* in the 40 uses of ETCs is $0.99999 * 0.58^3 = 0.195$.

A learner in the next generation uses Bayesian reasoning to infer the most likely grammar G that generates O, the observed uses of ETCs from the previous generation: $L(G \mid O) \propto \Pr(G) \cdot P(O \mid G)$. The two relevant candidate grammars are G_+ (ETCs are compatible with all types of DPs) and G_{-every} (ETCs are compatible with *a*-DPs and *no*-DPs but not *every*-DPs). The prior $\Pr(G)$ favors the simpler grammar G_+. The likelihood $P(O \mid G)$ implements learners' assumption that the observations are random samples of grammatical sentences.

For the 19.5% of the learners whose observations O_0 contain no instances of *every*-DPs, given the Bayse rule, we can calculate the ratio between their posterior probabilities of the two candidate grammars as the product of the ratio between the priors and the ratio between the likelihoods:

(20) $\frac{L(G_{-every} \mid O_0)}{L(G_+ \mid O_0)} = \frac{\Pr(G_{-every})}{\Pr(G_+)} \cdot \frac{P(O_0|G_{-every})}{P(O_0|G_+)} \approx \frac{\Pr(G_{-every})}{\Pr(G_+)} * 10^7$

((20) is derived as follows: Assuming that observation contains k instances of *no*, $40 - k$ instances of *a* (and 0 instances of *every*), the ratio between probabilities of encountering this observation under the G_{-every} as opposed to G_+ is the following (code in R):

(21) dmultinom(c(k,40-k,0), prob=c(0.5,0.5,0)) / dmultinom(c(k,40-k,0), prob=c(1/3,1/3,1/3))

No matter what k is, this ratio is 11057332, which is roughly 10^7. A more intuitive way to see this is to instead consider the ratio between seeing no instances of *every* under the two grammars (taking advantage of the fact that the probabilities of *no* and *a* are assumed to be the same). Under G_{-every}, this is guaranteed to be the case, whereas under G_+, the probability is $(2/3)^{40}$. Therefore the ratio is $1/(2/3)^{40} = 11057332$.)

In this case, the ratio between the likelihoods overwhelmingly favors G_{-every}. Therefore, even if there is a decent amount of prior preference for the simpler grammar G_+, essentially all the learners that observe O_0 will end up with G_{-every}. That is, 19.5% of the learners in the next generalization acquire a grammar according to which *every*-DPs are simply incompatible with ETCs. When they become speakers that produce ETCs for the next generalization, they will never use *every*-DPs. Consequently, their descendants will also similarly acquire G_{-every} from the observations.

As for the other 80.5% of the learners, they observe *every*-DPs in ETCs and acquire G_+ (since G_{-every} is incompatible with their input). However, when they become the speakers who produce ETCs for the next generation, 19.5% of their descendants will never observe *every*-DPs and those learners will switch to G_{-every}. That is, the proportion of G_+ in the population will be 19.5% less in each generation. As a result, in the long run virtually all speakers will acquire G_{-every}.

References

Abrusán, Márta. 2014. *Weak island semantics*, vol. 3. OUP Oxford.

Abrusán, Márta. 2019. Semantic Anomaly, Pragmatic Infelicity, and Ungrammaticality. *Annual Review of Linguistics* 5(1):329–351.

Abusch, Dorit. 2010. Presupposition triggering from alternatives. *Journal of Semantics* 27(1):37–80.

Barwise, Jon and Robin Cooper. 1981. Generalized Quantifiers and Natural Language. *Linguistics and Philosophy* 4(2):159–219.

Chierchia, Gennaro. 2006. Broaden your views: Implicatures of domain widening and the "logicality" of language. *Linguistic Inquiry* 37(4):535–590.

Chierchia, Gennaro. 2013. *Logic in Grammar: Polarity, Free Choice, and Intervention*. Oxford: Oxford University Press.

Culbertson, Jennifer and Simon Kirby. 2016. Simplicity and Specificity in Language: Domain-General Biases Have Domain-Specific Effects. *Frontiers in Psychology* 6.

Del Pinal, Guillermo. 2019. The Logicality of Language: A new take on Triviality, "ungrammaticality", and Logical Form. *Noûs* 53(4):785–818.

Fox, Danny. 2000. *Economy and Semantic Interpretation*. Cambridge, MA: The MIT Press.

Fox, Danny and Martin Hackl. 2006. The universal density of measurement. *Linguistics and Philosophy* 29:537–586.

Gajewski, Jon. 2002. *On Analyticity in Natural Language*.

Goodman, Noah D. and Michael C. Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences* 20(11):818–829.

Kadmon, Nirit and Fred Landman. 1993. Any. *Linguistics and Philosophy* 16:353–422.

Kirby, Simon, Hannah Cornish, and Kenny Smith. 2008. Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences* 105(31):10681–10686.

Kirby, Simon and James R Hurford. 2002. The emergence of linguistic structure: An overview of the iterated learning model. *Simulating the evolution of language* 121–147.

Kirby, Simon, Monica Tamariz, Hannah Cornish, and Kenny Smith. 2015. Compression and communication in the cultural evolution of linguistic structure. *Cognition* 141:87–102.

Krifka, Manfred. 1994. The semantics and pragmatics of weak and strong polarity items in assertions. In *Semantics and linguistic theory*, 195–219.

Mayr, Clemens. 2019. Triviality and interrogative embedding: context sensitivity, factivity, and neg-raising. *Natural Language Semantics* 27:227–278.

Pistoia-Reda, Salvatore and Uli Sauerland. 2021. Analyticity and modulation: Broadening the rescale perspective on language logicality. *International Review of Pragmatics* 13(1):1–13.

Romoli, Jacopo. 2013. A scalar implicature-based approach to neg-raising. *Linguistics and Philosophy* 36:291–353.

Smith, Kenny and Elizabeth Wonnacott. 2010. Eliminating unpredictable variation through iterated learning. *Cognition* 116(3):444–449.

Theiler, Nadine, Floris Roelofsen, and Maria Aloni. 2019. Picky predicates: Why believe doesn't like interrogative complements, and other puzzles. *Natural Language Semantics* 27(2):95–134.

Uegaki, Wataru. 2015. *Interpreting Questions under Attitudes*. Ph.D. thesis, Massachusetts Institute of Technology.

Uegaki, Wataru. 2023. *Question-Orientedness and the Semantics of Clausal Complementation*. Springer.

von Fintel, Kai. 1993. Exceptive Constructions. *Natural Language Semantics* 1(2):123–148.

Xu, Fei and Joshua B. Tenenbaum. 2007. Word learning as Bayesian inference. *Psychological Review* 114(2):245–272.

Zuber, Richard. 1982. Semantic restrictions on certain complementizers. In *Proceedings of the 12th International Congress of Linguists, Tokyo*, 434–436.