# Question bias as the pragmatics of context update<sup>\*</sup>

Deniz Rudin drudin@usc.edu

#### draft of December 28, 2024

### INTRODUCTION: CLAUSE TYPING AND CONTEXT UPDATE

There are well-known associations between clause types and context-update profiles. For instance, utterances of declarative sentences typically comprise assertions: the speaker endorses the truth of the sentence; and utterances of interrogative sentences typically comprise askings: the speaker requests that the addressee endorse the truth of an answer. Note that these are only associations, not deterministic absolutes, but as associations they're remarkably robust.

A common explanatory pipeline for deriving these associations goes as follows. (I) Distinctions in syntactic clause typing are mirrored by differences in semantic type: e.g. declarative clauses denote (singleton) propositions; interrogative clauses denote (non-singleton) sets of propositions (Hamblin 1973, cf. Biezma & Rawlins 2012). (II) A general context update rule derives different update effects from different semantic types (e.g. Roberts 1996, Portner 2004, Farkas & Roelofsen 2017, Ciardelli et al. 2018).

To make things concrete, take the account of Farkas & Roelofsen (2017) as an exemplary case. On this account, a declarative denotes a singleton set  $\{p\}$ , whereas the corresponding polar interrogative denotes the set  $\{p, \neg p\}$ . Any utterance commits the speaker to the 'informative content' (i.e., grand union) of the denotation of what they've uttered, and presents each member of that denotation as a potential future common ground for the addressee to choose

<sup>\*</sup>For discussion, thanks to Samir Alam, Pranav Anand, Diti Bhadra, Adrian Brasoveanu, Dylan Bumford, Ivano Caponigro, Emily Clem, Cleo Condoravdi, Chris Cummins, Amy Rose Deal, Karl DeVries, Jack Duff, Regine Eckardt, Donka Farkas, Steven Foley, Sunwoo Jeong, Beste Kamali, Magda Kaufmann, Andy Kehler, Natasha Korotkova, Manfred Krifka, Margaret Kroll, Dan Lassiter, Maura O'Leary, Travis Major, Sophia Malamud, Lisa Matthewson, Roumi Pancheva, Till Poppels, Jessica Rett, Tom Roberts, Floris Roelofsen, Hannah Rohde, Barry Schein, Carson Schütze, Beth Sturman, Wataru Uegaki, Matthijs Westera, Paul Willis, Danfeng Wu, the participants in a Spring 2019 graduate seminar at USC and a Summer 2019 course at Eastern Generative Grammar, and audiences at Santa Cruz, Stanford, Berkeley, Amsterdam, Konstanz, Bochum, MIT, UCLA, UCSD and USC. This paper was written *chez* Birdsall, with crucial logistical support provided by Maki, Miso, and Mogley. Any surviving problems are due entirely to my own personal mindviruses, which are sometimes strong enough to survive inoculation protocols as rigorous as the above.

from. An utterance of a declarative sentence, then, incurs substantive speaker commitment to p, while presenting the addressee with only one choice of a future common ground, one incorporating p. This captures the assertive force of utterances of declarative sentences. And an utterance of a polar interrogative sentence incurs only a trivial commitment to  $\bigcup\{p, \neg p\}$  (= W), while presenting the addressee with two choices of future common grounds, incorporating p or  $\neg p$ . This captures the asking force of utterances of interrogative sentences.

Such a theory needs to explain why not all declarative sentences are assertive, and not all interrogative sentences ask questions, but before we worry about that, let's think about the nature of this pipeline's explanation of the association between declaratives and assertion, and interrogatives and asking. There are two crucial properties of this pipeline that make it explanatorily satisfying. The first property is FORM-DRIVEN INTERPRETATION:

(1) FORM-DRIVEN INTERPRETATION

Context update effects follow compositionally from independently-motivated semantic ramifications of the form of the uttered sentence.

Without a difference in semantic type between declarative and interrogative clauses, this pipeline could not derive a difference in the typical force of their utterances.

The second property is a UNIVOCAL UPDATE PROCEDURE:

(2) UNIVOCAL UPDATE PROCEDURE Context update effects follow from a single mechanism that applies to all sentences, not from construction-specific stipulations.

A semantic type distinction between declaratives and interrogatives would not help explain differences in their force if we were also stipulating construction-specific discourse effects; only with a univocal update mechanism are we actually *deriving* different effects from those different denotations.

As already noted, there are limitations to the association between clause typing and context update, entailing limitations to the feasibility of maintaining those explanatorily satisfying properties of our pipeline. There are many such cases: rhetorical questions, for example, which despite their interrogative form do not ask questions (Rohde 2006, Caponigro & Sprouse 2007, Biezma & Rawlins 2017, Farkas to appear). But this paper is focused on case studies of two types of 'biased questions'.

The first type of 'biased question' that this paper focuses on is rising declaratives. A syntactically declarative sentence with a monotonically rising terminal contour has a non-clause-type-canonical context update of asking a question, but, unlike an ordinary polar interrogative, signals bias about the expected answer (Gunlogson 2001, 2008, Malamud & Stephenson 2015, Farkas & Roelofsen 2017, Jeong 2018, Rudin 2018a, 2022, Goodhue 2024, a.m.o.).

(3) Bizzy can suplex a refrigerator?

In uttering this sentence, the speaker doesn't present themself as though they believe that Bizzy can suplex a refrigerator; rather, they present themself as at the very least not quite sure, and potentially quite skeptical. However, they also signal, very roughly speaking, that there is some evidence that they'll receive an affirmative response. The sentence is natural both as a response to unexpected evidence in favor of the truth of the sentence, and as skeptical pushback to an interlocutor who has just claimed that the sentence is true.

The second type of 'biased question' dealt with here is elliptical tag questions. In these sentences, the form is of a declarative 'anchor' followed immediately by an elliptical interrogative 'tag'.

(4) Bizzy can suplex a refrigerator, can't he?

Because they have the form of a sequence of two sentences of different clause types, it's just not clear what the pipeline sketched above predicts about them. But the empirical fact of the matter is that they also comprise utterances that ask a question, but carry bias about the expected answer in a way quite distinct from ordinary polar interrogatives (Asher & Reese 2007, Krifka 2015, Malamud & Stephenson 2015, Farkas & Roelofsen 2017, Bill & Koev 2023, Scheffler & Malamud 2023, Roberts & Rudin 2024, a.m.o.).

Unlike in the rising declarative case, in uttering this sentence, the speaker does present themself as though they at the very least used to believe that Bizzy can suplex a refrigerator, and yet the effect of the utterance is to ask the addressee whether that is true, rather than to assert that it is. The sentence is natural both as a surprised reaction to unexpected evidence against the truth of the sentence, and in the absence of evidence against it as a way of suggesting that the sentence is true.

Biased questions like the two presented above are commonly taken to force a sacrifice of at least one of the explanatory properties of our pipeline. We could sacrifice FORM-DRIVEN INTERPRETATION, and not assign these sentences the denotations that match their clause types. In the case of rising declaratives, this could take the form of assigning them a non-singleton denotation in spite of their declarative form (e.g. Farkas & Roelofsen 2017, Jeong 2018); in the case of elliptical tag questions, this could take the form of treating the tag as an operator instead of as an interrogative clause (e.g. Farkas & Roelofsen 2017), or by positing covert operators relating the anchor and tag, instead of treating the sentence as the sequence of declarative clause and interrogative clause that it appears on the surface to be (e.g. Krifka 2015, Bill & Koev 2023).

Often, moving away from form-driven interpretation goes hand-in-hand with moving away from a UNIVOCAL UPDATE PROCEDURE. We could treat these cases as requiring construction-specific discourse effects in addition to the basic discourse effects that derive the fundamental association between declaratives and assertion, and interrogatives and asking (e.g. Malamud & Stephenson 2015, Farkas & Roelofsen 2017, Jeong 2018), perhaps justified by these sorts of questions being in some sense 'marked'.

Either maneuver undermines the explanatory power of our pipeline, and weakens our theories' explanation of the crucial, fundamental associations with which we started. But these explanatory sacrifices seem necessary: our hands are tied by the data, and the world can't be as beautiful as our most hopeful theories might suggest. My argument in this paper is that that pessimistic conclusion is made too hastily. Why exactly do these sorts of biased questions force our hand? The first reason is the nonobvious relationship between their clause type and their discourse effect, which motivates abandoning form-driven interpretation. The second reason is the bias inferences that they generate, which don't follow from the machinery we use to model ordinary questions, which motivates moving away from univocal update procedures and toward construction-specific stipulations of expression of bias. I think both issues can be dealt with without making such heavy explanatory sacrifices. We can in fact make headway on the analysis of question bias without sacrificing our theories' most explanatory properties. Two techniques make this possible:

- I. FORMALIZING MEANINGS ON THE LEVEL OF CONTEXT UPDATE, allowing for formdriven specification of non-default update configurations (e.g. Gunlogson 2001, Murray 2014, Faller 2019, Murray & Starr 2020, Farkas 2022, Rudin 2022)
- II. Considering the role of THE PRAGMATICS OF CONTEXT UPDATE in deriving bias inferences from update configurations, allowing for bias to follow from univocal update procedures (e.g. Rudin 2022, Scheffler & Malamud 2023, Goodhue 2024, Farkas to appear)

When it comes to form-driven interpretation, we can retain the role of clause typing while at the same time expanding the space of forms we're considering: there are aspects of linguistic form that operate compositionally on the level of context update, combining with the contributions of clause typing to produce novel update effects without sacrificing compositional form-driven interpretation. When it comes to the univocal update procedure, we can treat context update meanings as modifiers of a univocal update function; contextually-flexible inferences like question bias can be generated by pragmatic reasoning about why the speaker has chosen to update the context in the way that they have, avoiding the need to stipulate question bias as additional construction-specific discourse effects.

The rest of the paper makes that argument through presenting the two case studies as summaries of prior work. In order to make our way there, §1 gives background on default and non-default update configurations, and the pragmatics of context update, and §2 discusses the very idea of meaning on the level of context update. The first case study, a précis of Rudin's (2022) account of rising declaratives, is in §3; the second case study, a précis of Roberts & Rudin's (2024) account of elliptical tag questions, is in §4. §5 gives takeaways and prospects for future directions.

# 1 Context Update: The Details

I assume a 'conversational scoreboard' (Lewis 1979) model of dialogue with two crucial primitive components, in the vein of Farkas & Bruce's (2010) Table model.

- (5) PRIMITIVE COMPONENTS OF THE TABLE MODEL
  - a. INDIVIDUAL DISCOURSE COMMITMENTS For each discourse participant X,  $DC_X$  is the set of propositions that X has

publicly committed to (q.v. Hamblin 1971, Gunlogson 2001, Geurts 2019)

b. The Table

T is a push-down stack of Issues (sets of propositions), the uppermost element of which MAX(T) representing what is currently at issue (q.v. Roberts 1996, Ginzburg 1996)

Context updates in the Table model are defined as functions from contexts to contexts, and typically cashed out in terms of modifications to the primitive components: what they COMMIT the speaker to, and what they PROFFER by placing it on the Table: assertions commit the speaker to a proposition p, while pushing the singleton set containing it  $\{p\}$ onto the Table, proffering a single proposition for discussion. Polar question askings incur no speaker commitments, and push a non-singleton set of propositions  $\{p, \neg p\}$  onto the Table, proffering multiple propositions for discussion.

These updates can be generalized, rather than stipulated as the characteristic update of each clause type. For instance, Farkas & Roelofsen (2017) derive these updates as the characteristic update potential of declarative and interrogative sentences, respectively, using a general-purpose UTT(erance) function:<sup>1</sup>

(6) Let k be the type of a context.

$$\llbracket UTT \rrbracket = \lambda P_{(st)t} \cdot \lambda sp_e \cdot \lambda c_k \cdot \begin{bmatrix} DC_{sp} = DC_{sp,c} + \bigcup P \\ T = T_c + P \\ c' = c \text{ in all other respects} \end{bmatrix}^{c'}$$

This general update recipe works like so: What is pushed onto T is always the denotation of the uttered sentence—a singleton if declarative and a non-singleton if interrogative and the speaker always commits to the 'informative content' of the sentence (the union of propositions in its denotation); this commitment is trivial in the case of a polar question denotation  $\{p, \neg p\}$ .

We can visualize these updates like so:

(7) ASSERTION (= declarative update)



(8) ASKING (= polar interrogative update)

<i>C</i>					$c_1$	
$DC_A$	Table	$DC_B$	$\rightarrow$	$DC_A$	Table	$DC_B$
					$\{p, \neg p\}$	

<sup>&</sup>lt;sup>1</sup>Note that this function assumes that all sentential denotations are proposition-sets; declarative sentences denote singleton sets. Nothing crucial relies on this choice, which was made only for the purposes of simplification. Note also that I overload the interpretation of +, which represents both adding a proposition to a set, and pushing a set of propositions onto a stack

These basic declarative and interrogative updates represent natural output configurations: in the case of the basic declarative update, assertion, the speaker proffers just one proposition, and commits to it; in the case of the basic polar interrogative update, an asking, the speaker proffers multiple propositions, and refrains from (substantive) commitment.

At least, these update configurations *feel* natural. It's natural that the speaker doesn't proffer  $\neg p$  in the assertive case, because it's incompatible with their own commitment to p. And it's natural that the speaker doesn't make a substantive commitment in the asking case, because the propositions they've proffered form a cover of all logical space, ruling out no possible worlds. But these are just informal intuitions. The point of having a formal theory is to not be forced to rely on such informal intuitions. So why *exactly* are these configurations so natural? These informal intuitions rely on tacit assumptions about *why* speakers are making utterances, or under what circumstances making an utterance is *cooperative*. In other words, we need an articulation of the pragmatics of context update to explain these intuitions.

Table-mavens may have already noticed: I've left something out. The model is typically described as including two other, derived, components:

- (9) Derived Components of the Table Model
  - a. Common Ground

CG is the set of all propositions that all discourse participants are publicly committed to  $(= \bigcap \{ DC_X : X \text{ is a discourse participant } \})$ The CONTEXT SET  $CS = \bigcap CG$ 

b. The Projected Set PS is the set of all Common Grounds that could result from making a mutual commitment to an element of MAX(T) (= {  $CG + p : p \in MAX(T)$  })

Though these derived components are not necessary for defining the basic context updates above, they are useful for expressing the underlying pragmatics of the model. These derived components express the view that the goal of discourse, following Stalnaker (1978), is to build common ground, i.e., to establish shared commitments. Content is pushed onto the Table as a step toward achieving this goal: Tabled material is under discussion for potential mutual commitment. PS reflects this by showing the CG that would result from making a mutual commitment to any piece of content on the Table.

The problem with this machinery is that it does not actually formalize the pragmatics that it represents. Exactly what it means for a common ground to be a member of the projected set—under exactly what circumstances it is cooperative to project a common ground remains informal and intuitive. I propose that we dispense with *PS* in our representations, and instead directly formalize the pragmatics that it is intended to represent via Grice-style maxims expressing the conditions under which it is cooperative to include, or fail to include, a proposition in what is placed on the Table (Rudin 2022). These maxims are intended to specify what it means that a proposition placed on the Table is proffered as a candidate future common ground. There will be two such maxims, roughly analogous to Grice's QUALITY and QUANTITY. The first Table-maxim, VIABILITY, specifies the circumstances under which placing a proposition on the Table is cooperative. The second Table-maxim, COMPREHENSIVENESS, specifies the circumstances under which excluding possible worlds from the union of the propositions placed on the Table is cooperative.

- (10) a. VIABILITY (informal version):

  Do not place a proposition on the Table if it is incompatible with some interlocutor's commitment
  Do not place a proposition on the Table if you have reason to believe it is incompatible with some interlocutor's private beliefs
  - b. VIABILITY (formal version):

1) Violated by any MAX(T) that includes a proposition p, where  $\bigcap DC_X \cap p = \emptyset$  for some interlocutor X;

2) Violated by any MAX(T) that includes a proposition p, where  $DOX_{sp}$  entails that  $DOX_X \cap p = \emptyset$  for some interlocutor A

Here's the reasoning for why VIABILITY accurately expresses the tacit common-groundbuilding pragmatics of the Table model. Common Grounds represent mutual public commitments; by the standard assumption of Gricean QUALITY, public commitments (DC) are supposed to represent private beliefs (DOX). Including a proposition in what is placed on the Table puts it forward as a candidate for common grounding. A proposition is a VIABLE addition to the common ground only if it contradicts no interlocutor's individual public commitments—if not, it would not be a possible *shared* commitment—and, to the best of the speaker's knowledge, no interlocutor's private beliefs—if not, making it common ground would pull the conversation's public commitments away from the belief's they're meant to represent.

The set of propositions placed on the Table gives the space of options for common grounding currently under consideration. A countervailing pragmatic pressure is that options for common grounding should be withheld only if they are unviable.

- (11) a. COMPREHENSIVENESS (informal version): The options on the Table should exclude no worlds compatible with the Common Ground, modulo VIABILITY.
  - b. COMPREHENSIVENESS (formal version) Violated by any MAX(T) such that  $\bigcup MAX(T) \not\supseteq CS$ , and  $CS - \bigcup MAX(T)$  is viable

If the set of propositions proffered by an utterance excludes some possible worlds from its union, those worlds will be excluded by any common ground that results from accepting an answer to the QUD the speaker has raised. That's cooperative only if the speaker has reason to believe that those worlds are already ruled out by some interlocutor's public commitments or private beliefs, and therefore could not be legitimately common-grounded.

With this pragmatics in hand, our theory can *explain* why default declarative update and default polar interrogative update are natural update configurations, without reliance on informal intuitions. For simplicity's sake, let's consider a two-by-two grid of possible updates—nothing crucial relies on this, but it'll help us to see how things work. One factor in our grid

is whether the speaker commits to p, or nothing. The other factor is whether the speaker pushes  $\{p\}$  on the Table, or  $\{p, \neg p\}$ .

(12) ASSERTION (= declarative update)

	C					$c_1$	
ĺ	$DC_A$	Table	$DC_B$	$\rightarrow$	$DC_A$	Table	$DC_B$
					p	$\{p\}$	

In the case of an assertive update, the speaker commitments to p while projecting  $\{p\}$ . This is a pragmatically natural move: as long as the speaker has no reason to believe that the addressee believes  $\neg p$ ,<sup>2</sup> pushing p onto the Table is VIABLE. And the speaker's choice to withhold  $\neg p$  from the Table doesn't run afoul of COMPREHENSIVENESS, because their own commitment to p renders  $\neg p$  unviable. No violations!

(13) ASKING (= polar interrogative update)

<i>C</i> <sub>0</sub>				$c_1$			
$DC_A$	Table	$DC_B$	$\rightarrow$	$DC_A$	Table	$DC_B$	
					$\{p, \neg p\}$		

In the case of an asking update, the speaker projects  $\{p, \neg p\}$  and makes no (substantive) commitment. Again, this is a pragmatically natural move: as long as the speaker harbors no secret belief in p or in  $\neg p$ , and has no reason to believe that the addressee believe p, or believes  $\neg p$ ,<sup>3</sup> both options they've placed on the Table are VIABLE, and their union is comprehensive. The speaker's choice to avoid making a commitment to one of the propositions they've placed on the Table is the right one—if they had committed to p, that would've rendered  $\neg p$  unviable. Again, there are no maxim violations, and the assumption of the speaker's cooperativity delivers the fact that ordinary polar questions are unbiased as a pragmatic inference.

But the model doesn't rule out other conceptually possible configurations of commitment and Tabling:

(14) NON-CANONICAL UPDATE 1 (Singleton Table + no commitment)

$c_0$					$c_1$	
$DC_A$	Table	$DC_B$	$\rightarrow$	$DC_A$	Table	$DC_B$
					$\{p\}$	

(15) NON-CANONICAL UPDATE 2 (Non-singleton Table + substantive commitment)

$c_0$				$c_1$			
$DC_A$	Table	$DC_B$	$\rightarrow$	$DC_A$	Table	$DC_B$	
				p	$\{p, \neg p\}$		

These update configurations aren't ruled out; they're just pragmatically marked. In Update 1, the speaker must have reason to treat  $\neg p$  as unviable despite not being committing to p

 $<sup>^{2}</sup>$ See Rudin (2022) §7.2 on how the system deals with cases where this assumption doesn't hold.

<sup>&</sup>lt;sup>3</sup>See Rudin (2022) §7.2 on how the system deals with cases where these assumptions don't hold.

themself. And in Update 2, the speaker has proffered  $\neg p$  in spite of their own commitment to p rendering it unviable. We can visualize the full space we're thinking about like this:

		$\mathbf{commit} \ \mathbf{to} \ p$	don't commit
(16)	$\{p\}$	(default) declarative update	non-canonical update 1
	$\{p, \neg p\}$	non-canonical update 2	(default) polar interrogative update

So: a general theory of context update derives pragmatically natural output configurations as the default declarative and polar interrogative updates. But more pragmatically complex output configurations are also modellable.

In the rest of this paper, I argue that those output configurations are not just modellable in principle, but can be delivered compositionally as the update effects of the two problematic cases discussed in the previous section: rising declaratives (update 1) and elliptical tag questions (update 2). And the pragmatically marked properties of those output configurations deliver the distinctive bias profiles of those two forms as pragmatic inferences about why the speaker has put the context into such a configuration, obviating the need to hard-code bias-signaling into the meanings of the forms themselves. The resulting system preserves both the form-driven interpretation and the univocal update procedure that render theories of the relation between clause typing and context update explanatory.

But before we get to those case studies, a note on the very idea of context-update meanings, and how they're cashed out in English specifically.

# 2 Context update meanings and English intonation

Much recent work explores the hypothesis that some morphemes make their compositional contribution by modifying the context-update potential of sentences they accompany. This has proved especially useful in the analysis of sentence-level particles. For examples, see Cantonese ho2 and me1 (e.g. Law et al. 2024), Mandarin ba (e.g. Ettinger & Malamud 2015, Yang 2020), Japanese yo (e.g. Davis 2009), Romanian oare (e.g. Farkas 2022), German doch and ja (e.g. Kraus 2018), English alas (e.g. Rett 2021), a.m.o.

I assume here a distinction between the meaning of a SENTENCE and the meaning of an UTTERANCE. Sentences denote (sets of) propositions; utterances are functions from a sentential denotation, a speaker, and a context to an output context. In other words, I assume here a static sentential semantics, and a dynamic utterance semantics.

So to analyze a piece of linguistic form 'on the level of context update' is, for me, to claim that it modifies the utterance function, rather than modifying the denotation taken by that function as an input. That is to say, I am going to make the somewhat idiosyncratic terminological choice to say that such morphemes are not part of the sentences they accompany, but rather are part of the utterance of that sentence. Regardless of terminology, this distinction predicts that context-update-level morphemes should not have truth-conditional effects, and predicts that context-update-level morphemes should be unembeddable (except, perhaps, under predicates of communication). Context update function modifiers can be formalized as functions from functions from contexts to contexts to functions from context to contexts (Rudin 2022, Farkas 2022).

(17) CONTEXT UPDATE MODIFIERS (schema) Let K be an abbreviation for type ekk (a function from contexts to contexts).  $[MOD] = \lambda K_K \cdot \lambda sp_e \cdot \lambda c_k \cdot \begin{bmatrix} specification of modification \\ c' = K(sp, c) \text{ in all other respects} \end{bmatrix}^{c'}$ 

This is a fruitful framework for analyzing intonational meaning in English: English intonation makes compositional contributions to context-update potential. I'll note here that there's no reason to think that 'intonational meaning' is a stable category crosslinguistically. I caution against taking the term seriously, for reasons of Anglocentrism. English happens to not be using variation in f0 to do things like specifying lexical items or indicating clause typing, but obviously many languages use f0 for lexical or grammatical tone, and many languages mark neutral polar questions entirely intonationally, e.g. Italian, Brazilian Portuguese, a.m.o.

Even the general correlation between rising intonation and 'questionhood' is less robust than often assumed, e.g. San Martín Peras Mixtec marks polar questions by overall pitch range expansion + final glottalization (Eischens 2021), and even within English, *wh*-questions have falling intonation, and are usually infelicitous with rises (Bartels 1999). So the specific analyses of English intonation here are not meant to be cross-linguistically portable, just as an analysis of English negation makes no predictions about words pronounced *not* in other languages. But comparable meanings may be assigned to other aspects of linguistic form cross-linguistically, as in the cases of (segmental) particle meanings cited above.

### 3 The case of rising declaratives

(N.B.: this section is a précis of Rudin 2022, where many more details can be found.)

A rising declarative (RD) is a syntactically declarative sentence accompanied by a terminal contour rising monotonically from a low pitch accent ( $L^*$  H-H%):

(18) Bizzy can suplex a refrigerator?

Breaking from the typical assertive force of declarative sentences, this sentence asks a question of the addressee.<sup>4</sup> But it also conveys a bias not present in ordinary polar interrogatives, roughly that the speaker has some reason to anticipate a 'yes' answer.

RDs are crosslinguistically attested, in e.g. Dutch, German, modern Hebrew, and Bulgarian, but hardly universal—as mentioned above, many languages use rising intonation to form neutral polar questions, not biased questions.

<sup>&</sup>lt;sup>4</sup>Rising declaratives also have assertive uses, but there is active ongoing debate about whether they comprise the same construction or a different construction from the one discussed in the main text (Westera 2017, Jeong 2018, Goodhue 2024), so I set them aside for the purposes of this paper.

However, there are languages that seem to express identical or comparable meanings with segmental material. Matthewson (2023) argues that the N4e?kepmxcín clitic n contributes the same meaning as English rising declaratives; Ettinger & Malamud (2015), Yang (2020), Yuan (2021) analyze the Mandarin particle ba as contributing a nearby, though not quite identical, meaning; likewise Law et al. (2024) for Cantonese me1. So understanding this construction is of crosslinguistic interest.

The bias profile of RDs is compatible with a range of speaker epistemic biases toward the denotation of the sentence, running from positively biased (19) to negatively biased (20):

- (19) [Context: The speaker has just seen her coworker enter the office wearing a wet raincoat. She says to him:]
   It's raining?
- (20) [Context: George Stephanopoulos is interviewing Donald Trump.]
  DT: I think I've made a lot of sacrifices. I work very, very hard. I've created thousands of jobs, tens of thousands of jobs, built great structures. I've had tremendous success. I think I've done a lot.
  GS: Those are sacrifices?

What is consistent across all uses of RDs is an addressee-oriented bias, as evidenced by their infelicity in a modified version of (20) in which the speaker has no reason to suspect that the addressee is biased in favor of the RD:

(21) DT: I work very, very hard. I've created thousands of jobs, tens of thousands of jobs, built great structures. I've had tremendous success. I think I've done a lot. GS: #Those are sacrifices?

An idea running through much of the literature on RDs is that they are the result of rising intonation modulating the speaker commitment component of context update (Gunlogson 2001, 2008, Truckenbrodt 2006, Malamud & Stephenson 2015, Rudin 2018a,b, 2022, Goodhue 2024, cf. Farkas & Roelofsen 2017). One specific implementation of the way in which rising intonation modulates speaker commitment is that it simply obviates it entirely (Truckenbrodt 2006, Rudin 2018a, 2022).

(22) 
$$[L^* H-H\%] = \lambda K_K \cdot \lambda sp_e \cdot \lambda c_k \cdot \begin{bmatrix} DC_{sp} = DC_{sp,c} \\ c' = K(sp,c) \text{ in all other respects} \end{bmatrix}^c$$

A context update modified with the  $L^*$  H-H% tune will alter the context in the same ways it would without the tune, with the caveat that the speaker's discourse commitments remain unaltered.

Recall the standard declarative update that follows from our general-purpose UTTERANCE function: commitment to p; Tabling of  $\{p\}$ .

(23) DECLARATIVE UPDATE

<i>C</i> <sub>0</sub>					$c_1$	
$DC_A$	Table	$DC_B$	$\rightarrow$	$DC_A$	Table	$DC_B$
				p	$\{p\}$	

The modification expressed by the  $L^*$  H-H% tune results in an update that involves no commitment, just Tabling:

(24) RISING DECLARATIVE UPDATE



In other words, if you combine our general-purpose utterance function with a commitmentobviating account of rising intonation, the result is the first non-canonical update configuration discussed above!

As discussed above, this update configuration is pragmatically marked. The speaker has chosen to exclude a potentially quite large swath of worlds from the options on the Table: the speaker must view  $\neg p$  as unviable, or else it would be uncooperative for them to exclude all  $\neg p$  worlds from the space of possible future common grounds reachable by committing to something they've proffered. In the case of ordinary declarative update, it's clear why they've done so: their own commitment to p renders  $\neg p$  unviable as an addition to the common ground. But in this case, it can't be that the speaker's own beliefs are incompatible with  $\neg p$ ; if they were, they should have committed to p. So it must be that they suspect their *interlocutor's* beliefs to be incompatible with  $\neg p$ .

This pragmatics explains why RDs are compatible with a wide range of speaker biases (falling short of full belief in p), but uniformly associated with addressee-oriented biases. The account delivers question bias without sacrificing either of the desirable properties of our pipeline from clause typing to update effects. It meets the criterion of form-driven composition: being a declarative clause results in having a singleton denotation, and lack of commitment is the compositional contribution of the intonational form of the sentence. And it meets the criterion of a univocal update procedure: no context update effects have been assigned to RDs qua construction, just a general update function interacting compositionally with the contributions of sentence form.<sup>5</sup>

#### 4 The case of elliptical tag questions

(N.B.: this section represents joint work with Tom Roberts, and is a précis of Roberts & Rudin 2024, where many more details can be found.)

Elliptical tag questions (TQs) look like a declarative sentence followed by an interrogative sentence that has undergone VPE:

<sup>&</sup>lt;sup>5</sup>For a fuller exploration of the compositionality of the rising tune meaning across polar interrogatives and wh-interrogatives, see Rudin (2022) §8.

(25) Bizzy can suplex a refrigerator, can't he <suplex a refrigerator>?

This pattern is crosslinguistically attested (Sailor 2009, 2014):

- $\triangleright$  Aux-stranding VPE languages
- (26) TAIWANESE

A-Ying u thak cit-pun che, (i) kam b-o A-Ying perf read one-class book s/he Q neg-perf

"A-Ying read the book, didn't he?"

(27) Danish

han havde slået Udelvik, havde han ikke he have.past slav-part Udelvik have.past he not

"He had slain Udelvik, hadn't he?"

- $\triangleright$  Verb-stranding VPE languages
- (28) IRISH

cheannaigh said teach, nár cheannaigh bought they house neg.Q bought

"They bought a house, didn't they?"

(29) PERSIAN

Naysan ketaab-o khoond, na-khoond Naysan book-obj read neg-read

"Naysan read the book, didn't he?"

The pattern is robust enough that TQs deserve an explanation in terms of their form: being a sequence of a declarative sentence followed by an interrogative sentence. A common intuition in the literature is that TQs somehow combine a declarative update with an interrogative update (q.g. Asher & Reese 2007, Krifka 2015, Bill & Koev 2023, Scheffler & Malamud 2023, cf. Farkas & Roelofsen 2017). A novel idea for how to think about this compositionally: this combination is legislated by the characteristic *prosody* of tag questions.

TQs are differentiated from ordinary sequences of declarative and interrogative sentences via prosodic incorporation into a single Intonational Phrase (IP). IPs delineate (prosodically) independent utterances and host terminal contours: a nuclear pitch accent (\*), followed by a phrase accent (-) and boundary tone (%), plus optional leading or trailing tones.

(30) Intonation and the Prosodic Hierarchy



IPs can be empirically identified both by intonation (presence of a terminal contour) and by prosody (separated from other IPs by a large juncture—Jun 2022 ex.2). By both diagnostics, TQs occupy a single IP:

(31) Pitch track of *I* met Mary, didn't *I*?, read by the author



TQ intonation falls over the declarative anchor and rises over the interrogative tag, just like in ordinary independent falling declaratives and rising polar interrogatives. However, TQs have a single terminal contour, with a pitch accent only on the anchor, and there is no sizeable juncture between clauses. In other words, TQs are prosodically "packaged" into a single IP: TQs comprise two sentences syntactically, packaged into one utterance prosodically.

Because IPs delineate prosodically independent utterances, we propose that the IP is the unit of context update: UTT applies to IPs. If an IP contains multiple sentences, this is cashed out as context update function composition:

(32)  $\operatorname{UTT}(\operatorname{IP}) = \operatorname{UTT}(S_1) \circ \ldots \circ \operatorname{UTT}(S_n)$ , where  $S_1, \ldots S_n$  is the sequence of sentences contained within IP

That is to say, prosodic packaging composes two updates into a single update, moving to the output context of the update sequence in a single step. (33) a. Declarative update

<i>c</i> <sub>0</sub>					$c_1$	
$DC_A$	Table	$DC_B$	$\rightarrow$	$DC_A$	Table	$DC_B$
				p	$\{p\}$	

b. Subsequent polar interrogative update

$c_0$				$c_1$			
$DC_A$	Table	$DC_B$	$\rightarrow$	$DC_A$	Table	$DC_B$	
p	$\{p\}$			p	$\{p, \neg p\}$		

(34) Declarative Update  $\circ$  Polar Interrogative Update (= TQ update)

C_				$c_1$			
$DC_A$	Table	$DC_B$	$\rightarrow$	$DC_A$	Table	$DC_B$	
				p	$\{p, \neg p\}$		

In other words: analyzing prosodic packaging as context update function composition delivers our second non-canonical update configuration from above as the output of a tag question. Again, this update configuration has a non-default pragmatics: speaker commitment to p renders the Tabled proposition  $\neg p$  unviable. Either the speaker does judge it viable, in which case their commitment must be weak and potentially rescindable (Scheffler & Malamud 2023), or the speaker's commitment is strong, and there is merely the pretense of proffering  $\neg p$  as a potential mutual commitment.

This predicted ambiguity matches the attested range of interpretations of TQs:

- (35) LEGITIMATE QUESTION, LESS-THAN-FULL COMMITMENT: [Context: the speaker thought that their addressee was in Bucharest, but they appear at the speaker's doorstep in Pasadena] You're in Bucharest, aren't you?
- (36) PRETENSE-QUESTION, FULL COMMITMENT:
   [Context: The addressee has just asked the speaker why they're watching a 40-minute youtube video about the battle of Midway]
   I'm a 40-year-old man, aren't I?

The takeaway is the same as the previous case study: the account delivers question bias without sacrificing either of our desirable properties. It meets the criterion of form-driven composition, because TQ updates are derived by composing an ordinary declarative update with an ordinary interrogative update, just as TQs appear syntactically to be a matrix declarative followed by a matrix interrogative (with VPE); context update composition follows from the prosodic form of the utterance. And it meets the criterion of a univocal update procedure: no construction-specific update effects are involved, just a general update function interacting compositionally with the two sentences packaged together in its scope.

## 5 TAKEAWAYS

This paper has taxonomized one small corner of the space of context-update configurations: a 2x2 grid of committing to p vs making no (substantive) commitment, and proffering  $\{p\}$ or proffering  $\{p, \neg p\}$ . I've argued that all cells in this taxonomy are attested:

		$\mathbf{commit} \ \mathbf{to} \ p$	don't commit
(37)	$\{p\}$	(default) declarative update	rising declarative update
	$\{p, \neg p\}$	tag question update	(default) polar interrogative update

So-called 'default' declarative and polar interrogative updates occupy the pragmatically most natural top left and bottom right cells. Rising declaratives and tag questions instantiate the more pragmatically marked top right and bottom left cells. Just as the 'default' update configurations can be derived as the update effects of ordinary declarative and polar interrogative sentences without construction-specific stipulations, the 'non-default' output configurations are also achievable without sacrificing form-driven interpretation, given an analysis of (English) intonational meaning on the level of context update.

Most crucially to the point of this paper: bias inferences are derivable without sacrificing a univocal update procedure, falling out of the pragmatics of proffering or not proffering content. There is no need to conventionalize bias or stipulate construction-specific update effects. Rather, we can think of question bias as falling out of the *pragmatics* of context update.

Both RDs and TQs combine aspects of 'default' declarative and polar interrogative update, but different ones. RDs involve singleton proffering, like assertions, but lack substantive commitment, like askings. Their bias follows pragmatically downstream from the combination of singleton proffering and lack of commitment to the proffered proposition. TQs involve non-singleton proffering, like askings, but make a substantive commitment to one of the proffered propositions, like assertions. Their bias follows pragmatically downstream from the clash between the speaker commitment to p and proffering of  $\neg p$ .

These differences result in different bias profiles for RDs and TQs. RDs are characterized by highly variable speaker-oriented bias, but uniformly positive addressee-oriented bias; TQs are characterized by uniformly positive speaker-oriented bias, and variable addressee-oriented bias.

These case studies exemplify a program for thinking about question bias: let's try to preserve the explanatory properties of our pipeline from clause typing to context update effects.

- I. We can preserve FORM-DRIVEN INTERPRETATION by relying on CONTEXT-UPDATE-LEVEL MEANINGS
- II. We can preserve UNIVOCAL UPDATE PROCEDURE by formalizing the PRAGMATICS OF CONTEXT UPDATE

The argument that these two types of question bias can be derived in this way is, of course, hardly a proof that all kinds of question bias can. But note that this paper has explored

only a small corner of the space of possible output configurations, and a small corner of the pragmatics.

There is a much larger space of output configurations explored in recent literature. For instance, some utterances may incur additional speaker commitments to propositions other than the one they've proffered (e.g. Rett 2021); some may move the locus of commitment away from the speaker to some third party (e.g. Faller 2019, Pancheva & Rudin 2019); some may distinguish between commitment as source and as dependent (e.g. Gunlogson 2008, Bhadra 2020).

There is also a much larger space of pragmatics to explore. For instance, the role of relevance (Roberts 1996). Some have argued, for instance, that particles may instantiate relevance relations (e.g. Rojas-Esponda 2014), or that relevance may be legislated by clause typing (e.g. Kaufmann & Kaufmann 2021).

My hope in presenting these two case studies is to show the value of thinking in this way about question bias in particular, and about context update in general. As soon as we step into a world in which aspects of context update are grammaticalized as the meanings of aspects of linguistic form, we need to be thinking about pragmatics. Because as soon as context update is grammaticalized, we have to contend with the context update/pragmatics interface. By their choice of form, the speaker has carried out a particular kind of update; their interlocutors will then draw inferences about why they chose to update the context in that way, which will lead to utterances licensing inferences that go beyond their grammaticalized meanings, often in robust and patterned ways. This is old hat in the domain of truth-conditional semantics, but is disappointingly underexplored in the domain of context update. Formalizing the pragmatics of context update can allow us to give explanatory derivations of various phenomena, including but hardly limited to bias inferences, as following from pragmatic reasoning about the output configurations the speaker has chosen to put the context into. And that can in turn lead to a cleaner semantic theory by offloading construction-specific effects, including but hardly limited to question bias, into the pragmatics. The more construction-specific effects we propose, the fewer predictions our accounts make about anything other than the constructions they analyze; we're in danger of having a collection of formal descriptions of phenomena instead of having a theory.

In getting concrete about the pragmatics of context update, we are not making our theories more baroque or complicated. That pragmatics is already implicit in the design of our theories of contextual representations, and our analyses of meanings as update instructions. In thinking about whether to formalize the pragmatics of context update, our choice isn't between keeping our theories simple or making them more complex. Our choice is only between formally articulating our theories or allowing them to remain handwavy and implicit. To a serious theoretician, that is no choice at all.

#### References

Asher, Nicholas & Brian Reese. 2007. Intonation and discourse: Biased questions. In S. Ishihara, S. Jannedy & A. Schwarz (eds.), *Interdisciplinary studies on information structure*  08, 1–38. Universitätsverlag Potsdam.

- Bartels, Christine. 1999. The intonation of English statements and questions: A compositional interpretation. Routledge.
- Bhadra, Diti. 2020. The semantics of evidentials in questions. *Journal of Semantics* 37. 367–423.
- Biezma, María & Kyle Rawlins. 2012. Responding to alternative and polar questions. Linguistics and Philosophy 35. 361–406.
- Biezma, María & Kyle Rawlins. 2017. Rhetorical questions: Severing asking from questioning. In Proceedings of SALT 27, .
- Bill, Cory & Todor Koev. 2023. Bias in tag questions. To appear in *Perspectives on Biased Questions*.
- Caponigro, Ivano & Jon Sprouse. 2007. Rhetorical questions as questions. In *Proceedings of* Sinn und Bedeutung 11, .
- Ciardelli, Ivano, Jeroen Groenendijk & Floris Roelofsen. 2018. *Inquisitive semantics*. Oxford University Press.
- Davis, Christopher. 2009. Decisions, dynamics and the Japanese particle yo. Journal of Semantics 26. 329–366.
- Eischens, Ben. 2021. Polar question formation in San Martín Peras Mixtec. In Proceedings of the Workshop on Structure and Constituent in Languages of the Americas 25, .
- Ettinger, Allyson & Sophia Malamud. 2015. Mandarin utterance-final particle ba in the conversational scoreboard. In Eva Csipak & Hedde Zeijlstra (eds.), Proceedings of Sinn und Bedeutung 19, .
- Faller, Martina. 2019. The discourse commitments of illocutionary reportatives. *Semantics* and *Pragmatics* 12(8). 1–46.
- Farkas, Donka. 2022. Non-intrusive questions as a special type of non-canonical questions. Journal of Semantics.
- Farkas, Donka. to appear. Rhetorical questions revisited. In Proceedings of Chicago Linguistic Society 60, .
- Farkas, Donka & Kim Bruce. 2010. On reacting to assertions and polar questions. *Journal* of Semantics 27. 81–118.
- Farkas, Donka & Floris Roelofsen. 2017. Division of labor in the interpretation of declaratives and interrogatives. *Journal of Semantics*.
- Geurts, Bart. 2019. Communication as commitment sharing: speech acts, implicatures, common ground. *Theoretical Linguistics* 45(1-2). 1–30.
- Ginzburg, Jonathan. 1996. Dynamics and the semantics of dialogue. In J. Seligman & Dag Westerståhl (eds.), Language, logic, and computation, vol. 1, 1–16. CSLI.

- Goodhue, Daniel. 2024. Everything that rises must converge: Toward a unified account of inquisitive and assertive rising declaratives. In A. Benz, D. Goodhue, M. Krifka, T. Trinh & K. Yatsushiro (eds.), *Biased questions: Experimental results & theoretical modelling*, Language Science Press.
- Gunlogson, Christine. 2001. True to form: Rising and falling declaratives as questions in English: University of California, Santa Cruz dissertation.
- Gunlogson, Christine. 2008. A question of commitment. *Belgian Journal of Linguistics* 22. 101–136.
- Hamblin, C.L. 1971. Mathematical models of dialogue. *Theoria* 37. 130–155.
- Hamblin, C.L. 1973. Questions in Montague English. Foundations of Language 10(1). 41–53.
- Jeong, Sunwoo. 2018. Intonation and sentence type conventions: Two types of rising declaratives. *Journal of Semantics* 35(2). 305–356.
- Jun, Sun-Ah. 2022. The ToBI transcription system: Conventions, strengths, and challenges. In Jonathan Barnes & Stefanie Shattuk-Hufnagel (eds.), *Prosodic theory and practice*, 151–181. The MIT Press.
- Kaufmann, Magdalena & Stefan Kaufmann. 2021. Iffy endorsements. *Journal of Semantics* .
- Kraus, Kelsey. 2018. *Great intonations*: University of California, Santa Cruz dissertation.
- Krifka, Manfred. 2015. Bias in commitment space semantics: Declarative questions, negated questions, and question tags. In *Proceedings of SALT 25*, 328–345.
- Law, Jess, Haoze Li & Diti Bhadra. 2024. Force shift: a case study of Cantonese *ho2* particle clusters. *Natural Language Semantics*.
- Lewis, David. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8. 339–359.
- Malamud, Sophia & Tamina Stephenson. 2015. Three ways to avoid commitments: Declarative force modifiers in the conversational scoreboard. *Journal of Semantics* 32. 275–311.
- Matthewson, Lisa. 2023. Two types of polar question in N4e?kepmxcín. In D. K. E. Reisinger, Laura Griffin, Gloria Mellesmoen, Sander Nederveen, Julia Schillo & Bailey Trotter (eds.), Papers for the international conference on salish and neighbouring languages 58, UBC Working Papers in Linguistics.
- Murray, Sarah. 2014. Varieties of update. Semantics and Pragmatics 7.
- Murray, Sarah E. & William B. Starr. 2020. The structure of communicative acts. *Linguistics* and *Philosophy*.
- Pancheva, Roumyana & Deniz Rudin. 2019. Speaker discourse roles and the discourse profile of reportative evidentials. In Proceedings of the 22nd Amsterdam Colloquium, 327–336.

- Portner, Paul. 2004. The semantics of imperatives within a theory of clause types. In Robert B. Young (ed.), *Proceedings of SALT 14*, 235–252. Evanston, IL: Northwestern University.
- Rett, Jessica. 2021. The semantics of emotive markers and other illocutionary content. *Journal of Semantics* 38(2). 305–340.
- Roberts, Craige. 1996. Information structure in discourse. In J.H. Yoon & A. Kathol (eds.), OSU working papers in linguistics 49: Papers in semantics, 1–53. Ohio State University.
- Roberts, Tom & Deniz Rudin. 2024. Polarity reversal questions and the semantics of prosodic integration. In *Proceedings of Sinn und Bedeutung 28*, .
- Rohde, Hannah. 2006. Rhetorical questions as redundant interrogatives. San Diego Linguistic Papers 2.
- Rojas-Esponda, Tania. 2014. A discourse model for *überhaupt*. Semantics and Pragmatics.
- Rudin, Deniz. 2018a. *Rising above commitment*: University of California, Santa Cruz dissertation.
- Rudin, Deniz. 2018b. Rising imperatives. In *Proceedings of SALT 28*, 100–119.
- Rudin, Deniz. 2022. Intonational commitments. Journal of Semantics 39(2). 338–383.
- Sailor, Craig. 2009. Tagged for deletion: A typological appraach to VP ellipsis in tag questions: UCLA MA thesis.
- Sailor, Craig. 2014. The variables of VP ellipsis: UCLA dissertation.
- Scheffler, Tatjana & Sophia Malamud. 2023. "won't you?" reverse-polarity question tags in American English as a window into the semantics-pragmatics interface. *Linguistics and Philosophy*.
- Stalnaker, Robert. 1978. Assertion. In Peter Cole (ed.), Syntax and semantics 9, 315–32. Academic Press.
- Truckenbrodt, Hubert. 2006. On the semantic motivation of syntactic verb movement to C in German. *Theoretical Linguistics* 32(3). 257–306.
- Westera, Matthijs. 2017. Exhaustivity and intonation: A unified theory: Institute for Logic, Language and Computation, Universiteit van Amsterdam dissertation.
- Yang, Yu'an. 2020. Represent an issue as open: Mandarin discourse particle ba. In Joseph Rhyne, Kaelyn Lamp, Nicole Dreier & Chloe Kwon (eds.), Proceedings of SALT 30, 523– 541.
- Yuan, Xuetong. 2021. Extracting commitments: rising intonation and ba in Mandarin. In *Proceedings of Sinn und Bedeutung 25*, .