

The negatively biased Mandarin belief verb *yǐwéi*

draft of April 2022, accepted at *Studia Linguistica*

Lelia Glass¹

Abstract. The Mandarin belief verb *yǐwéi* strongly suggests that the belief it embeds is wrong or questionable. Based on original data, I propose that this sense of negative bias stems from a *postsupposition* that the reported belief must not be accepted in the Common Ground following an update with *yǐwéi*. When a belief is reported using a neutral, non-factive verb such as *rènwéi* ‘think,’ it is possible for the content of that belief to become Common Ground if the belief or belief-holder are considered reliable; but the postsupposition of *yǐwéi* explicitly blocks such an inference, giving rise to its negative bias. By heading off a potential inference, this postsupposition is further argued to serve a function common to other proposed postsuppositions in the literature. Zooming out, this investigation illuminates the complex calculations triggered by belief reports in discourse, and the linguistic resources used to guide them.

1. Introduction

As a window into the pragmatic complexity of belief reports in discourse, this paper analyzes the sense of negative bias associated with the Mandarin belief verb *yǐwéi*.

In many contexts, *yǐwéi* strongly suggests that the belief it embeds is false, as in (1). In an out-of-the-blue context (joining an ongoing conversation of acquaintances), (1) by default strongly suggests that the speaker is *not* sick. In the same context, (1) – with the neutral verb *rènwéi* substituted for *yǐwéi* – would leave the issue open.

(1) Sick

Māma yǐwéi wǒ bìng le
Mother yǐwéi I sick ASP

¹This paper has existed in various forms since 2015, and I am deeply grateful to everyone who has helped to improve it along the way. Thanks to Christopher Potts, Cleo Condoravdi, Beth Levin, Daniel Lassiter, James N. Collins, and Sven Lauer for helping me see the narrative of this project, and to Jinyi Chu, Robert Henderson, Floris Roelofsen, Michael Yoshitaka (Mitcho) Erlewine, María Biezma, and Valentine Hacquard for helpful discussions. Thanks also to audiences at California Universities Semantics & Pragmatics, the East Asian Linguistics Workshop at the Stanford Department of East Asian Languages & Cultures, Stanford SemFest, LSA 2016, Semantics and Philosophy in Europe 2018, the University of Konstanz, the National University of Singapore; and to all the reviewers (at *Journal of Semantics*, *Glossa*, and *Studia Linguistica*) who have offered very constructive comments on this project over the years. For Chinese data, I express gratitude to Tom Qitong Cao, Delphine Chang, Jia’ao Chen, Jinyi Chu, Xiaochuang Han, Nanjiang Jiang, Chi-Shuen Vince Lee, Peipei Li, Sylvia Li, Zhaoran Ma, Longlu Qin, Ciyang Qing, Joseph Wu, Rob Xu, Fan Zhang, and Chunhao Zhou.

‘Mom is under the impression that I’m sick.’

(LineDictionary²)

§2 presents the inferences of wrongness or skepticism triggered by *yǐwéi* in different discourse contexts, contrasting it with neutral alternatives such as *rènwéi* ‘think.’ (A long-form Appendix, available at <https://osf.io/pr3vt/>, lays out the elicited data in more detail, fully describing discourse contexts for each sentence.)

To derive the effect of *yǐwéi*, §3 analyzes sentences of the form *x yǐwéi p* using an at-issue meaning of *x believes p* and a projecting, non-at-issue *postsupposition* (Lauer 2009, Brasoveanu 2009) requiring the *output context* to be consistent with not-*p* – the mirror image of a presupposition, which instead places a requirement on the *input* context. This postsupposition is argued to prevent a potential inference arising from the pragmatic calculations involved in belief reports.

Zooming out, belief reports present a coordination problem for speakers and hearers (Karttunen 1973, Heim 1992, Simons 2007, Chemla 2008, Beaver 2010, de Marneffe, Manning, & Potts 2012, Anand & Hacquard 2014, Hunter 2016, Özyıldız 2017, Lauer 2017, Lewis, Hacquard, & Lidz 2017, Tonhauser, Beaver, & Degen 2018). Semantically, a non-factive belief report of the form *x believes p* conveys no information about *p* except that *x* believes it. Pragmatically, however, hearers reason about how the speaker views *p*, how reliable the belief-holder *x* is, and how *p* should be overlaid with the Common Ground. If the belief is implausible, if the belief-holder is unreliable, or if the hearer senses skepticism in the speaker’s choice to report a belief rather than making a more direct claim, then the reported belief will not enter the Common Ground, as in (2).

(2) My daughter thinks she’s a mermaid. [web example]

In contrast, if the belief is plausible, if the belief-holder is informed, and if the speaker seems to cite the belief-holder’s belief in *p* as evidence for it (what Simons 2007 calls an ‘evidential’ use, building on Urmson 1952), then the belief may be taken as true, as in (3) – where the ensuing context shows that the author goes on to take it for granted that the motorcyclist hit something.

(3) The investigators think [the motorcyclist] hit something in the road. They don’t know what he hit. Whatever he hit, he flipped the bike. [web example]

More generally: although *x believes p* is itself silent on whether *p* is true or not, it can convey *p* in combination with other contextual assumptions about the credibility of the belief-holder and the belief (Simons 2007, de Marneffe, Manning, & Potts 2012). On the proposed analysis, this potential inference is what *yǐwéi*’s postsupposition serves to prevent.

By semantically blocking one inference (from *x*’s belief that *p* to *p*) via the postsupposition attributed to *yǐwéi*, a speaker’s choice to use *yǐwéi* over a neutral alternative (such as *rènwéi* ‘think’) triggers further, context-sensitive pragmatic reasoning about why the

²This example sentence comes from the Chinese/English web resource LineDictionary. I have collected judgments from native speakers of Mandarin about how it can be used and what inferences would be drawn from it in different contexts.

speaker made this choice. Depending on who the belief-holder x is and what the speaker knows about the belief p , hearers infer different reasons that the speaker wishes to prevent p from being taken up: perhaps the speaker thinks the belief is false or unreliable, or perhaps the speaker questions the belief-holder's credibility, giving rise to the context-dependent notions of wrongness and skepticism associated with *yǐwéi* (§4).

§5 steps back to situate the proposed analysis of *yǐwéi* among other uses of postsupposition in the literature. §5 argues that *yǐwéi* fits into a larger class of postsuppositions used to semantically block an inference which might otherwise arise from the main assertion, thereby giving rise to pragmatic reasoning about why the speaker chose to do so. The result is a unified, intuitive understanding of several otherwise heterogeneous uses of this device. §6 concludes.

2. Data

Yǐwéi is one of many different Mandarin verbs describing an attitude holder's doxastic relation to an embedded proposition (cf. Anand & Hacquard 2014 for English): *rènwéi* 'think,' *xiǎng* 'believe/want,' *juéde* 'feel that,' *zhīdào* 'know,' *jiǎshè* 'assume,' *wàng* 'forget,' *dāng* 'mistake,' and others³.

In descriptive work on Mandarin, it is observed that *yǐwéi* is often used for false beliefs, while *rènwéi* 'think' is more neutral (Lü 1999)⁴. In research on child language, it has been found that the sense of wrongness associated with *yǐwéi* is so striking that children perform better in false-belief tasks when *yǐwéi* is used instead of a neutral alternative such as *xiǎng* 'believe/want' (Lee, Olson, & Torance 1999).

To pinpoint the source of this negative bias, *yǐwéi* must be further tested in contexts involving different assumptions about the status of the reported belief, the credibility of the belief-holder, and the speaker's reason for reporting a belief in the first place. To that end, this section presents a theory-neutral description of original data on *yǐwéi*, collected in consultation with fifteen native speakers of Mandarin Chinese (all young-adult international students educated in Mainland China through high school or college) at two different United States universities.

Yǐwéi is interesting precisely because it sheds light on the pragmatic complexity of belief reports, and this complexity must be taken into account when eliciting and interpreting data. The discourse context should offer a reason to discuss people's beliefs about the world rather than the world itself. It should be clear whether the speaker has an opinion about the truth or falsity of the reported belief or remains uncommitted; whether the belief itself is already taken as true in the Common Ground, already taken as false, or left open;

³Lee, Olson, & Torance (1999) claim that *dāng* reports false beliefs and note that its use-conditions overlap with those for *yǐwéi*. But they find that adults almost never produce *dāng* spontaneously (whereas they readily produce *yǐwéi*). In my own elicitations (described further in the Appendix at <https://osf.io/pr3vt/>), consultants do not want to use *dāng* as a mono-morphemic belief verb.

⁴My brief literature review is regrettably limited by my inability to read the rich portion that is written in Mandarin.

and in some cases, what evidence grounds the belief-holder's belief. As a result, all of the sentences that I present must be understood (and were elicited) in rich discourse contexts (set up using English as a metalanguage; Matthewson 2004) attempting to control these important contextual factors.

Yǐwéi is also interesting precisely because it gives rise to different inferences in different discourse contexts. The reported belief is sometimes taken as false, sometimes questionable or poorly-evidenced, sometimes hedged. These different inferences do not contradict each other, but must be understood together to illustrate different facets of the complex reality.

This presentation distills the data to drive towards an analysis thereof. I also offer an Appendix (<https://osf.io/pr3vt/>) presenting long-form accounting of all of the discourse contexts where the felicity of each utterance was judged, along with consultants' judgments and comments. Each class of examples is given a label ((4): 'Sick') for cross-reference with the Appendix.

2.1. *Non-first-person*

Since *yǐwéi* is a belief verb, its effect on the discourse depends on who the belief-holder is, particularly whether the belief-holder is the speaker (as in the case of first-person *yǐwéi*) or someone else. Therefore, I first focus on *yǐwéi* in non-first-person contexts.

When the belief has not been settled either way in the discourse, *yǐwéi* strongly suggests that the belief is wrong. If it's not known whether the speaker is sick or not (for example, joining an ongoing conversation of acquaintances), but it can be assumed that the speaker has an opinion one way or the other (Bartsch 1973), plausible because people generally have opinions regarding their own health), (4) – repeated from above – by default conveys that the speaker is *not* sick.

(4) **Sick**

Māma yǐwéi wǒ bìng le
Mother yǐwéi I sick ASP

'Mom is under the impression that I'm sick.' (=1); from LineDictionary)

When the reported belief is already established as true, *yǐwéi* cannot be used. If all interlocutors know that the speaker *is* sick, a sentence such as (4) is rejected. The best verb in this context would be *zhīdào* 'know.' A neutral verb such as *rènwéi* 'think' is less preferred (perhaps relating to the 'Maximize Presupposition' principle of Heim 1991, requiring speakers to use the presuppositionally strongest lexical item compatible with the context), but still much better than *yǐwéi*.

When the reported belief is known to be false (when it's known that I'm not sick; if I tell you that I've been faking an illness to skip school), *yǐwéi* is perfectly felicitous. (4) is natural if all interlocutors already know that the speaker is *not* sick. In this context, *zhīdào* 'know' is incoherent, because it requires its complement to be true when the context already

establishes that it is false. *Rènwéi* ‘think’ and *juéde* ‘feel that’ are felicitous here, but said to be less critical towards the mother than *yǐwéi*, describing the situation from her perspective.

Because *yǐwéi* so saliently conveys that the reported belief is wrong, it cannot be used to report a third party’s belief impartially. If a politician feels optimistic about an upcoming election, (5) could not be used by an impartial journalist, but only by a highly biased pundit. In contrast, if *yǐwéi* were replaced with *rènwéi* ‘think,’ (5) could be used impartially.

(5) **Presidential candidate**

tā yǐwéi tā huì yíng
3sg yǐwéi 3sg will win

‘She is under the impression that she’s going to win.’

As another consequence of its strong negative bias, *yǐwéi* often comes across as rude or aggressive in the second person. Like (5), (6) could not be used by a detached interviewer, but only by someone willing to antagonize the addressee. Again, (6) would become neutral if *yǐwéi* were replaced with *rènwéi* ‘think.’

(6) **Presidential candidate**

A presidential candidate says, ‘I’m feeling very confident about the election.’ A reporter follows up:

suǒyǐ nǐ yǐwéi nǐ huì yíng
so you yǐwéi you will win

‘So you’re under the impression that you’re going to win.’

While *yǐwéi* often suggests that the speaker thinks the reported belief is false, it can also be used in a context in which the speaker does not know whether the reported belief is true or not, but wants to signal that the belief-holder’s evidence for it is somehow incomplete or defective. (7) could be used in a scenario in which an American football player catches the ball in the end-zone right on the sideline, and begins celebrating – oblivious that the officials are congregating to debate whether the catch counted or not. The athlete may indeed have scored, but since he doesn’t know that the officials are debating the catch, his information is incomplete. (7) still conveys a sense of negative bias towards the reported belief, but the effect is no longer to convey that the reported belief is false – only that it is not fully informed.

(7) **Football player**

wǒ bù zhīdào yǒu-méi-yǒu défēn, dànshì zhège qiúyuán yǐwéi défēn le
I not know have-not-have score, but this-CL ball-player yǐwéi score ASP

‘I don’t know whether the player scored or not, but he’s under the impression that

he did.’

If the speaker doesn’t know whether the reported belief is true or not, but has no reason to question the belief-holder’s reasoning, then it does not make sense to use *yǐwéi*. If we just see the athlete catch the ball on the sideline and begin celebrating, Mandarin consultants reject (7), saying things like, ‘if you don’t know, why are you saying he’s wrong?’ (7) only makes sense with some reason to doubt the athlete’s reasoning.

Since (7) already shows that the belief embedded by *yǐwéi* does not have to be false, it is perhaps not surprising that this inference of falseness behaves like a conversational implicature, in that it can be reinforced without redundancy (Grice 1989, Hirschberg 1985, Levinson 2000, Potts 2014). In (8) (in an out-of-the-blue context, joining the conversation of acquaintances), the second clause confirms the expectation raised by the first clause: that the reported belief is wrong.

(8) **Billionaire**

rénmén yǐwéi tā shì yìwànfùwēng, dànshì tā bú shì
person-PL yǐwéi 3sg be billionaire but 3sg not be

‘People are under the impression that she’s a billionaire, but she’s not.’

The sense of falseness evoked by *yǐwéi* also behaves like a conversational implicature in that it can be cancelled, although the effect of doing so is more surprising than the effect of reinforcing it. The first clause of (9) alone would convey that this person is *not* a billionaire; but the second clause is a coherent, non-contradictory (though marked) continuation.

(9) **Billionaire**

rénmén yǐwéi tā shì yìwànfùwēng ...ér tā díquè shì
person-PL yǐwéi 3sg be billionaire ...and 3sg indeed be

‘People are under the impression that she’s a billionaire ... and she actually is.’

(9) is an unusual discourse move, in that it first leads hearers to infer that this person is not a billionaire, and then abruptly tells them that she is. For this reason, sentences like (9) are generally only felicitous with a veracity emphasizer such as *díquè* ‘indeed’ or *zhēnde* ‘really’ in the second clause. Such sentences also only make sense when the speaker has a specific rhetorical goal. A speaker might say, ‘I have a friend who invented a really famous app. People *yǐwéi* she’s a billionaire – and she actually is (=9)), but she never made a cent on that app. She just inherited a fortune from her parents.’ The person is indeed a billionaire – but not for the reason you might think. *Yǐwéi* makes sense here as a way of highlighting the tension between the appearance and the reality of how this person acquired her wealth. In such a context, (9) suggests that it would be incorrect to conclude that this person is a billionaire on the grounds that other people think she is – but that she is a billionaire anyway.

Although sentences like (9) subvert expectations, they are not contradictory, nor are they self-corrections. Truly contradictory sentences such as (10) are judged incoherent in a way that (9) is not.

(10) **Billionaire**

#tā shì ge yìwànfùwēng ... ér tā díquè bú shì
3sg be CL billionaire ... and 3sg indeed not be

‘#She’s a billionaire ... and she’s actually not.’

The only way to make sense of (10) is to posit that the third-person pronouns in each clause do not refer to the same person; or to imagine some way that a person could be a billionaire in one sense but not another. In contrast, (9) makes sense without such rescue strategies.

To recap, non-first-person *yǐwéi* strongly suggests that the speaker views the reported belief with skepticism. Sometimes, *yǐwéi* conveys that the speaker rejects the reported belief; but other times, the speaker simply finds it questionable or unwarranted, even if it may be true.

2.2. *First-person*

So far, we have only considered beliefs that the speaker attributes to another party (second or third person), in which case it makes sense for these beliefs to be doubted or rejected by the speaker. The situation is different when the belief-holder and the speaker are one and the same, as in first-person belief reports (*I think*). Given that it is incongruous to both believe something and want to flag it as mistaken⁵, it is perhaps surprising that *yǐwéi* can be used in the first person at all. But in fact, *yǐwéi* has two distinct first-person uses, each involving a different way of reconciling the speaker’s reported belief with the skepticism signaled by *yǐwéi*.

Most commonly, first-person *yǐwéi* communicates that the speaker *previously* believed the embedded proposition, but now either thinks it is false or is confused as to whether it is true, as in (11).

(11) **The talk today**

wǒ yǐwéi jīntiān yǒu ge jiǎngzuò
I yǐwéi today have CL talk

‘I thought there was a talk today.’

⁵This observation has a long history in the philosophy literature: Hintikka (1962) works to explain why *It’s raining but I don’t believe it’s raining* is an absurd assertion (Moore’s Paradox, Moore 1993 [1942]); Wittgenstein (1953) observes that ‘If there were a verb meaning ‘to believe falsely,’ it would not have any significant first person, present indicative’ (para. 470), brought to my attention by Daniel Lassiter.

If it is already established that there is no talk, (11) would simply convey that the speaker used to think there was a talk and now knows there isn't. If it is not established whether there is a talk or not, (11) would suggest that the speaker used to think there was a talk, and now either realizes there isn't one, or has become confused as to whether there is one or not – for example, upon finding the conference room empty. The speaker might continue: 'Maybe there is a talk today, but I just have the wrong room.'

This past interpretation of *yǐwéi* can be understood in light of the way temporal information is conveyed in Mandarin more generally. Past is not morphologically distinct from the present; atelic verbs such as *yǐwéi* are interpreted as present by default, but can be shifted around depending on the context (Lin 2003, Smith & Erbaugh 2005, Lin 2006). In (11), it seems that the meaning of the sentence itself provides enough context for it to be understood as past even though it is atelic. Without backshifting, (11) would describe a conflicted mental state: the speaker both believes there is a talk and wants to suggest that there is not. But when (11) is understood as past, it is much more coherent: the speaker used to think there was a talk, and now realizes there isn't or has become confused. Interestingly, *yǐwéi* stands out from its alternatives in giving rise to this backshifting. With *rènwéi* 'think,' *zhīdào* 'know' or *juéde* 'feel that' substituted for *yǐwéi*, (11) would by default describe the speaker's *current* belief.

In addition to the past interpretation of first-person *yǐwéi*, there is also a present interpretation, in which the speaker holds the belief at speech-time but welcomes the hearer to disagree. In (12) (in a context where I am suggesting that you go to graduate school), the speaker is understood to currently believe that "you should do this," but wants to frame the advice as a suggestion which might be wrong, and which the hearer is free to disregard. (In comparison to the past understanding of *yǐwéi*, the hedged-present understanding is less common, requires more contextual support, and can be perceived as more formal.)

(12) **Personal advice**

wǒ gèrén yǐwéi nǐ yīnggāi zhèyàng zuò
I personally yǐwéi you should this-way do

'Personally, I would think you should do this.'

(12) is described as tentative, conservative, and hedged, since this use of *yǐwéi* seems to diminish the speaker's confidence in the assertion. Consultants say that such sentences would be stronger if *rènwéi* 'think' were used instead.

Summing up again, first-person *yǐwéi* can be understood as past or hedged – two different ways for the speaker to hold (or to have held) the reported belief while also flagging it as questionable.

2.3. Projection

To diagnose the source of *yǐwéi*'s negative bias, it is also important to consider its behavior under entailment-cancelling operators such as questions, conditional antecedents, possibil-

ity modals, and negation (Langendoen & Savin 1971, Karttunen 1973). These operators are said to suspend entailments (at-issue content; Potts 2005, Simons, Tonhauser, Beaver, & Roberts 2010) while allowing presuppositions and conventional implicatures (non-at-issue content) to survive.

Looking first at questions, (13) most saliently conveys that the speaker believes there is no test. (13) would make sense if all interlocutors already know that there is no test (say we're going on a field trip tomorrow, and wondering if the hardworking Lili is so paranoid as to be imagining a pop quiz on the field trip). If it is not yet settled in the discourse whether there is a test or not (in an out-of-the-blue context), (13) suggests that the speaker thinks there is no test, or that the speaker has some reason to question the belief-holder's credibility on the matter. If *rènwéi* 'think' were used instead, the issue would be open.

(13) **Test tomorrow**

Lìlì yǐwéi míngtiān yǒu kǎoshì ma?
Lili yǐwéi tomorrow have test QUESTION

'Is Lili under the impression that there's a test tomorrow?'

As another example, (14) is understood as aggressive, suggesting that the hearer is mistaken in thinking that their actions are acceptable. (15) (using *rènwéi* 'think') could come across as similarly aggressive, but could also be a neutral question about one's self-concept.

(14) **Your self-concept**

nǐ yǐwéi nǐ shì shéi?
you yǐwéi you be who

'Who the hell do you think you are?'

(15) **Your self-concept**

nǐ rènwéi nǐ shì shéi?
you think you be who

'Who do you think you are?'

Turning to conditional antecedents, (16) (in an out-of-the-blue context) conveys that in the speaker's view, the belief-holder would be incorrect to think there is a test tomorrow. If *rènwéi* were used instead, the issue would be open. Here as well, *yǐwéi*'s sense of wrongness survives.

(16) **Test tomorrow**

rúguǒ tā yǐwéi míngtiān yǒu kǎoshì de-huà, tā yīnggāi zài xuéxí
if 3sg yǐwéi tomorrow have test in-that-case, 3sg should PROG study

‘If she thinks there’s a test tomorrow, she should be studying.’

As for possibility modals, (17) is taken to convey that there is no test, or at least that this belief is somehow questionable or defective. It cannot be used when we know for sure that there is a test. Again, *yǐwéi*’s sense of wrongness persists. Again, if *rènwéi* were used, (17) would be much more neutral.

(17) **Test tomorrow**

Lìlì kěnéng yǐwéi míngtiān yǒu kǎoshì
Lili may yǐwéi tomorrow have test

‘Lili might be under the impression that there is a test tomorrow.’

While (13)–(17) are quite clear, the data become more elusive when we turn to the question of whether the sense of wrongness signaled by *yǐwéi* also projects out of negation.

For some reason, *yǐwéi* is generally rejected under the most common Mandarin negation morpheme, *bù*⁶: (18) is judged unacceptable, no matter whether there is a test or not.

(18) **Test tomorrow**

#Lìlì bù yǐwéi míngtiān yǒu kǎoshì
Lili not yǐwéi tomorrow have test

‘#Lili doesn’t think there’s a test tomorrow.’

Some consultants⁷ suggests that (18) is rejected because the string *bù yǐwéi* constitutes the first three characters of a very common four-character phrasal idiom⁸, *bùyǐwéirǎn*, ‘to take objection to’ – so that the first three characters of this common four-character phrase constitutes a ‘garden path’ in the sense of Bever (1970). *Yǐwéi* can much more easily be

⁶Apart from *bù*, the other major sentential negator in Mandarin is *méi*. *Méi yǐwéi* is much more acceptable than *bù yǐwéi* is, although about a third of my Mandarin consultants still find it marked and say they never use it. But these data must be interpreted carefully. *Méi* argued to comprise not just negation, but also perfective aspect (for example, *méi* is in complementary distribution with the perfective marker *le*; Huang 1988). Therefore, *méi yǐwéi* is predicted to trigger complex inferences arising from the interaction of negation and perfective aspect (which often signals information about tense as well as aspect, triggering ‘cessation implicatures’ that past states do not continue into the present; Altshuler & Schwarzschild 2013), along with whatever negative bias is conveyed by *yǐwéi*. Perhaps thanks to this blend of aspect/tense, negation, and negative bias, Mandarin consultants disagree vehemently on when *méi yǐwéi* can be used and what it conveys. (Please see the Appendix, <https://osf.io/pr3vt/>, for data).

⁷Thanks to Jinyi Chu and Jia’ao Chen for this insight.

⁸Mandarin has many such four-character *chéngyǔ* idioms, which blur boundaries between words and phrases; please see Jiao, Kubler, & Zhang (2013) for discussion.

negated by the compound negation *bìng-bù* ‘not at all’ – perhaps because the common bigram *bìng-bù* forestalls the garden path. But even then some consultants would prefer to substitute *rènwéi*, or to re-word the sentence with negation in the lower clause rather than the matrix clause. (Any study of negated attitude verbs of course implicates neg-raising; Horn 1971, Bartsch 1973, Horn 1978, 1989) – a topic of ongoing debate in Mandarin; please see Xiao & McEnery 2008, Xiang 2013, Popp 2016).

(19)–(20) are presented with full discourse contexts, both because negation is rare out-of-the-blue (Horn 1989, Matthewson 2004), and because it is valuable to make the truth or falsity of the reported belief explicit. (The Appendix, <https://osf.io/pr3vt/>, provides further examples of negated *yǐwéi* in various contexts.) In (19), the belief reported by negated *yǐwéi* – that there’s a test tomorrow – is explicitly false, while in (20) it is true. Interestingly, modulo some slight discomfort with negating *yǐwéi*, both (19) and (20) are accepted.

(19) **Test tomorrow**

We all know that there is no test tomorrow because we are going on a field trip. Lili also knows that there is no test; that’s why she is out seeing a movie tonight instead of studying. Explaining Lili’s whereabouts, I say:

Lìlì bìng-bù yǐwéi míngtiān yǒu kǎoshì ér quèshí meíyǒu
Lili definitely-not yǐwéi tomorrow have test and actually not-have

‘Lili definitely doesn’t think there’s a test tomorrow, and there’s not.’

(20) **Test tomorrow**

We all know that there definitely is a test tomorrow. But Lili has been absent, so she doesn’t know about the test; that’s why she is out seeing a movie tonight instead of studying. Explaining Lili’s whereabouts, I say:

Lìlì bìng-bù yǐwéi míngtiān yǒu kǎoshì dànshì yǒu
Lili definitely-not yǐwéi tomorrow have test but have

‘Lili definitely doesn’t think there’s a test tomorrow, but there is.’

In sum, while the negative bias of *yǐwéi* projects out of questions, possibility modals, and conditional antecedents, it puzzlingly seems to disappear under negation. We might have expected that a sentence of the form *x bìng-bù yǐwéi p* would convey skepticism towards *p*; and yet such a sentence can be used when *p* (that there’s a test) is Common Ground, just as it can be used when not-*p* is Common Ground (where there’s no test).

2.4. Discussion

Table 1 recaps the effect of *yǐwéi* in different contexts.

Facing these data, the question is: What blend of semantics and pragmatics derives *yǐwéi*’s negative bias?

Context	Effect
1st person <i>yǐwéi p</i>	understood as past or hedged
2rd person <i>yǐwéi p</i>	typically rude because it suggests hearer is in error
3rd person <i>yǐwéi p</i>	typically conveys that speaker is skeptical towards <i>p</i>
3rd person <i>yǐwéi p</i> and in fact <i>p</i>	subverts expectations
conditionals, questions, <i>might</i>	sense of wrongness persists
negation	sense of wrongness seems to disappear

Table 1: Effects of *yǐwéi* in different contexts.

On the one hand, at least some part of its effect clearly arises pragmatically, because it is the context which determines whether the speaker rejects the belief, finds it unwarranted, previously believed it, wants to hedge it, and so on, and because the inference that the belief is false can be reinforced and cancelled like a conversational implicature (Grice 1989, Hirschberg 1985, Levinson 2000, Potts 2014).

On the other hand, at least some part of *yǐwéi*'s negative bias seems tied to its conventional meaning. When multiple near-synonyms give rise to the same inference (for example, *Alice {quit/stopped/ceased} smoking* all convey that Alice previously smoked), that inference may be considered 'non-detachable' from the message itself (Simons 2006 drawing on Grice 1989) and thus arguably not part of the conventional meaning of any particular form. But we have seen that other attitude verbs such as *rènwéi* 'think' and *juéde* 'feel that' behave differently from *yǐwéi*, so its negative bias does in fact seem conventional to the word itself.

Moreover, *yǐwéi*'s projection behavior (except for the puzzling negation facts) is reminiscent of various types of non-at-issue content, such as presuppositions and conventional implicatures (Potts 2014). Its negative bias also seems to be 'backgrounded' (Simons 2006), distinct from the 'main point' of an utterance that answers the Question Under Discussion in the sense of Roberts (2012). When we are wondering what Lili is up to, a sentence of the form 'Lili *yǐwéi* there's a test tomorrow' addresses the QUD by describing Lili's belief that there's a test (thus, she's studying); the unreliability of this belief is a side note or perhaps already assumed (if it's already Common Ground that there is no test). In providing/reinforcing background information, the negative bias of *yǐwéi* behaves similarly to a presupposition.

These two pieces of *yǐwéi*'s meaning – its at-issue description of belief, and its backgrounded negative bias – are also semantically independent from one another (Potts 2005, Simons, Tonhauser, Beaver, & Roberts 2010): the claim that *x believes p* is of course logically independent from the status of *p* as true, false, questionable, or unknown. Here, the negative bias of *yǐwéi* appears similar to many conventional implicatures and other projective content (Potts 2005, Simons, Tonhauser, Beaver, & Roberts 2010).

In sum, although *yǐwéi*'s negative bias is pragmatically flexible, it seems that these pragmatic calculations must ultimately be grounded in some sort of backgrounded conventional meaning, which distinguishes *yǐwéi* from its alternatives.

3. Analysis

To account for these data, I propose a two-part analysis of *yǐwéi* invoking both a semantic element and a pragmatic element: (i) a conventional, semantic requirement that *p* must not be Common Ground following an assertion of *x yǐwéi p* (a postsupposition); and (ii) a series of context-sensitive pragmatic calculations triggered by a speaker’s choice to use *yǐwéi* over a neutral alternative such as *rènwéi* ‘think.’

The analysis is situated in a framework in which sentences serve as updates to the Common Ground (the set of propositions mutually agreed on by interlocutors), or equivalently to the context set, the set of worlds consistent with the propositions in the Common Ground (Karttunen 1974, Stalnaker 1979). Formally (drawing on Heim 1982: Chapter 3), when an assertion *S* is accepted in a context *c*, the new context *c + S* is restricted to only the worlds in *c* in which *S* is true – adding information to the Common Ground by narrowing the context set of worlds considered possible.

3.1. Definedness conditions on input and output contexts

In this setup, a *presupposition* can be stated as a definedness condition (Strawson 1950) on input contexts (Stalnaker 1979).

To capture the longstanding idea that *know* presupposes its complement (Kiparsky & Kiparsky 1970, Stalnaker 1974), a sentence of the form *x knows p* is analyzed to update the context with *x believes p* (its entailed, at-issue content; Potts 2005, Simons, Tonhauser, Beaver, & Roberts 2010), and is defined only if the input context already entails *p* (its presupposed, projecting content). According to (21), a speaker uses *know* (Mandarin *zhīdào*) to signal that they take its complement *p* to already be Common Ground⁹.

- (21) at-issue and presupposed content of *zhīdào* ‘know’
- a. $c + x \textit{zhīdào } p = c + x \textit{believes } p$
 - b. defined only if $\forall w[w \in c \rightarrow p(w) = 1]$

Stepping back, there is a debate in the presupposition literature (e.g., Strawson 1950, Stalnaker 1974, Abusch 2002, Simons 2006, Abusch 2010) about whether presuppositions should be considered semantic properties of sentences (if presuppositions are viewed as restrictions on the usage of particular constructions or lexical items), pragmatic actions on the part of speakers (if presuppositions are viewed as facts that a speaker takes for granted), or perhaps sometimes both (“using a [lexical] presupposition trigger is an excellent way to achieve the speaker action of presupposing,” from Potts 2014: §2.1.2). In framing the presupposition of *know* as a definedness conditions on its input context, (21) assumes that at least some presuppositions are semantically encoded in lexical items. My goal here is

⁹Of course, speakers often bend the rules, using *know* to signal that they think *p* should be accommodated as Common Ground even if it is not yet so (Stalnaker 1979, Lewis 1979, von Stechow 2008, Schlenker 2012), especially if the speaker’s belief in *p* (signaled by their use of a presuppositional lexical item) serves as evidence that the hearer should take it up too.

not to defend this analysis of *zhīdào* ‘know’ (21), but rather to use it as a comparison to the analysis proposed for *yǐwéi*.

To capture the negative bias associated with *yǐwéi*, I propose that *yǐwéi* has a different sort of definedness condition: that there exists at least one not-*p* world in the context set, so that speakers and hearers must mutually entertain the possibility that *p* is false. Moreover, I argue, it is not enough for this definedness condition to be stated on the input context, as a presupposition. If a sentence of the form *x yǐwéi p* simply required its *input* context to be compatible with not-*p*, then – recalling the pragmatic complexity of belief reports – in a situation where *x* is authoritative and *p* is plausible, *following* the assertion, *p* might become Common Ground on the evidence that *x* believes it. On my analysis, *yǐwéi* explicitly prevents that outcome.

3.2. *Yǐwéi* and its proposed postsupposition

To derive the effect of *yǐwéi*, I argue that speakers and hearers must continue to entertain the possibility that *p* is false not just prior to the assertion, but also afterwards. More formally: not just the input context, but also the output context, must be compatible with not-*p*. In an update framework, such a definedness condition is called a *postsupposition*, using a term from Brasoveanu (2009) and Lauer (2009) which builds on an idea from Farkas (2002a,b). In other words, *yǐwéi* is analyzed as in (22). *Yǐwéi* updates the context with *x believes p*, and postsupposes that the *output* context is consistent with not-*p*.

- (22) at-issue and postsupposed content of *yǐwéi*
- a. $c + x \text{ yǐwéi } p = c + x \text{ believes } p$
 - b. defined only if $\exists w \in (c + x \text{ believes } p) : p(w) = 0$

Of course, since an output context is always a subset of its input context, then if there is a not-*p* world in the output context, there must have been one in the input context as well. In that sense, the postsupposition of *yǐwéi* also serves as a presupposition, requiring a not-*p* world in the input (which is why, I argue, *yǐwéi* is rejected if *p* is already taken to be true; §2). But in a discourse where *x*’s belief in *p* serves as evidence for *p*, a presupposition requiring a not-*p* world in the input context would still allow *p* to become Common Ground following the assertion, as in (23).

- (23) **Initial context:** Not sure what happened to the motorcyclist; maybe he hit something, maybe not
(Context includes p worlds and $\neg p$ worlds, where p is the proposition that the motorcyclist hit something.)
Utterance: The investigators think the motorcyclist hit something in the road.
Output context: We assume the motorcyclist hit something.
(p is taken up on the grounds that the investigators believe it; now the context includes only p worlds.)

In contrast, the proposed postsupposition of *yǐwéi* serves to explicitly prevent this effect.

No matter how credible the belief, or how authoritative the belief-holder, p is not to become Common Ground. Informally, (22) can be paraphrased as: ‘ x believes p – but we won’t take their word for it.’

In contrast to *yǐwéi*, a neutral belief verb such as *rènwéi* ‘think’ is given no definedness condition at all: x *rènwéi* p simply yields the worlds in c where it is true that x believes p (just (22a), without (22b)). Therefore, *yǐwéi* can be used in a subset of the contexts where *rènwéi* can: *rènwéi* but not *yǐwéi* can be used when p is taken as true (the postsupposition of *yǐwéi* cannot be satisfied if the input context contains only p worlds); and *rènwéi* but not *yǐwéi* can be used where p is considered true if x believes it (the postsupposition of *yǐwéi* cannot be satisfied if x believes p is taken to contextually entail p). The speaker’s choice to use *yǐwéi* over *rènwéi* therefore provides information about what they take to be Common Ground both before and after the utterance.

3.3. *Updating with yǐwéi and rènwéi in an uncertain Common Ground*

In order for *yǐwéi* to provide information about what the speaker takes to be Common Ground, it is important to establish that both speakers and hearers may be at least slightly uncertain as to what the Common Ground contains. (Belief-holders may also be uncertain about the beliefs that they hold; but my analysis focuses only the uncertainty of interlocutors about the state of the conversation that they are in.)

Common Ground is defined as the propositions that all interlocutors believe and believe that they all believe; but beliefs can be uncertain, and people may mistake, forget, or fail to pay attention to what they have previously agreed on. This idea has a long precedent: Walker (1996) observes that there can be uncertainty as to whether a speaker’s assertions have been accepted by other interlocutors as true. Beaver (2001) suggests that, instead of assuming that all propositions in the Common Ground are equally definitive, some should be considered more or less likely than others, introducing uncertainty by making the CG probabilistic. Stalnaker (2002) discusses how defective or misaligned assumptions between interlocutors may come to light or be corrected. Horton & Gerrig (2005) show empirically that Common Ground depends on fallible human memories of prior discourse. Lewis (1979), Rawlins (2010), and Klecha (2014) explore how speakers can be misaligned on the possibilities under consideration. Schlenker 2012 derives the principle of Maximize Presupposition (Heim 1991, Percus 2006) in Gricean terms of quantity on the grounds that there is a non-zero chance of people forgetting previously agreed-upon information (Fallibility). Crone (2018) explains seemingly-redundant reminders of what’s already Common Ground because people may forget these commitments.

In light of this uncertainty, lexical items with definedness conditions (such as *zhīdào* ‘know’ and *yǐwéi*) do not just reflect established information, but may also add information by reducing uncertainty regarding the Common Ground. Even if definedness conditions are considered non-at-issue (in that they project and don’t directly address the Question Under Discussion in the sense of Roberts 2012), they can still provide information about what the

speaker takes as Common Ground¹⁰.

To illustrate the proposal, I go through several different possible states of the Common Ground before and after it is updated with the postsuppositional *x yǐwéi p* versus the neutral *x rènwéi p*, summarized in Table 2.

When *p* is already in the Common Ground before the assertion (top row of Table 2), *x rènwéi p* is a consistent update, but *x yǐwéi p* would create a contradiction: *p* is incompatible with *yǐwéi*'s requirement that the Common Ground be consistent with not-*p*. Therefore, a speaker who uses *yǐwéi* signals that they do not take *p* to already be Common Ground. This signal is redundant if there is no uncertainty about the CG; but if there is any uncertainty, it can be informative.

When it is Common Ground that *p* is true if *x* believes it (second row of Table 2), then *x rènwéi p* leads the Common Ground to be updated with *p* (as illustrated in the evidential uses of belief verbs from §1, like *the investigators think the motorcyclist hit something*). In contrast, *x yǐwéi p* again yields a contradiction (requiring both *p* and possibly not-*p* to be Common Ground). Therefore, a speaker who uses *yǐwéi* signals (again, redundantly or informatively) that they do not take *if x believes p, then p* to be part of the Common Ground.

When not-*p* is already in the Common Ground before the assertion (third row of Table 2), then both *x rènwéi p* and *x yǐwéi p* are consistent updates, but *yǐwéi*'s postsupposition reiterates the possibility that *p* is false. Again, given that people may mistake or forget or fail to attend to what's in the Common Ground, *yǐwéi*'s signal that *p* is not Common Ground may be redundant or informative.

When the Common Ground is compatible with both *p* and not-*p* before the assertion (fourth row of Table 2), then both *x rènwéi p* and *x yǐwéi p* are consistent updates. Redundantly or informatively, *yǐwéi* flags that *p* may be false, and thus also that *x*'s belief in *p* cannot constitute definitive evidence for it; while the neutral *rènwéi* does not signal any such skepticism.

When it is Common Ground that the speaker either believes *p* or believes not-*p* (opinionatedness; Bartsch 1973, in the fifth row of Table 2), then *x rènwéi p* leaves both options open. However, a speaker who uses *yǐwéi* signals that they want the Common Ground to be compatible with 'possibly not-*p*,' most likely because their own beliefs are compatible with *p* being false. Assuming that they either believe *p* or believe not-*p*, a speaker who signals that they believe 'possibly not-*p*' ultimately conveys not just that they believe *p* might be false, but that they believe *p* actually *is* false. (To reiterate: if a speaker either believes *p* or believes not-*p*, then if they believe possibly not-*p*, it follows that they believe not-*p*).

Finally (last row of Table 2), imagine that the hearer is bewildered about the speaker's assumptions, so the Common Ground is fully uncertain. Here, *yǐwéi* informs the hearer that the speaker thinks neither *p* nor *x*'s authority should be taken for granted (because then *yǐwéi*'s postsupposition would create a contradiction), while *rènwéi* leaves those possibilities open.

In sum, *yǐwéi* can be used in a subset of the contexts where its neutral alternative *rènwéi* 'think' can (Table 2). With no definedness condition on its input or output contexts, *x*

¹⁰Please see Schlenker (2012) for arguments that the presupposition of *know* can be informative for this reason.

Initial Common Ground	Utterance	Updated Common Ground
p	x <i>rènwéi</i> p	p , x believes p [initial CG] [utterance]
	x <i>yǐwéi</i> p	p , x believes p , possibly not p → Contradiction [initial CG] [utterance] [postsupposition]
if x believes p , then p	x <i>rènwéi</i> p	if x believes p , then p x believes p , p [initial CG] [utterance] [initial CG + utterance]
	x <i>yǐwéi</i> p	if x believes p , then p x believes p , p , possibly not p → Contradiction [initial CG] [utterance] [initial CG + utterance] [postsupposition]
not p	x <i>rènwéi</i> p	not p , x believes p [initial CG] [utterance]
	x <i>yǐwéi</i> p	not p , x believes p , possibly not p [initial CG] [utterance] [postsupposition]
possibly p , possibly not p	x <i>rènwéi</i> p	possibly p , possibly not p , x believes p [initial CG] [initial CG] [utterance]
	x <i>yǐwéi</i> p	possibly p , possibly not p , x believes p , possibly not p [initial CG] [initial CG] [utterance] [postsupposition]
Speaker believes p OR believes not- p	x <i>rènwéi</i> p	S believes p OR believes not- p , x believes p [initial CG] [utterance]
	x <i>yǐwéi</i> p	Speaker believes p OR believes not- p , x believes p , possibly not p [initial CG] [utterance] [postsupposition]
??	x <i>rènwéi</i> p	x believes p [utterance]
	x <i>yǐwéi</i> p	x believes p , possibly not p [utterance] [postsupposition]

Table 2: Updating with *rènwéi* versus *yǐwéi* in different contexts on the proposed analysis. As illustrated in the first two rows of the table where *yǐwéi* but not *rènwéi* creates a contradiction, *yǐwéi* can be used in a subset of the contexts where *rènwéi* can be used.

rènwéi p provides no information about p except that x believes it, allowing that p may be or may become Common Ground. Even if it is not definitively agreed that p is true or that x is informed about p , interlocutors might still entertain those possibilities, since *rènwéi* does not reduce any uncertainty about this. So x *rènwéi* p leaves open the possibility that x and/or p may be considered reliable.

In contrast, a speaker’s choice to use x *yǐwéi* p signals that we are definitely not in a context where p is taken as true, nor in a context where x is considered authoritative – because *yǐwéi* would create a contradiction in such contexts, as illustrated in the first two rows of Table 2. A speaker’s choice to use *yǐwéi* reduces uncertainty about both the input and output contexts in a way that signals skepticism towards x and p . If everyone already knows not- p or agrees that p or x are questionable, then *yǐwéi* just reiterates that information; but if anyone was confused or forgetful, *yǐwéi* raises awareness that p and x ’s belief in it are to be treated with skepticism.

I argue that these effects explain the negative bias associated with *yǐwéi*.

4. Explaining the data

4.1. Inference that p is false

Recall that when a reported belief is not settled in the discourse, *yǐwéi* strongly suggests that it is false, as in (24).

(24) **Sick**

Māma yǐwéi wǒ bìng le
 Mother yǐwéi I sick ASP

‘Mom is under the impression that I’m sick.’ (= (1), (4))

The proposed analysis explains why. Presumably, the speaker has an opinion as to whether they are sick or not; they either believe p or believe not- p (fifth row of Table 2). In such a context, ‘possibly not- p ’ (signaled by the speaker’s choice to use *yǐwéi*) is strengthened to ‘not- p ’ when combined with the assumption that the speaker either believes p or its negation – deriving the inference that the speaker is *not* sick (that the mother is wrong). (25) spells out this reasoning step by step.

- (25)
- a. *Contextual premise*: Speaker knows (or has an opinion as to) whether they are sick or not (Sp believes $p \vee$ Sp believes not- p).
 - b. *Contextual premise*: If Speaker indeed believed that they are sick, they would have no objection to that information becoming Common Ground.
 - c. Assume that Speaker is cooperative (Grice 1989) and is following Grice’s conversational maxims (quality, quantity, manner, relevance).
 - d. By the postsupposition of *yǐwéi*, Speaker signals that they do not want p (that Speaker is sick) to become Common Ground following an assertion of (24).
 - e. By (b) and (d): It must not be the case that Speaker believes they are sick, because then they would have no objection to this becoming Common Ground.
 - f. Since Speaker either believes they are sick or believes they are not sick (a), and since it is not the case that they believe they are sick (e), then Speaker must believe that they are *not* sick – that p is false, and that the mother’s belief

is incorrect.

The proposed analysis also explains why (24) is nonsensical if the speaker *is* known to be sick. Following the first row of Table 2, it is a contradiction for *yǐwéi* to require a not-*p* world in the output context if the input context contains only *p*-worlds.

4.2. Inference that *p* is questionable even if possibly true

In many cases, the speaker can be assumed to have an opinion about the reported belief, leading to the inference that the speaker who uses *yǐwéi* thinks the belief is false (25). But sometimes this assumption of opinionatedness is called off, as in (26), in which the speaker explicitly claims not to know whether the athlete scored or not (as in the fourth row of Table 2: the Common Ground is compatible with both *p* and not-*p*).

(26) Football player

wǒ bù zhīdào yǒu-méi-yǒu défēn, dànsì zhège qiúyuán yǐwéi défēn le
I not know have-not-have score, but this-CL ball-player yǐwéi score ASP

‘I don’t know whether the player scored or not, but he’s under the impression that
he did.’ (= (7))

Empirically (§2), (26) is only felicitous if there is at least *some* reason to question the football player’s reasoning. If the speaker just sees the athlete catch the ball on the sideline and begin celebrating, (26) is unusual (although it would be fine if *rènwéi* ‘think’ were used), since it seems to cast doubt on the athlete’s belief without justification. In contrast, if the speaker sees the officials congregating to discuss whether the catch counted, (26) is felicitous, since the speaker has reason to suspect that the athlete’s belief is not fully informed.

I argue that the proposed analysis explains these facts. First, *yǐwéi* explicitly signals that the speaker takes neither *p* nor *if x believes p, then p* to be Common Ground (because *yǐwéi* would create a contradiction there, as in the first two rows of Table 2). Second, when the Common Ground is compatible with both *p* and not-*p* (as indicated by ‘I don’t know whether he scored or not’), then the postsupposition of *yǐwéi* reiterates and highlights the possibility of not-*p* (fourth row of Table 2). The effect is to flag *p* as questionable rather than simply unknown.

One might object that if the Common Ground is compatible with *p* and not-*p* (fourth row of Table 2), the postsupposition of *yǐwéi* should be vacuous: it simply restates the possibility that not-*p*. So why does *yǐwéi* have a different effect than *rènwéi* in such a context? Again, the explanation begins with the claim that people may mistake, forget, or fail to attend to what is Common Ground (even within the same sentence; Schlenker 2012); and that the definedness condition of *yǐwéi* reduces such uncertainty by drawing attention to the possibility that *p* may be false and/or that the belief-holder *x* may be unreliable – in turn triggering pragmatic inferences about why the speaker chose to do so. That,

I argue, is why (26) is rejected when the speaker simply sees the athlete catch the ball and begin celebrating (with no reason to question him), but is accepted when the speaker sees the officials congregating to debate the catch (and thus sees his belief as insufficiently informed).

4.3. Reinforcement and cancelation

Moving forward, the inference that p is false is derived pragmatically, by the way the post-supposition of *yǐwéi* interacts with the assumptions that the speaker has an opinion regarding p and wants the Common Ground to be consistent with not- p . Therefore, it is no surprise that this inference can be reinforced without redundancy (27), and cancelled without contradiction (28) – two hallmarks of pragmatic inferences (Grice 1989, Hirschberg 1985, Potts 2014). In (28), the speaker suggests that ‘she is a billionaire’ should not be taken up on the grounds that other people believe it, because their reasoning is faulty (*if x believes p , then p is not Common Ground*) – but that it is true anyway¹¹.

(27) Billionaire

rénmén yǐwéi tā shì yìwànfùwēng, dànshì tā bú shì
 person-PL yǐwéi 3sg be billionaire but 3sg not be

‘People are under the impression that she’s a billionaire, but she’s not.’ (= (8))

(28) Billionaire

rénmén yǐwéi tā shì yìwànfùwēng ...ér tā díquè shì
 person-PL yǐwéi 3sg be billionaire ...and 3sg indeed be

‘People are under the impression that she’s a billionaire ... and she actually is.’ (= (9))

4.4. First-person *yǐwéi*

When the belief-holder and the speaker are the same, as in first-person *yǐwéi*, hearers infer different reasons that the speaker chose to use *yǐwéi*. Rather than deciding that the speaker both believes p and wants to flag it as false or questionable (an incongruous mental state), hearers may determine that the speaker *used to* believe p and *now* finds it false or questionable (an available understanding of (29), because past and present are not morphologically distinguished in Mandarin).

¹¹Schematically: to start, the Common Ground is compatible with both ‘she’s a billionaire’ (p) and ‘she’s not a billionaire’ (not- p); then the speaker uses *yǐwéi* to add the information that *people believe she’s a billionaire* while signaling that *if they believe she’s a billionaire, she is one (if x believes p , then p) is not Common Ground* (conveying that other people are not to be taken as authorities on her wealth); then the speaker proceeds to update the Common Ground with *she is a billionaire* – a consistent series of updates.

(29) **The talk today**

wǒ yǐwéi jīntiān yǒu ge jiǎngzuò
I yǐwéi today have CL talk

‘I thought there was a talk today.’ (= (11))

Or the hearer may infer that the speaker does currently believe p , but does not want it to become Common Ground in case the hearer disagrees (or because they to flag their opinion as fallible by signaling that *if I believe p , then p* should not be considered Common Ground), as in (30). Normally, if a speaker wants to prevent p from being Common Ground, it is because the speaker disbelieves it; but in the hedged case, I argue that it is because the speaker does not want to presume that the hearer accepts p even if the speaker does.

(30) **Personal advice**

wǒ gèrén yǐwéi nǐ yīnggāi zhèyàng zuò
I personally yǐwéi you should this-way do

‘Personally, I would think you should do this.’ (= (12))

4.5. *Projection*

The proposed analysis also explains why the negative bias associated with *yǐwéi* projects. As definedness conditions, postsuppositions project – but not exactly the same way as presuppositions, behaving uniquely with respect to negation (Lauer 2012).

Looking first at questions, I adopt for concreteness the analysis of Groenendijk & Stokhof (1984). On this analysis, a polar question such as (31) (repeated from above) partitions the context set into worlds in which the proposition (*Lili yǐwéi there is a test*) is true, and those in which it is false. Regardless of whether it is true or false, if that proposition is defined, the postsupposition of *yǐwéi* requires there to be some worlds in the context set in which there is no test. Thus, both among the worlds in which Lili believes there is a test, and among the worlds in which Lili does not believe there is a test, there are required to be some worlds in which there is no test: if Lili thinks there is a test, she might be wrong. The postsupposition of *yǐwéi* projects, consistent with the data.

(31) **Test tomorrow**

Lìlì yǐwéi míngtiān yǒu kǎoshì ma?
Lìlì yǐwéi tomorrow have test QUESTION

‘Is Lili under the impression that there’s a test tomorrow?’ (= (13))

As a result, (31) signals that we are not in a context where it’s known that there is a test, nor in a context where Lili is considered authoritative on the matter. If *rènwéi* were used

instead, then (31) could seek Lili’s opinion as definitive, but *yǐwéi* signals that it is not.

Turning to conditionals, it is common in the literature (Karttunen 1974, Heim 1983) to assume that the antecedent of a conditional is added to a version of the context set first, and then the consequent of the conditional is added to a context set that has already been updated with the antecedent. In (32) (repeated from above), a version of the context set would first be updated with the antecedent *she yǐwéi there is a test*. The result of that update is required by the postsupposition of *yǐwéi* to contain some worlds in which there is no test: in other words, if she thinks there is a test, she might be wrong. Thus, in a conditional antecedent, *yǐwéi* continues to suggest that the belief-holder is not a reliable source regarding the reported belief.

(32) **Test tomorrow**

rúguǒ Lìlì yǐwéi míngtiān yǒu kǎoshì de-huà, tā yīnggāi zài xuéxí
 if Lili yǐwéi tomorrow have test in-that-case, 3sg should PROG study

‘If she thinks there’s a test tomorrow, she should be studying.’ (= (16))

As for negation, negation in dynamic semantics is generally analyzed as a two-step process (Heim 1983, Beaver 2001): first the assertion *A* is added to the context *c* in the usual way (by intersecting them); then the result of that update is subtracted from the original context *c*.

$$(33) \quad c + \neg A =_{def} c - (c + A) \quad \text{Heim (1983)}$$

If *A* has a postsupposition, this postsupposition applies to the intermediate context set, *c + A* (Lauer 2012). In this way, postsuppositions behave differently under negation than presuppositions. A presupposition applies directly to *c*, while a postsupposition applies only to *c + A*.

To see how this analysis applies to *yǐwéi*, imagine that the full, negated sentence ($\neg A$) is $\neg(x \text{ yǐwéi } p)$. To update the context with this sentence, the first step is to calculate $c+(x \text{ yǐwéi } p)$: the worlds in *c* in which *x yǐwéi p* is true. Thanks to the postsupposition of *yǐwéi*, some of these worlds are required to be not-*p* worlds: in some of the counterfactual worlds in which *x* were to believe *p*, *p* is false. These worlds (the *c + A* worlds) are then subtracted from the original context set, leaving only worlds in which it is not true that *x* believes *p*. The postsupposition of *yǐwéi* does not directly apply to these remaining worlds, so its negative bias persists only in a very weak, counterfactual sense (Lauer 2012): if *x* were to believe *p* (in the hypothetical *c + A* worlds), *p* would not necessarily be true.

I argue that the *yǐwéi*’s negative bias disappears under negation precisely because postsuppositions interact with negation in this unique manner. (Lauer 2012 uses these insights to explain the effect of negation on German *irgend-ein* ‘some or other,’ which he also analyzes with a postsupposition, as discussed below in §5.)

(34)–(35) show empirically that *x does not yǐwéi p* can be used both when *p* is false and when *p* is true. To explain these data, we first update the context with ‘Lili *yǐwéi* there’s a

test' [$c + A$]: here, the postsupposition of *yǐwéi* requires that in at least some world where Lili thinks there's a test, she is wrong. Next, we remove all the $c + A$ worlds from the actual context [$c - (c + A)$], leaving only the worlds where it is not the case that Lili thinks there's a test. The worlds where Lili *does* think there's a test [$c + A$] are removed and exist only as counterfactuals. At least one of these counterfactual worlds is one in which there is no test: if Lili were to believe there is a test (which she doesn't), there might not be one. But that is a very weak condition because it is contrary to fact. At least one counterfactual world may satisfy any condition at all, with no consequences for our actual state of affairs: there are counterfactual worlds where unicorns rule Mars.

Thus, under negation, the postsupposition of *yǐwéi* amounts only to a very weak requirement that p might be false in a counterfactual scenario, but that tells us nothing about whether p is true or false here and now: explaining why x *does not* *yǐwéi* p can be used in either case.

(34) **Test tomorrow**

We all know that there is no test tomorrow because we are going on a field trip. Lili also knows that there is no test; that's why she is out seeing a movie tonight instead of studying. Explaining Lili's whereabouts, I say:

Lìlì bìng-bù yǐwéi míngtiān yǒu kǎoshì ér quèshí meiyǒu
Lili definitely-not yǐwéi tomorrow have test and actually not-have

'Lili definitely doesn't think there's a test tomorrow, and there's not.' (= (19))

(35) **Test tomorrow**

We all know that there definitely is a test tomorrow. But Lili has been absent, so she doesn't know about the test; that's why she is out seeing a movie tonight instead of studying. Explaining Lili's whereabouts, I say:

Lìlì bìng-bù yǐwéi míngtiān yǒu kǎoshì dànshì yǒu
Lili definitely-not yǐwéi tomorrow have test but have

'Lili definitely doesn't think there's a test tomorrow, but there is.' (= (20))

In sum (Table 3), the proposed analysis is argued to derive the effects of *yǐwéi* in different contexts.

4.6. *Alternative analyses*

I advocate for this analysis over alternatives which have been proposed for other negatively-biased belief verbs.

***Yǐwéi* as 'falsely believe'** One alternative, proposed by Hsiao (2017) for Southern Min *liah-tsun* and by Anvari, Maldonado, & Soria Ruiz (2019) for Spanish *creer se*, would be to analyze *yǐwéi* to mean 'falsely believe.' To capture the projection behavior of *yǐwéi*,

Context	Effect	Explanation
1st person <i>yǐwéi p</i>	understood as past or hedged	two different ways of reconciling speaker's belief in <i>p</i> with signal that <i>p</i> is not to be taken up
2rd person <i>yǐwéi p</i>	typically rude because it suggests hearer is in error	speaker indicates that <i>p</i> is not to become Common Ground, perhaps because it is wrong/questionable
3rd person <i>yǐwéi p</i>	typically conveys that speaker is skeptical towards <i>p</i>	speaker indicates that <i>p</i> is not to become Common Ground, perhaps because it is wrong/questionable
3rd person <i>yǐwéi p</i> and in fact <i>p</i>	subverts expectations	by signaling that <i>p</i> should not become Common Ground, speaker implicates that <i>p</i> is questionable; then goes on to endorse it
questions, conditionals, <i>might</i>	sense of wrongness persists	postsuppositions project
under negation	sense of wrongness seems to disappear	postsuppositions project differently than presuppositions (Lauer 2012)

Table 3: Effects of *yǐwéi* in different contexts, and how these effects are explained on the proposed analysis.

the falseness of *p* should be formalized not as an at-issue entailment, but rather as some sort of non-at-issue projective meaning – a presupposition, a conventional implicature (see Potts 2007, 2014), or perhaps even a postsupposition. On such an analysis, *yǐwéi* might be represented as in (36): identical to the meaning that I propose above, except that now, not just *some*, but *all* of the worlds in the output context are required to be not-*p* worlds.

- (36) at-issue and postsupposed content of *yǐwéi* on an alternative analysis where *yǐwéi* conveys ‘falsely believe’
- $c + x \text{ } yǐwéi \text{ } p = c + x \text{ believes } p$
 - defined only if $\forall w [w \in (c + x \text{ believes } p) \rightarrow p(w) = 0]$

Proponents would argue that (36) immediately explains the headline example of *yǐwéi*: that *Mom yǐwéi I'm sick* (1) conveys that the speaker is *not* sick, because (36) requires not-*p* (‘I’m not sick’) to be Common Ground following the assertion. But, I argue, (36) fails to capture several important pieces of data (§2):

- Yǐwéi* can be used when the speaker does not commit to *p* being false, but rather expresses ignorance regarding *p* in situations where the belief-holder’s evidence is incomplete (‘I don’t know whether he scored or not, but he *yǐwéi* he did’ (7)); or when the speaker is confused about *p*, but continues to entertain it as a possibility (‘I *yǐwéi* there was a talk today; maybe there is one, but I have the wrong room’ (11)) – not predicted if *yǐwéi* commits the speaker to *p* being false.
- Yǐwéi* can be used when the speaker claims to believe *p* (the hedged suggestion ‘I per-

sonally *yǐwéi* you should do this’ (12)) – not predicted if *yǐwéi* commits the speaker to *p* being false.

3. *Yǐwéi* can be used when the speaker goes on to endorse *p* (‘People *yǐwéi* she’s a billionaire . . . and she actually is’ (9)) – not predicted if *yǐwéi* commits the speaker to *p* being false.

Thus, even though (36) may be a simpler way of handling the prototypical case where *x yǐwéi p* conveys not-*p*, it captures significantly less of the data compared to the analysis that I advocate.

The same critique extends to an alternative analysis (proposed by Kierstead 2013 for Tagalog *akala*) whereby *yǐwéi* conventionally implicates that the speaker ‘doubts’ *p* – not consistent with data showing that the speaker can entertain *p* as a live possibility or go on to endorse it (1–3).

Moreover, if *yǐwéi* meant ‘falsely believe,’ then it should be redundant with the adverb *wù* ‘falsely/wrongly.’ But actually, *wù yǐwéi* ‘falsely *yǐwéi*’ can only be used in cases where the reported belief is actually taken to be false. *Wù yǐwéi* ‘falsely *yǐwéi*’ cannot be used in any of the situations (1–3) where unmodified *yǐwéi* can be used to describe a belief that might be true. Thus, *wù yǐwéi* ‘falsely *yǐwéi*’ and *yǐwéi* are not synonymous, and *yǐwéi* does not entail falsity as part of its lexical meaning.

***Yǐwéi* as non-commitment** Another alternative takes inspiration from Potts (2005) (Chapter 5) on the German Konjunktiv I verb ending (von Stechow 2002). This verb ending distances the speaker from an embedded proposition (37) and is commonly used in detached journalism, roughly paraphrased as ‘purportedly.’

- (37) Sheila behauptet dass sie krank sei
Sheila maintains that she sick be.KONJ
‘Sheila maintains that she is (purportedly) sick.’ adapted Potts (2005): 21

Potts (2005) analyzes propositions marked with the Konjunktiv I using a (projecting) conventional implicature that the speaker makes no commitment regarding the truth of that proposition. But while the Konjunktiv I is favored by neutral journalists for describing a third party’s belief without committing to it, *yǐwéi* comes across as polemical rather than neutral in such contexts (§2: (5)), and so requires a different analysis.

***Yǐwéi* with a presupposition rather than a postsupposition** What if the proposed postsupposition of *yǐwéi* were instead just a regular presupposition (38)?

- (38) at-issue and presupposed content of *yǐwéi* on an alternative analysis
- a. $c + x \text{ } yǐwéi \text{ } p = c + x \text{ believes } p$
 - b. defined only if $\exists w \in c \rightarrow p(w) = 0$

(38) seems appealing because it uses a vanilla presupposition rather than a newfangled postsupposition. But (38) cannot capture the data. (38) would allow a discourse like (39) (inspired by Simons 2007, Lauer 2009):

(39) **Lili’s arrival**

- a. **A:** Lìlì dào le ma? / Has Lili arrived?
- b. **B:** Zhangsan yǐwéi tā dào le. / Zhangsan yǐwéi she has.
- c. **A:** #Hén hǎo, wǒ qù gēn tā shuōhuà! / #Great, then I’ll go talk to her!

(39b) is perfectly compatible with (38), because the input context is compatible with the proposition that ‘Lili has not arrived’ (since the matter of Lili’s arrival is unsettled); the presupposition of (38) is satisfied. But if Zhangsan is Lili’s best friend and can be trusted regarding Lili’s whereabouts (if it’s Common Ground that Lili has arrived if Zhangsan believes it), then the utterance (39b) could lead to Lili’s arrival becoming Common Ground on the evidence that Zhangsan believes it (even though the presupposition of (38) requires that this wasn’t Common Ground previously). Nothing in (38) prevents this discourse pattern.

And yet, discourses such as (39) are *not* available for *yǐwéi* (whereas (39) would make sense with *rènwéi* instead). (39b) is not compatible with the assumption that Lili has arrived if Zhangsan believes it, and (39b) cannot license the inference that Lili has arrived on those grounds. (38) cannot explain why. In contrast, my proposed analysis specifically requires that the *output* context of (39b) *remains* compatible with not-*p* – thereby preventing ‘Lili has arrived’ from becoming Common Ground on the evidence that Zhangsan believes it.

***Yǐwéi* as weak belief with a scalar implicature** A final alternative analysis would take *yǐwéi* to simply mean ‘weakly believe’ or ‘have an inkling that,’ in contrast to stronger alternatives such as *rènwéi* ‘think’ and *zhīdào* ‘know.’ Just as *she thinks there’s a test* in English may implicate that she doesn’t *know* there’s a test (Levinson 1983, Hirschberg 1985, Horn 1989, Lauer 2017), the idea would be that *x yǐwéi p* (*x* weakly believes *p*) would implicate *it’s not the case that x rènwéi p* (*x* doesn’t [strongly] believe *p*). *Yǐwéi* would have no postsupposition, just a scalar competition with an alternative analyzed to be stronger.

If *yǐwéi* just meant ‘weakly believe,’ then we would expect first-person *yǐwéi* (just like other verbs of weak or fallible belief in Mandarin, such as *juéde* ‘feel that’) to describe the speaker’s current belief, leaving it unexplained why first-person *yǐwéi* differs from those other verbs in favoring a past interpretation. We would expect third-person *yǐwéi* to describe an agent’s inconclusive or weakly held belief in a neutral manner, and to be usable when the belief is established as true. But as §2 shows, *yǐwéi* strongly suggests that either the belief or the belief-holder cannot be trusted, and is rejected when the belief is assumed to be true – not explained if it is just a verb of weak belief.

In sum, I maintain my analysis that *x yǐwéi p* updates the Common Ground with the information ‘*x* believes *p*,’ and postsuppositionally requires not-*p* to remain a live possibility.

5. Postsuppositions as inference-blockers

On the proposed analysis, *yǐwéi*'s postsupposition blocks a potential inference from x 's belief in p to $p -$ and in doing so, triggers pragmatic reasoning about why the speaker wished to prevent p from becoming Common Ground. While this analysis is argued to capture the data, the central device of postsupposition may not seem intuitive.

5.1. *Unifying (some) postsuppositions in the literature*

At first glance, the phenomena handled in terms of postsuppositions may appear rather miscellaneous. Most commonly, postsuppositions have been used to analyze some (but not all) types of noun phrases in various languages that are characterized as nonspecific or related to free choice (Farkas 2002a,b, Lauer 2009, 2012, Condoravdi 2015) – those that can be roughly paraphrased as ‘some or other.’ Here, the proposed postsupposition ensures that the Common Ground is compatible with multiple different referents for the noun phrase following the update.

Also in the domain of noun phrases, Kuhn (2022) analyzes negative concord (in some varieties of English, *I didn't see nothing*) via a postsupposition requiring no discourse referent for the variable introduced by the word *nothing*. He proposes to unify his treatment of negative concord with a postsupposition-like analysis of definites from Bumford (2017) which requires the cardinality of a definite's referent to be equal to one, and with post-suppositional analyses of distributive markers (*one book each*) from Henderson (2014) and Champollion (2015) requiring multiple distinct referents for *one book*. In all these cases, a postsupposition serves as a backgrounded element of meaning which out-scopes other elements of the sentence to guarantee the cardinality of a noun phrase's referent.

Moving from noun phrases into the realm of numerals and additivity, Brasoveanu (2009, 2012), and Charlow (to appear) use postsuppositions to ensure the correct scope of various numerals in sentences involving multiple plurals. Brasoveanu & Szabolsci (2013) invoke a postsupposition to explain why the presuppositional word *too* (or its equivalent, in various languages) can occur before the information that satisfies its presupposition.

In a different vein, Constant (2012) and Rudin (2018) use postsuppositions to capture inferences associated with marked intonation contours in English: for Constant, rise-fall-rise is argued to convey that the speaker is not willing to assert any alternative proposition that would be consistent and informative in the Common Ground following the assertion uttered with rise-fall-rise intonation; for Rudin, steep rising intonation is argued to prevent the usual outcome that speakers commit to the content of what they say. Cohen & Krifka (2014) use postsupposition to analyze *Jo pets at most three rabbits* so that a speaker's utterance *declines* to commit to Jo petting any number of rabbits greater than three.

Finally, this paper uses a postsupposition to explain why the Mandarin belief verb *yǐwéi* conveys skepticism towards the reported belief, the first use of postsupposition to handle an open-class lexical item rather than a function word or an intonation pattern.

I argue that several of these uses of postsupposition can be grouped together as ways of blocking a potential inference which might result from a neutral, non-postsuppositional

alternative – thereby triggering further pragmatic reasoning about why the speaker wished to prevent that inference. (40) distills the empirical properties of the phenomena that have been analyzed in terms of such inference-blocking postsuppositions, and (41) sketches the theoretical core shared by these analyses, using *yǐwéi* to illustrate this broader category.

- (40) Empirically, an item *b* may be analyzed in terms of an (*yǐwéi*-like) inference-blocking postsupposition if . . .
- a. An utterance using *b* gives rise to strong, context-dependent effects which set it apart from utterances using a salient, seemingly more “neutral” alternative *a*. If it is possible to test for projection (appropriate for lexical items but not utterance-level phenomena such as intonation contours), then these inferences project – but not out of negation.
 - (i) An utterance *x yǐwéi p* gives rise to strong, context-dependent, projecting inferences that *p* is wrong, questionable, or hedged which set it apart from utterances using the neutral alternative *rènwéi*.
 - b. Unlike the case of scalar implicatures, there is *no* asymmetric entailment relation whereby *a* entails *b* but *b* doesn’t entail *a*.
 - (i) *x rènwéi p* certainly does not asymmetrically entail *x yǐwéi p*.
 - c. An utterance using the neutral alternative *a* (potentially) gives rise to an inference or effect that is not available when *b* is used instead.
 - (i) An utterance *x rènwéi p* may give rise to the potential inference that *p* is true on the grounds that *x* believes it – not available when *yǐwéi* is used instead.
- (41) Theoretically, an item *b* is analyzed in terms of an (*yǐwéi*-like) inference-blocking postsupposition if . . .
- a. *b* (argued to have a postsupposition) competes with a salient neutral alternative *a* (argued to have no postsupposition).
 - (i) *Yǐwéi* (argued to have a postsupposition) competes with a salient neutral alternative *rènwéi* (argued to have no postsupposition).
 - b. The postsupposition attributed to *b* semantically blocks a (potential) inference or effect which may result from an alternative utterance using the neutral alternative *a*.
 - (i) The postsupposition attributed to *yǐwéi* is used to block the potential inference – from *x*’s belief in *p* to *p* – which may result from an utterance of *x rènwéi p*.
 - c. A speaker’s choice to use the marked, postsuppositional form *b* is argued to trigger context-sensitive pragmatic reasoning about why the speaker wished to prevent the effects that might have arisen from the neutral alternative *a*.
 - (i) A speaker’s choice to use *x yǐwéi p* is argued to trigger pragmatic inferences (that *p* is false, questionable, unsupported, or hedged) arising from the speaker’s choice to signal that *p* should not be taken up on the grounds that *x* believes it.

Thus, the present analysis of *yǐwéi* is used to unify a class of proposed postsuppositions in the literature, and to provide a map for future researchers to identify further phenomena which may be fruitfully analyzed in the same terms.

For concreteness, please consider the English data of Lauer (2009) as an example of the use of postsuppositions to handle nonspecific/free choice noun phrases built from *wh + ever* (also proposed, for other data and languages, by Farkas 2002a,b, Lauer 2012, Condoravdi 2015). The *wh+ever* form¹² (42b) evokes a (projecting) sense of indeterminacy over possible referents for the thing Arlo is cooking (perhaps he is cooking stew, perhaps he is cooking stir-fry), not shared by its alternative (42a).

- (42) a. What Arlo is cooking smells delicious.
(allows that the might be a unique/known referent for the thing Arlo is cooking)
- b. Whatever Arlo is cooking smells delicious.
(signals that the referent cannot be uniquely identified – triggering further inferences about why not)

While some researchers have analyzed *wh+ever* constructions to explicitly encode speaker ignorance and/or indifference in a modal framework (Dayal 1997, von Stechow 2000), Lauer’s analysis aims to derive those effects as secondary. Following the theoretical frame (41), (42a) is argued to place no requirements on its input or output contexts, while (42b) is argued to have a postsupposition requiring the output context to be compatible with multiple distinct referents (variable assignments) for the thing that Arlo is cooking. In the right context, an utterance of (42a) might lead hearers to believe that this noun phrase has a unique agreed-upon referent in the context. The postsupposition of (42b) explicitly prevents this outcome. Moreover, in choosing the postsuppositional variant (42b) over the neutral alternative, the speaker triggers further reasoning about why they explicitly do not want the hearer to infer that there is a unique referent for the thing Arlo is cooking – because they don’t know, don’t care, or don’t want to say. Lauer’s analysis of *wh+ever* thus follows the same framework as the proposed analysis of *yǐwéi*.

The same picture also encompasses the postsupposition proposed by Constant (2012) for the English rise-fall-rise intonation pattern. Again following the empirical pattern set out in (40), Constant observes that (43b) (with rise-fall-rise intonation) signals an elusive sense of finality. While (43a) leaves it open whether the speaker can go on to list more people beyond Alex who liked the film, Constant argues, (43b) signals that the speaker has answered the question as fully as they can, and cannot (or do not wish to) add more.

- (43) **A:** Who liked the film?
- a. **B:** Alex_{neutral} liked it.
(allows that the speaker might go on to identify more people who liked it)
- b. **B:** Alex_{RFR} liked it.

¹²Lauer sets aside free choice items built from *any*, which have different properties from those built from *-ever*.

(signals that the speaker cannot identify more people who liked it – triggering further inferences about why not)

Following the theoretical frame of inference-blocking postsuppositions (41), Constant gives (43b) a postsupposition requiring that no alternative answer to the question would be assertable (true and informative) in the output context following the assertion marked with rise-fall-rise. (43a) is said to place no restrictions on its input or output context. Hearers may infer from (43a) that the speaker has more information to add. The postsupposition of (43b) explicitly prevents this outcome. Moreover, the speaker's choice to use (43b) over its neutral alternative cues pragmatic reasoning about why the speaker cannot say more – because they don't know who else liked the film, because they know that no one else liked it, because they have some reason to be cagey about other people's opinions. These context-sensitive inferences parallel the ones triggered by a speaker's choice to use *yǎwéi* over *rènwéi*.

Also in the realm of intonation, Rudin (2018) uses what is essentially a postsupposition to analyze steep rising (question-like) intonation in English in both declaratives and imperatives (*it's fun?*, *don't go?*), which evokes a sense of uncertainty or open-endedness. Following the empirical pattern (40), steep-rising intonation gives rise to strong, context-dependent effects distinguishing it from a neutral alternative (falling intonation).

Rudin's analysis also follows the theoretical frame (41). For Rudin, falling (unmarked) declaratives and imperatives commit the speaker to the content of the utterance: the doxastic truth of a declarative, the teleological desirability of the addressee carrying out the action described by an imperative. With steep rising intonation, the speaker does not take on these commitments. Using a discourse model inspired by Farkas & Bruce (2010), Rudin takes all other discourse effects to remain consistent with either sort of intonation; but argues that steep rising intonation requires the speaker's commitments to remain as they were prior to the utterance.

- (44) a. It's fun.
(speaker commits doxastically to this proposition)
b. It's fun?
(speaker does not commit doxastically to this proposition)
- (45) a. Don't go.
(speaker commits teleologically to a preference for the addressee not going)
b. Don't go?
(speaker does not commit teleologically to a preference for the addressee not going)

Although Rudin does not use the word *postsupposition*, his formalization has the effect of one: steep rising intonation restricts the output context so that the speaker does not end up committed to the content of their utterance. Like *yǎwéi*, steep rising intonation competes with an unmarked alternative (falling intonation), and explicitly prevents its outcome (speaker commitment). The speaker's choice to use a steep rise triggers further inferences

about why they did so: because they want to hedge their utterance, elicit confirmation, reject an utterance previously made by their interlocutor, and so on – giving rise to the rich, diverse pragmatic and social inferences associated with such intonation.

Finally, the same picture of postsupposition further extends to the work of Cohen & Krifka (2014) on superlative quantifiers (*at most*). Following the empirical pattern (40), superlative quantifiers give rise to strong, context-dependent inferences – that Jo does pet at least some rabbits (perhaps two?), the speaker doesn't know the number of rabbits that Jo pets, or (if ?? is taken as habitual) that the number varies across situations but is never greater than three. These effects distinguish ?? from a neutral alternative such as the regular numeral in ??.

Parallel to (41), the analysis of Cohen & Krifka (2014) restricts the output context so that, following the assertion, the speaker is *not* committed to Jo petting any number of rabbits greater than three. As always, explicitly preventing one inference gives rise to other inferences about why the speaker chose to make such a move. Why does the speaker refuse to commit to Jo having pet four or more rabbits, and why the speaker is still open to the commitment that Jo pet three rabbits or fewer? The effects of ?? are explained using such reasoning. Therefore, Cohen and Krifka's proposed postsupposition parallels the one proposed here.

Although there are many differences between belief reports, nonspecificity, intonation, and superlatives, these phenomena all make use of postsupposition in similar ways (41). This understanding helps to profile further linguistic items to be analyzed postsuppositionally (40). Such an item will trigger inferences which seem semantic in their strength (because they are grounded in a conventional semantic postsupposition) but pragmatic in their context-sensitivity (because they arise pragmatically as hearers reason about why the speaker wanted to prevent some potential discourse effect of its neutral alternative); and will only be understood in light of the inferences which might result from its alternative.

Some other proposed postsuppositions from the literature do not fit this picture, particularly those that do not subvert potential inferences arising from a neutral alternative in the way as *yǐwéi*, free-choice items, and rise-fall-rise are argued to. It may not be possible to connect all of these uses of postsuppositions, but we have made progress by unifying some of them.

5.2. *Comparing presupposition and postsupposition*

To wrap up this section, it is worth briefly comparing and contrasting presuppositions and postsuppositions.

Of course, presuppositions are a moving target for comparison, since the phenomena grouped under this heading are empirically heterogeneous and theoretically debated (Stalnaker 1974, Kadmon 2001, Abusch 2002, Simons 2006, Beaver & Geurts 2014, Potts 2014, Karttunen 2016, Tonhauser, Beaver, & Degen 2018). But ideally, both presuppositions and postsuppositions should have a recognizable empirical character; an intuitive function; and a theoretical explanation.

Presuppositions are empirically recognized by their projection behavior and by the intuition that they flag information that the speaker views as old and/or uncontroversial. Their function is to signal such information. Theoretically, presuppositions may stem from different sources; some may be derived purely from shared assumptions and conversational reasoning (Stalnaker 1974, Kadmon 2001, Abusch 2002, Simons 2006), while others may be triggered by particular lexical items or constructions. If a lexical item is taken to conventionally trigger a presupposition, it may be said to place a definedness condition on its input context requiring that its presupposition is (or becomes) Common Ground, meant to explain its empirical properties and its discourse function.

Postsuppositions, by comparison, should be empirically recognized by the properties laid out above in (40): their competition with a neutral alternative, the fact that they trigger (projecting) inferences which seem semantic in their strength but pragmatic in their context-sensitivity. Their function is to prevent a effect which might have arisen from a neutral alternative, thereby triggering further context-sensitive inferences about why the speaker wanted to do that.

As strategies for preventing a potential effect of an utterance, it is notable that all of the proposed postsuppositions I discuss restrict *changes* to the context rather than restricting the context itself. Whereas *zhīdào* ‘know’ (with a presupposition) is argued to restrict the context to contain only *p* worlds, *yǐwéi* (with a postsupposition) is argued to restrict the context from *changing* to include only *p* worlds. That is why the timing of postsuppositions is so significant, and why they must inherently make reference to the output context rather than the input context.

Theoretically, postsuppositions are attributed to the conventional meaning of words or intonation patterns, as definedness conditions on output contexts which prevent certain changes to the context (on this framing, postsuppositions differ from presuppositions in that no postsupposition could be derived pragmatically). This characterization is argued to explain their empirical properties and their discourse function.

Stepping back, this section has situated the proposed analysis of *yǐwéi* within a larger class of postsuppositions. While there are many open questions about this relatively new-fangled device, this paper has offered the seeds of a unified, intuitive conception which may help to give postsuppositions a home in semantic and pragmatic theory.

6. Conclusion

Using Mandarin *yǐwéi* as a case study, this paper has argued for a unified and intuitive understanding of postsuppositions as ways of blocking potential contextual inferences.

As for *yǐwéi*, its negative bias is semantically grounded in a postsupposition requiring the output context to be compatible with not-*p*; and pragmatically derived as hearers reason about why the speaker chose to explicitly signal that *x*'s belief is not to be taken up. This exploration engages the complex reasoning involved in deciding what to think about what other people think, and the linguistic resources used to guide it. It is well known that some belief reports are factive, conveying that the speaker endorses the reported belief, while

others are nonfactive, silent on what the speaker thinks about it. Enriching this picture, *yǐwéi* exemplifies a relatively less-studied class of strategies for reporting beliefs that the speaker views with skepticism.

Other negatively biased belief verbs have also recently entered the discussion – Kierstead (2013) on Tagalog *akala*, Hsiao (2017) on Southern Min *lih-tsun*, Anvari, Maldonado, & Soria Ruiz (2019) on Spanish *creer se*. These words appear quite similar to *yǐwéi*, although they are analyzed to mean ‘believe falsely.’ In the philosophy literature, Holton (2017) claims that no lexical item should have that meaning, so I look forward to further discussion of these words. They will not necessarily be identical to *yǐwéi*. It seems likely that different languages may use different resources to achieve a shared discourse function (e.g., Deal 2011): English *be under the impression that* roughly paraphrases *yǐwéi*, but could perhaps be analyzed using a Gricean manner implicature based on the speaker’s choice to describe the belief as an *impression*, rather than using a postsupposition.

Turning to postsuppositions, these are cast as ways of semantically blocking one inference, thereby pragmatically triggering others. Just as the postsupposition of *yǐwéi* can only be understood in light of the inference it prevents, which in turn can only be understood within the complex calculations involved in reported beliefs, postsuppositions in general can only be understood in the context of the inferences that would arise without them. There is a mutually illuminating symbiosis between the study of postsuppositional items, their non-postsuppositional alternatives, and the discourse effects of each one.

Lelia Glass
School of Modern Languages
Georgia Institute of Technology
613 Cherry Street NW
Atlanta, GA 30313
United States of America
lelia.glass@modlangs.gatech.edu

References

- ABUSCH, D. 2002. Lexical alternatives as a source of pragmatic presuppositions. B. Jackson, ed., *Proceedings of SALT (Semantics and Linguistic Theory) 12*, 1–19.
- ABUSCH, D. 2010. Presupposition triggering from alternatives. *Journal of Semantics* 27(1):37–80.
- ALTSHULER, D. & SCHWARZSCHILD, R. 2013. Moment of change, cessation implicatures and simultaneous readings. E. Chemla, V. Homer, & G. Winterstein, eds., *Proceedings of Sinn und Bedeutung 17*. Paris: École Normale Supérieure.
- ANAND, P. & HACQUARD, V. 2014. Factivity, belief and discourse. L. Crnić & U. Sauerland, eds., *The art and craft of semantics: A Festschrift for Irene Heim*, 69–90. Cambridge: MIT Working Papers in Linguistics.
- ANVARI, A., MALDONADO, M., & SORIA RUIZ, A. 2019. The puzzle of ‘reflexive belief’ construction in Spanish. M. Espinal, E. Castroviejo, M. Leonetti, L. McNally, & C. Real-Puigdollers, eds., *Proceedings of Sinn und Bedeutung 23*. Cerdanyola del Vallès: Universitat Autònoma de Barcelona.
- BARTSCH, R. 1973. ‘Negative transportation’ gibt es nicht [‘Negative transportation’ does not exist]. *Linguistische Berichte* 27:1–7.
- BEAVER, D. I. 2001. *Presupposition and assertion in dynamic semantics*. Stanford: CSLI (Center for the Study of Language and Information) Publications.
- BEAVER, D. I. 2010. Have you noticed that your belly button lint colour is related to the colour of your clothing? R. Bäuerle, U. Reyle, & T. E. Zimmermann, eds., *Presuppositions and discourse: Essays offered to Hans Kamp*, 65–99. Bingley: Emerald Group Publishing Limited.
- BEAVER, D. I. & GEURTS, B. 2014. Presupposition. E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2014 edition.

- BEVER, T. G. 1970. The cognitive basis for linguistic structures. J. R. Hayes, ed., *Cognition and the development of language*, 279–362. New York: Wiley.
- BRASOVEANU, A. 2009. Modified numerals as post-suppositions. M. Aloni, H. Bastiaanse, T. de Jager, & K. Schulz, eds., *Logic, language and meaning: Proceedings of the Amsterdam Colloquium 17*, 203–212. Berlin: Springer.
- BRASOVEANU, A. 2012. Modified numerals as post-suppositions. *Journal of Semantics* 30(2):155–209.
- BRASOVEANU, A. & SZABOLSCI, A. 2013. Presuppositional *too*, postsuppositional *too*. M. Aloni, M. Franke, & F. Roelofsen, eds., *The dynamic, inquisitive, and visionary life of ϕ , $?\phi$, and $\diamond\phi$: A festschrift for Jeroen Groenendijk, Martin Stockhof, and Frank Veltman*, 55–64. Amsterdam: Institute for Logic, Language, and Computation.
- BUMFORD, D. 2017. Split-scope definites: Relative superlatives and Haddock descriptions. *Linguistics and Philosophy* 40(6):549–593.
- CHAMPOLLION, L. 2015. Every boy bought two sausages each: Distributivity and dependent numerals. U. Steindl, T. Borer, H. Fang, A. G. Pardo, P. Guekguezian, B. Hsu, C. O’Hara, & I. C. Ouyang, eds., *Proceedings of the West Coast Conference on Formal Linguistics (WCCFL)*, volume 32, 103–110. Somerville: Cascadilla Proceedings Project.
- CHARLOW, S. to appear. Post-suppositions and semantic theory. *Journal of Semantics* .
- CHEMLA, E. 2008. An epistemic step for anti-presuppositions. *Journal of Semantics* 25(2):141–173.
- COHEN, A. & KRIFKA, M. 2014. Superlative quantifiers and meta-speech acts. *Linguistics and Philosophy* 37(1):41–90.
- CONDORAVDI, C. 2015. Ignorance, indifference, and individuation with *wh-ever*. L. Alonso-Ovalle & P. Menéndez-Benito, eds., *Epistemic indefinites: Exploring modality beyond the verbal domain*, 213–243. Oxford: Oxford University Press.
- CONSTANT, N. 2012. English rise-fall-rise: A study in the semantics and pragmatics of intonation. *Linguistics and Philosophy* 35(5):407–442.
- CRONE, P. 2018. Assertions of clarity and raising awareness. *Journal of Semantics* 36(1):53–97.
- DAYAL, V. 1997. Free relatives and ‘ever’: Identity and free choice readings. A. Lawson, ed., *Proceedings of SALT (Semantics and Linguistic Theory)* 7, 99–116.
- DEAL, A. R. 2011. Modals without scales. *Language* 87(3):559–585.

- FARKAS, D. & BRUCE, K. 2010. On reacting to assertions and polar questions. *Journal of Semantics* 27:81–118.
- FARKAS, D. F. 2002a. Extreme non-specificity in Romanian. C. Beyssade, ed., *Romance languages and linguistic theory*, 127–151. Amsterdam: John Benjamins.
- FARKAS, D. F. 2002b. Varieties of indefinites. B. Jackson, ed., *Semantics and Linguistic Theory (SALT)*, volume 12, 59–83. Ithaca: Cornell Linguistics Club Publications.
- VON FINTEL, K. 2000. Whatever. B. Jackson & T. Matthews, eds., *Proceedings of SALT (Semantics and Linguistic Theory) 10*, 27–40.
- VON FINTEL, K. 2008. What is presupposition accommodation, again? *Philosophical perspectives* 22(1):137–170.
- GRICE, H. P. 1989. Logic and conversation. *Studies in the way of words*. Cambridge: Harvard University Press. Originally published as Grice (1975), in *Syntax and semantics* (edited by Peter Cole and Jerry L. Morgan), volume 3: Speech acts, 41–58. New York: Academic Press.
- GROENENDIJK, J. & STOKHOF, M. 1984. *Studies on the semantics of questions and the pragmatics of answers*. Ph.D. thesis, University of Amsterdam, Amsterdam.
- HEIM, I. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst.
- HEIM, I. 1983. On the projection problem for presuppositions. M. Barlow, D. P. Flickinger, & M. T. Westcoat, eds., *Proceedings of the West Coast Conference on Formal Linguistics (WCCFL)*, volume 2, 114–125. Stanford: Stanford Linguistics Association.
- HEIM, I. 1991. Artikel und definitheit. A. von Stechow & D. Wunderlich, eds., *Handbuch der Semantik*, 487–535. Berlin: de Gruyter.
- HEIM, I. 1992. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics* 9(3):183–221.
- HENDERSON, R. 2014. Dependent indefinites and their post-suppositions. *Semantics and Pragmatics* 7(6):1–58.
- HINTIKKA, J. 1962. *Knowledge and belief*. Ithaca: Cornell University Press.
- HIRSCHBERG, J. 1985. *A theory of scalar implicature*. Ph.D. thesis, University of Pennsylvania.
- HOLTON, R. 2017. Facts, factives, and contrafactuals. *Aristotelian Society Supplementary Volume* 91(1):245–266.

- HORN, L. 1971. Negative transportation: Unsafe at any speed? *Chicago Linguistic Society (CLS)*, volume 7, 120–133. Chicago: Chicago Linguistic Society.
- HORN, L. 1978. Remarks on neg-raising. P. Cole, ed., *Pragmatics*, volume 9 of *Syntax and semantics*, 129–220. New York: Academic Press.
- HORN, L. 1989. *A natural history of negation*. Chicago: University of Chicago Press.
- HORTON, W. S. & GERRIG, R. J. 2005. Conversational common ground and memory processes in language production. *Discourse Processes* 40:1–35.
- HSIAO, P.-Y. K. 2017. On counterfactual attitudes: A case study of Taiwanese Southern Min. *Lingua Sinica* 3(4).
- HUANG, C.-T. J. 1988. Wǒ pǎo de kuài and Chinese phrase structure. *Language* 64(2):274–311.
- HUNTER, J. 2016. Reports in discourse. *Dialogue & Discourse* 7(4):1–35.
- JIAO, L., KUBLER, C. C., & ZHANG, W. 2013. *500 common Chinese idioms: An annotated frequency dictionary*. Routledge.
- KADMON, N. 2001. *Formal pragmatics: Semantics, pragmatics, presupposition, and focus*. Oxford, U.K.: Blackwell.
- KARTTUNEN, L. 1973. Presuppositions of compound sentences. *Linguistic Inquiry* 4(3):169–193.
- KARTTUNEN, L. 1974. Presupposition and linguistic context. *Theoretical Linguistics* 1(1-3):181–194.
- KARTTUNEN, L. 2016. Presupposition: What went wrong? *Proceedings of Semantics and Linguistic Theory (SALT) 26*, volume 26, 705–731.
- KIERSTEAD, G. 2013. Shifted indexicals and conventional implicature: Tagalog *akala* ‘falsely believe’. Talk presented at SALT 23, University of California, Santa Cruz.
- KIPARSKY, P. & KIPARSKY, C. 1970. Fact. M. Bierwisch & K. E. Heidolph, eds., *Progress in Linguistics*, 345–369. The Hague: Mouton. Reprinted in L. Jakobovits and D. Steinberg (eds.), 1976: *Semantics: An Interdisciplinary Reader*. Cambridge: Cambridge University Press.
- KLECHA, P. 2014. *Bridging the divide: Scalarity and modality*. Ph.D. thesis, University of Chicago.
- KUHN, J. 2022. The dynamics of negative concord. *Linguistics and Philosophy* 45:1–46.

- LANGENDOEN, D. T. & SAVIN, H. 1971. The projection problem for presuppositions. C. Fillmore & D. T. Langendoen, eds., *Studies in Linguistic Semantics*, 373–388. New York: Holt, Rinehart, and Winston.
- LAUER, S. 2009. Free relatives with -ever: Meaning and use. Manuscript, Stanford University.
- LAUER, S. 2012. Some news about *irgend-* and *algún*. Manuscript, Stanford University.
- LAUER, S. 2017. *I believe* in a ranking-theoretic analysis of *believe*. A. Cremers, T. van Gessel, & F. Roelofsen, eds., *Proceedings of the Amsterdam Colloquium*, 335–344. Amsterdam: ILLC.
- LEE, K., OLSON, D. R., & TORANCE, N. 1999. Chinese children’s understanding of false beliefs: The role of language. *Journal of Child Language* 26(1):1–21.
- LEVINSON, S. C. 1983. *Pragmatics*. Cambridge: Cambridge University Press.
- LEVINSON, S. C. 2000. *Presumptive meanings: The theory of generalized conversational implicature*. Cambridge: MIT Press.
- LEWIS, D. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 3(8):339–359.
- LEWIS, S., HACQUARD, V., & LIDZ, J. 2017. ‘Think’ pragmatically: Children’s interpretation of belief reports. *Language Learning and Development* 13(4):395–417.
- LIN, J.-W. 2003. Temporal reference in Mandarin Chinese. *Journal of East Asian Linguistics* 12(3):259–311.
- LIN, J.-W. 2006. Time in a language without tense: The case of Chinese. *Journal of Semantics* 23:1–53.
- LÜ, S. 1999. *Xiàndài Hànyǔ Bābái Cì [800 Words in Modern Chinese]*. Beijing: Shāngwù yìn shūguǎn [Commercial Press].
- DE MARNEFFE, M.-C., MANNING, C., & POTTS, C. 2012. Did it happen? The pragmatic complexity of veridicality assessment. *Computational Linguistics* 38(2):301–333.
- MATTHEWSON, L. 2004. On the methodology of semantic fieldwork. *International Journal of American Linguistics* 70(4):369–415.
- MOORE, G. 1993. Moore’s paradox. T. Baldwin, ed., *G. E. Moore: Selected Writings*, 207–212. London: Routledge.
- ÖZYILDIZ, D. 2017. Attitude reports with and without true belief. D. Burgdorf, J. Col-lard, S. Maspong, & B. Stefánsdóttir, eds., *Semantics and Linguistic Theory (SALT)*, volume 27. Washington, D.C.: Linguistic Society of America.

- PERCUS, O. 2006. Anti-presuppositions. A. Ueyama, ed., *Theoretical and empirical studies of reference and anaphora: Toward the establishment of generative grammar as an empirical science. Report of the Grant-in-Aid for Scientific Research (B), Project No. 15320052*, 52–73. Japan Society for the Promotion of Science.
- POPP, M.-L. 2016. Neg-raising in crosslinguistic perspective. Master's Thesis, University of Leipzig.
- POTTS, C. 2005. *The logic of conventional implicatures*. Oxford Studies in Theoretical Linguistics. Oxford: Oxford University Press.
- POTTS, C. 2007. Into the conventional-implicature dimension. *Philosophy Compass* 2(4):665–679.
- POTTS, C. 2014. Presupposition and implicature. S. Lappin & C. Fox, eds., *Handbook of Contemporary Semantics*, 168–202. Wiley-Blackwell.
- RAWLINS, K. 2010. Conversational backoff. N. Li & D. Lutz, eds., *Proceedings of SALT (Semantics and Linguistic Theory) 20*, 347–365.
- ROBERTS, C. 2012. Information structure: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5.
- RUDIN, D. 2018. *Rising above commitment*. Ph.D. thesis, University of California, Santa Cruz, Santa Cruz, CA.
- SCHLENKER, P. 2012. Maximize Presupposition and Gricean reasoning. *Natural Language Semantics* 20(4):391–429.
- SIMONS, M. 2006. Foundational issues in presupposition. *Philosophy Compass* 1(4):357–372.
- SIMONS, M. 2007. Observations on embedding verbs, evidentiality, and presupposition. *Lingua* 117(6):1034–1056.
- SIMONS, M., TONHAUSER, J., BEAVER, D., & ROBERTS, C. 2010. What projects and why. N. Li & D. Lutz, eds., *Semantics and Linguistic Theory (SALT)*, volume 20, 309–327. Ithaca: Cornell Linguistics Club Publications.
- SMITH, C. S. & ERBAUGH, M. S. 2005. Temporal interpretation in Mandarin Chinese. *Linguistics* 43:713–756.
- STALNAKER, R. 1974. Pragmatic presuppositions. M. K. Munitz & P. Unger, eds., *Semantics and philosophy*, 197–213. New York: New York University Press. Reprinted in Stalnaker (1999). *Context and content: Essays on intentionality in speech and thought*. Oxford: Oxford University Press.

- STALNAKER, R. 1979. Assertion. P. Cole, ed., *Pragmatics*, volume 9 of *Syntax and Semantics*, 315–332. New York: Academic Press. Reprinted in Stalnaker (1999). *Context and content: Essays on intentionality in speech and thought*. Oxford: Oxford University Press.
- STALNAKER, R. 2002. Common ground. *Linguistics and Philosophy* 25:701–721.
- VON STECHOW, A. 2002. Binding by verbs: Tense, person and mood under attitudes. M. Kadowaki & S. Kawahara, eds., *Proceedings of North East Linguistics Society (NELS)* 33, 379–403.
- STRAWSON, P. F. 1950. On referring. *Mind* 59(235):320–334.
- TONHAUSER, J., BEAVER, D. I., & DEGEN, J. 2018. How projective is projective content? Gradience in projectivity and at-issueness. *Journal of Semantics* 35:495–542.
- URMSON, J. O. 1952. Parenthetical verbs. *Mind* 61(244):480–496.
- WALKER, M. A. 1996. Inferring acceptance and rejection in dialogue by default rules of inference. *Language and Speech* 39(2):265–304.
- WITTGENSTEIN, L. 1953. *Philosophical Investigations*. Oxford: Basil Blackwell.
- XIANG, Y. 2013. Neg-raising and aspect: Evidence from Mandarin. *Proceedings of the Penn Linguistics Colloquium*, volume 19, 257–266. Philadelphia: University of Pennsylvania Working Papers in Linguistics.
- XIAO, R. & MCENERY, T. 2008. Negation in Chinese: A corpus-based study. *Journal of Chinese Linguistics* 36(2):274–320.