1. Intro   Using original elicited data, I analyze why Mandarin *yǐwéi* conveys that the belief it embeds is false/questionable. I argue that *x yǐwéi p* has an at-issue meaning of *x believes p* and a projecting requirement that the Common Ground (CG) remains compatible with not-p **after** it's updated with *x believes p*: a **postsupposition** (Farkas '02, Brasoveanu '09, Lauer '09: mirror of a presupposition). Speaker's choice to use *yǐwéi* thus triggers context-dependent pragmatic inferences: that Speaker thinks p is false/dubious and/or questions x's evidence.

2. Data   Belief reports trigger complex reasoning (e.g., Simons '07, Anand & Hacquard '14): Is the content of the belief credible? Is the belief-holder reliable? Why did Speaker report a belief vs. make a direct claim? Belief verbs help guide this process: in Mandarin, factive *zhīdào* 'know', neutral *rènwéi* 'think' (1), or the negatively biased ***yǐwéi*** (2), which signals skepticism.

(1)   Māma rènwéi wǒ bìng le
       Mother think I   sick ASP
       'Mom thinks I'm sick'
       (leaves open whether sick or not)

(2)   **Māma yǐwéi wǒ bìng le**
       Mother yǐwéi I sick ASP
       '**Mom has the impression that I'm sick**'
       (suggests: not sick)

Discourse-initially, (2) implies that Speaker is not sick; when it's already CG that Speaker is *not* sick, (2) reinforces that. If it's CG that Speaker *is* sick, (2) is nonsense.

   Yǐwéi can also be used in a context in which Speaker does not know whether the reported belief is true or not, but wants to signal that the belief is not fully evidenced. (3) could be used if an American football player catches the ball and begins visibly celebrating – oblivious that the officials are congregating to debate whether the catch counted. (3) still conveys skepticism, but the effect is no longer to convey that the belief is false – just insufficiently informed. (But if we just see the athlete catch the ball on the sideline and begin celebrating – with no reason to question his belief – then (3) is rejected: 'if you don't know, why say he's wrong?')

(3)   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.'

The inference that the belief is false behaves like a conversational implicature (Grice 1975), in that it can be reinforced without redundancy, and cancelled without contradiction (4). The latter 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, requiring a specific rhetorical goal: 'I have a friend who invented a famous app. People *yǐwéi* she's a billionaire – and she actually is, but she never made money on that app. Her fortune is inherited.'

(4)   rénmén yǐwéi tā shì yìwànfùwēng ... { dànshì tā bú shì/ér tā díquè shì }
       people yǐwéi 3sg be billionaire   ... { but   3sg not be/and 3sg actually be }
       'People are under the impression that she's a billionaire ... { but she's not/and she actually is} '

1st-person *yǐwéi* has 2 uses, each reconciling in a different way Speaker's stated belief in p with the skepticism of *yǐwéi*. Most commonly, it conveys that Speaker previously believed p but now rejects/questions it (5) (since Mandarin tense is inferred pragmatically). Less commonly, 1st-person *yǐwéi* signals that Speaker currently believes p but invites hearers to disagree (6). (With *rènwéi* instead of *yǐwéi*, (5) would report Speaker's current belief; (6) would be less hedged.)

(5)   wǒ yǐwéi jīntiān yǒu ge jiǎngzuò
       I   yǐwéi today have CL talk
       'I used to think there was a talk today.'

(6)   wǒ gèrén   yǐwéi nǐ yīnggāi qù
       I personally yǐwéi you should go
       'Personally, I would think you should go.'

The negative bias of *yǐwéi* **projects** (Chierchia & McConnell-Ginet 1990) out of conditionals, possibility modals, questions, etc (though can't directly combine with negation). **In sum**, *yǐwéi* conveys projecting, context-dependent skepticism towards the belief p: p may be false, questionable, hedged, or improperly evidenced. What semantics/pragmatics explains these facts?

3. Analysis Assume an update semantics (Stalnaker '79), where an accepted assertion of sentence S shrinks the CG to include only the worlds where S is true. Here, a **presupposition** is a definedness condition on the CG prior to the assertion: *x knows p* updates the CG with *x believes p* and is only defined if p is already CG (7). We can also place a definedness condition on the CG *after* the assertion – a **postsupposition**. I propose that *yǐwéi* updates the CG with *x believes p* and is only defined if, **after** that update, it is compatible with the CG that not-p (8). Informally, no matter how reliable x is/no matter how plausible p is, we are not to take it up. If *yǐwéi* only *pre*supposed that the CG is compatible with not-p, then p could be taken up after the assertion on the grounds that x believes it (even if it was unknown before that); but the *post*-supposition prevents that scenario. Informally, *yǐwéi* conveys, 'x believes p, but don't take their word for it'. Finally, I take *x rènwéi* 'thinks' p to just update the CG with *x believes p*, with no definedness condition; so it can be used in a wider variety of contexts.

(7) update effect and presupposition of *know*
c + *x knows p* = c + *x believes p*
defined iff $\forall w[w \in c \to p(w) = 1]$

(8) update effect and postsupposition of *yǐwéi*
c + *x yiwei p* = c + *x believes p*
defined iff $\exists w \in c + x\ believes\ p: p(w) = 0$

Since *yǐwéi* signals that p can't become CG, it cannot be used if p were already CG prior to the assertion; nor if it were CG that *if x believes p, then p*. As a result, a speaker's choice to use *yǐwéi* triggers pragmatic inferences that p is false/questionable and/or that x is uninformed. E.g., in (2), Speaker signals that Mom's belief (that they're sick) should not be taken up. Speaker presumably has an opinion whether they're sick or not ($p \vee \neg p$); by post-suppositionally signaling $\Diamond \neg p$, *yǐwéi* comes to convey $\neg p$: that Speaker is **not** sick. The inference that p is false is therefore derived pragmatically from the combination of *yǐwéi*'s semantics with background assumptions about whether Speaker has an opinion re: p. As a pragmatic inference, it's not surprising that it can be reinforced and canceled (4).

When Speaker claims not to have an opinion whether p (3), the use of *yǐwéi* (over the more common, neutral *rènwéi* 'think') signals that the CG can't contain *if x believes p, then p*, thus questioning the athlete's reasoning. As for the 'past' 1st-person use (5), Speaker's choice of *yǐwéi* (perhaps combined with an assumption that they have an opinion whether p) conveys that Speaker currently questions or disbelieves p: sensible if (5) describes their **past** belief. In (6), Speaker believes p 'you should go', but explicitly prevents it from becoming CG in case the hearer disagrees (hedged). As non-at-issue definedness conditions, postsuppositions **project** much like presuppositions (Lauer '12), explaining why *yǐwéi*'s effect persists in projection contexts. The analysis (8) thus explains its context-dependent negative bias.

4. Significance Postsuppositions have been used for seemingly heterogeneous phenomena – a.o., the scope of numerals (Brasoveanu), nonspecificity (Farkas, Lauer), particular intonation (Constant '12), and now *yǐwéi*. I propose that *yǐwéi*'s postsupposition fits into a larger class of those used to prevent pragmatic inferences that may otherwise arise from the main assertion. *Yǐwéi* postsuppositionally prevents a potential pragmatic inference from *x believes p* to p, just as Farkas/Lauer use a postsuppositional (in)definite to prevent a potential inference that the speaker can identify its referent. Thus, I offer a unified, intuitive understanding of (some) uses of this device. **More broadly**, this paper engages the complex reasoning involved in deciding what to think about what others think, and the cross-linguistic resources used to guide it. Some belief reports are known to be factive, conveying that Speaker endorses the reported belief, while others are nonfactive, silent on what Speaker thinks about it. Enriching this picture, *yǐwéi* exemplifies a relatively less-studied class of strategies (recently discussed: Holton '17, Hsiao '17, Anvari *et al* '18) for reporting beliefs that Speaker views with skepticism.

Refs P. Anand and V. Hacquard (2014): Factivity, belief & discourse ● A. Brasoveanu (2009): Modified numerals as post-suppositions ● G. Chierchia & S. McConnell-Ginet (1990): Meaning & grammar ● N. Constant (2012): English Rise-Fall-Rise ● D. Farkas (2002): Varieties of indefinites ● R. Holton (2017): Facts, factives, and contrafactives ● S. Lauer (2009): Free relatives with *-ever* ● S. Lauer (2012): Some news about *irgend-* & *algún* ● M. Simons (2007): Observations on embedding verbs ● R. Stalnaker (1979): Assertion