Biased Questions Project Answers to Resolve Issues

Cory Bill (ZAS) and Todor Koev (BGU)

Abstract

Previous literature has distinguished between two types of bias in polar questions: "original bias", which reflects the speaker's prior belief regarding the truth of the answers, and "contextual bias", which conveys freshly introduced evidence for or against a given answer. Building on these two fundamental biases and employing a novel diagnostic based on bare polar response particles, we argue for the existence of a third type of bias, which we call "projected bias", indicating the answer that the speaker considers most likely to enter the Common Ground. We posit this flavor of bias on account of its distinct alignment pattern, which cannot be attributed to merely mirroring one of the two fundamental biases or a random choice between them. Specifically, while projected bias requires the presence of a fundamental bias in order to arise, it may align with either of them and displays a strong affinity with contextual bias when the two fundamental biases are in conflict. We capture this alignment pattern by appealing to differences in evidential strength between original and contextual bias, thereby highlighting the manner in which projected bias coordinates on the discourse goal shared by the speaker and the addressee of resolving the issue raised by the question. We make our proposal explicit by implementing it within the Scoreboard Model.

Keywords: question bias, projected answers, evidence strength, resolving issues

1 The plot

According to the popular Question Under Discussion framework (Roberts 2012), discourse revolves around raising and resolving issues. This dynamics is manifested by sequences of questions and answers, where the questions are typically left implicit as they cohere with and can be derived from the focus structure of the answers. The Question Under Discussion framework is often implemented within the Scoreboard Model (Farkas and Bruce 2010), which represents a record of the conversation at a given stage and includes various components: the Table containing the issue currently under discussion, the current and projected commitments of discourse participants, the current and projected Common Grounds, and so on. The Scoreboard Model has also been employed to capture the effect of **biased questions**, or information-seeking questions that present one of the answers as privileged in some sense (e.g., Malamud and Stephenson 2015). Building on this prior work, this paper studies how different *flavors* of question bias interact and how this interaction plays into the discourse dynamics made explicit by the Scoreboard Model.

We begin by introducing the question types that form the empirical basis of our investigation. In the absence of relevant contextual information, canonical polar questions like (1) aim to elicit information from the addressee in a neutral way, maintaining an equal balance between the two polar answers corresponding to *yes* and *no*. This is as predicted by standard accounts of question meaning, which partition the logical space into a set of equally accessible alternatives (Hamblin 1973; Groenendijk and Stokhof 1984; Ciardelli et al. 2019; a.o.). At the other end of the spectrum are polar questions which effectively lack an interrogative force and serve as rhetorical devices (Heim 1984; Krifka 1995; Han 2002; van Rooy 2003; Rohde 2006; Caponigro and Sprouse 2007; Biezma and Rawlins 2017; Farkas 2026). For example, this rhetorical effect is triggered by questions with minimizers like *lift a finger* or the discourse marker *after all*, as in (2).

(1) Is it raining outside? (positive polar question)

(2) a. Did he lift a finger to help when you were in trouble? (polar question with a minimizer)
b. After all, did he help you when you were in trouble? (polar question with *after all*)

Between these two extremes, there are numerous non-canonical polar questions that, while seeking to elicit information from the addressee, also convey some additional attitude or information related to the resolution of this question. This extra, non-interrogative aspect of the meaning is known as **question bias**. Interrogative forms which systematically have this effect include at least the following: polar questions with high (preposed) negation,[1] accented low negation, a verum accent, or conversational *really*, as in (3); various tag questions, including nuclear and postnuclear reverse-polarity tag questions, like (4); questions with a nonstandard pairing between clause type and prosodic contour, such as rising declaratives[2] and falling interrogatives, see (5); and finally,

[1] For concreteness, we are only considering the "outer" negation reading of high negation questions, in which the speaker tries to confirm the positive alternative (Ladd 1981). The existence of an "inner" negation reading, in which the speaker tries to confirm the negative alternative, has been disputed (Sailor 2013; Goodhue 2022b), although there is also evidence in its support (Romero et al. 2017; Arnhold et al. 2021; Jeong 2021).

[2] Jeong (2018) distinguishes between two types of rising declaratives in English, which differ both prosodically and interpretationally. "Inquisitive" rising declaratives end in a low-rising terminal contour (L* H-H%) and present the addressee as more knowledgeable than the speaker, whereas "assertive" rising declaratives end in a high-rising terminal contour (H* H-H%) and present the speaker as more knowledgeable than the addressee. Only the former inquisitive type is relevant here, as only this type displays a proper interrogative force.

questions in which *even* associates with the low endpoint of some scale, as in (6).[3]

(3) a. Isn't Zelda from Australia? (high negation question)
b. Is Zelda NOT from Australia? (accented low negation question)
c. IS Zelda from Australia? (verum question)
d. Is Zelda really from Australia? (*really*-question)

(4) a. It's snowing, ISN'T it. (nuclear reverse-polarity tag question)
b. It's snowing isn't it? (postnuclear reverse-polarity tag question)

(5) a. You are the new manager? (rising declarative)
b. Are you a terrorist. (falling interrogative)

(6) Can Jack even solve the EASIEST problem? (*even*-question)

Notably, alternative questions, as in (7), are the mirror image of the questions in (3)–(6). These interrogative forms do not seem able to convey bias at all.

(7) Did ALFONSO or JOANNA give you a ride. (alternative question)

The empirical focus of this paper is on the biased questions in (3)–(6). As already mentioned, such questions do two things simultaneously. First, just like canonical questions, they raise an issue by partitioning the set of possibilities and letting the addressee select one of the partition cells. Second, they lend a privileged status to one of the answers, thus giving rise to the intuition of bias. The precise nature of this bias, as well as how it is derived, has been explored in previous literature (for representative accounts, see Abels 2003; van Rooy and Šafářová 2003; Guerzoni 2004; Romero and Han 2004; Reese 2007; Repp 2012; Krifka 2015; Malamud and Stephenson 2015; AnderBois 2019; Goodhue 2022b; a.o.).

Two varieties of question bias have been distinguished in such literature (Romero and Han 2004; Roelofsen et al. 2012; Sudo 2013; a.o.). The first is **original bias** (OB), which conveys the speaker's prior belief regarding what they take to be the true answer (Ladd 1981). For example, (3d) expresses the speaker's belief that Zelda is not from Australia, and hence conveys a positive OB. A second variety is **contextual bias** (CB), which reflects the shared evidence for or against a particular answer (Büring and Gunlogson 2000). Sticking with (3d), in addition to its obligatory OB effect just mentioned, the use of this question also requires contextual evidence supporting the conclusion that Zelda is from Australia. For example, this question would fit a context where someone has just claimed that Zelda is from Australia (positive contextual bias), while the speaker thought that Zelda grew up in New Zealand and therefore doubts the claim (negative original bias).

Importantly, OB and CB as described above do not directly engage with how the issue raised by the question itself is to be resolved. That is, none of these bias flavors directly indicates which answer is expected by the speaker to enter the Common Ground. On the one hand, this is unsurprising, as biased questions are information-seeking and the addressee can freely choose either polar answer depending on their epistemic state. At the same time, it would be puzzling if such biases—which reflect the speaker's beliefs and the shared contextual evidence—did not serve to facilitate in some way the addressee's response. Yet given the way they are currently understood in

[3]Notation in linguistic examples: Capital letters indicate prosodic prominence, a comma indicates a prosodic break, a question mark indicates a rising boundary tone, and a full stop indicates a falling boundary tone.

the literature, the notion of question bias (OB or CB) and that of expected answer could in principle be unrelated.

To bridge the gap between these two notions, we motivate and spell out a third—truly distinct yet not fully independent—variety of question bias, which we call **projected bias** (PB). Its role is already anticipated by the projection component of the Scoreboard Model. That is, among other things, this component signals which answer is flagged as being more likely to be accepted into the Common Ground. By employing a new empirical diagnostic based on bare polar response particles, we will demonstrate that PB arises as the cumulative effect of the two fundamental biases, OB and CB. Specifically, PB requires the presence of OB or CB in order to emerge. In the absence of a conflict between the two fundamental biases, PB aligns with either (or both). However, when the two are in conflict, PB aligns with CB and against OB.

We propose an explanation of these facts by appealing to a difference in the evidential basis of the two fundamental biases. Specifically, we appeal to the fact that CB is based on knowledge whereas OB is based on belief, hence the former comes with a stricter evidential requirement than the latter. This distinction is important because PB, as suggesting which answer is anticipated to be added to the Common Ground, requires some evidence and can in principle align with either of these two bias flavors. However, in the case of an epistemic conflict, PB aligns with knowledge-based CB over belief-based OB.

The paper is organized as follows. Section 2 outlines the different flavors of question bias and includes a brief discussion of what previous accounts predict about these flavors. Section 3 establishes the alignment pattern between different bias flavors and presents our own proposal. Section 4 implements the proposal in the Scoreboard Model. Finally, Section 5 concludes and draws broader implications regarding how projected bias allows the two fundamental biases to contribute to the general discourse goal of resolving the current conversational issue.

2 Varieties of question bias

This section provides general background on the Question Under Discussion framework assumed throughout the informal discussion, defines three distinct notions of question bias, derives empirical diagnostics based on these definitions, and briefly discusses the predictions of prior accounts regarding the existence and interaction between different bias flavors.

2.1 Background on questions and answers

Following the classical view in formal semantics, the meaning of a declarative sentence is tied to its truth conditions. The basic idea is that to understand the meaning of a sentence is to know exactly what state of affairs must hold in the world of evaluation for that sentence to be true (Tarski 1956; Kripke 1963). One challenge for this view is that it does not immediately cover interrogative sentences, which intuitively constitute requests for information and as such lack truth conditions. To overcome this limitation, Hamblin (1973) proposed that any viable semantics for questions should be able to specify what counts as an answer to a given question, thus in effect reducing the semantics of interrogatives to that of declaratives. Specifically, Hamblin proposed that the meaning of a question is the set of its possible answers (i.e., answers that would fully resolve the issue raised by the question). For a polar question like *Does Mira like ice cream?*, for instance,

this approach yields the set of polar alternatives {'Mira likes ice cream', 'Mira doesn't like ice cream'}, corresponding to its positive and negative answers, respectively.

Against this static perspective, Stalnaker (1978) emphasized the dynamic aspect of linguistic meaning, placing it within the broader context of human communication. Rather than viewing the meaning of a declarative sentence in terms of its truth conditions, Stalnaker focused instead on how an assertion of a declarative sentence affects the linguistic context. Specifically, he proposed that the essential effect of an assertion is to enrich the Common Ground with the descriptive content of the sentence by conveying a commitment to that content. For example, when a speaker utters *Mira likes ice cream*, they present themselves as committed to the truth of the proposition expressed and as intending to add it to the Common Ground. Importantly, since the Common Ground is a set of propositions representing the shared discourse commitments of the speaker and the addressee (i.e., what they mutually take for granted in the conversation), assertion includes a negotiation component. An act of assertion will successfully update the Common Ground only if the addressee does not object to the conveyed semantic content.

Hamblin (1973)'s view of the meaning of questions as sets of answers not only extended the traditional truth-based semantics of declaratives to questions, but also provided a mechanism for determining how questions and answers fit together.[4] However, his semantics does not specify how questions and answers interact in conversation. Stalnaker (1978)'s approach to linguistic meaning has the opposite limitation. While it recognizes the dynamic nature of assertion, it does not link an assertion to the discourse question it is intended to answer. Roberts (2012) combines these two strands of prior work into a comprehensive framework that captures both the dynamic nature of meaning and the intuition that assertions are guided by (typically implicit) topic-setting questions.

More concretely, Roberts (2012) views language as a game for achieving certain conversational goals (e.g., sharing information about the world). In pursuing these goals, speakers make two types of moves: moves that raise an issue (i.e., questions) and moves that resolve a standing issue (i.e., assertions or answers that address these questions). When a question is accepted as a valid move, it partitions the space of possibilities into cells corresponding to its answers and becomes an immediate topic of discussion, called the **Question Under Discussion**. An assertion following this move is felicitous only if relevant to that question (i.e., if it answers the question). When the answer is accepted by all discourse participants, its content is added to the Common Ground and the issue is settled. Crucially, unlike previous approaches, the Robertsian framework gives a handle on how questions guide the conversation. Specifically, Questions Under Discussion are partially ordered by informativeness, giving rise to a tree structure in which more general questions are split into subquestions, which are then answered in a particular order. To illustrate, in (8) the declarative utterance is mapped onto the question structure as shown and fully resolves the topmost question.

(8) *Who*$_{\{Hilary,\ Robin\}}$ *ate what*$_{\{bagels,\ tofu\}}$*?* (modeled after Roberts 2012: p.16)

- a. *What did Hilary eat?*
  - i. *Did Hilary eat bagels? (Yes.)*
  - ii. *Did Hilary eat tofu? (No.)*
- b. *What did Robin eat?*
  - i. *Did Robin eat bagels? (Yes.)*
  - ii. *Did Robin eat tofu? (Yes.)*

[4] This is sometimes called "question–answer congruence" and is mediated by focus structure (e.g., Rooth 1992).

Hilary ate bagels but not tofu, while Robin ate both bagels and tofu.

In our informal discussion below (Sections 2 and 3), we will assume the Robertsian framework just outlined, treating biased questions as a special kind of Question Under Discussion, namely issue-setting questions that additionally convey bias. In Section 4, we will provide a formal analysis of the various alignments patterns of different bias flavors within the Scoreboard Model of Farkas and Bruce (2010).

2.2 Three bias flavors

As already mentioned in the introductory section, two well-established varieties of question bias are original bias (OB) and contextual bias (CB). Although the terminology varies somewhat across authors,[5] the distinction between the two is implicit in Ladd (1981), Romero and Han (2004), and Reese (2007), and is explicitly drawn in Roelofsen et al. (2012) and Sudo (2013). We now introduce each of these fundamental biases in turn.

Ladd (1981) characterized OB as the speaker believing or assuming that one of the possible answers is true (see p.164). Since OB is anchored to the internal state of the speaker rather than the context of utterance, this bias flavor must be explicitly signaled by the question form itself. Based on this characterization, we propose the more rigorous definition in (9).

(9) ORIGINAL BIAS

A question Q uttered in a context c conveys **original bias** for a partition cell p iff Q indicates that the prior doxastic state of the speaker in c supports p over other partition cells.

There is one important consequence of this definition that sheds light on the notion of OB and how it differs from the notion of CB (discussed below). Notice that a question carrying original bias conveys information about the speaker's doxastic state rather than about some external feature of the context. As such, uttering a question whose original bias implication is not supported by the speaker's doxastic state would generally not be infelicitous or uninterpretable. However, doing so would be *insincere*. That is, in the absence of the relevant belief, the speaker would be disingenuous in presenting information that they do not actually believe. This is a very different situation from the one of CB, where the question needs to obey certain contextual restrictions and failing to do so results in infelicity (see (12) and (13)).

The definition of OB in (9)—as well as those of CB and PB stated in (11) and (17) below—adopt the propositional approach to the semantics of questions, spearheaded in Hamblin (1973). According to this approach, a possible answer corresponds to one of the cells of the partition generated by the question's semantics. The possibility of a choice between different partition cells implies that the question is information-seeking rather than rhetorical. Relatedly and to maintain generality, the definitions allow for more than two possible answers (partition cells), given that wh-questions may also exhibit bias (e.g., Guerzoni 2003, ch.3). While our discussion is limited to polar questions, presumably the same distinctions between bias flavors apply to wh-questions as well.

Interrogative forms that have been claimed to systematically trigger OB include high negation questions, accented low negation questions, questions with epistemic *really*, questions with a

[5] More specifically, OB is also called "speaker" or "epistemic" bias, while CB is sometimes called "evidential" bias.

verum accent, and various tag questions (Romero and Han 2004; Gutzmann et al. 2020; Goodhue 2022a, 2022b; Scheffler and Malamud 2023; Bill and Koev 2025). For example, high negation questions obligatorily convey a positive OB, as illustrated in (10).

(10) Isn't Maureen coming to the party?
$\rightsquigarrow$ *The speaker thought that Maureen was coming to the party.*

A second type of bias is CB, first discussed in Büring and Gunlogson (2000). This bias flavor reflects the evidence available in the current speech situation for or against a possible answer. It is defined in (11).

(11) CONTEXTUAL BIAS
A question Q uttered in a context c conveys **contextual bias** for a partition cell p iff some piece of evidence shared in c provides sufficient support for p, and no other similar piece of evidence supports another partition cell of Q.

Before commenting on the above definition, we provide examples that demonstrate the effect of CB. This bias flavor is particularly evident in rising declaratives, which have been argued to require contextual support without also requiring a prior belief on the part of speaker (Gunlogson 2003; Malamud and Stephenson 2015; Jeong 2018; Rudin 2022). In (12), for instance, the rising declarative would be infelicitous in the absence of contextual evidence for rain (e.g., if the speaker's colleague had entered the office wearing sunglasses and a T-shirt).

(12) *Context: The speaker sees her coworker walk into their shared windowless office wearing a wet raincoat. The speaker has no prior information on the weather conditions outside.*
It's raining outside?

Even-questions display a similar pattern (cf. Abels 2003; Guerzoni 2004). In (13), Prof. Smith need not have approached Jack's exam sheet with a specific belief about pass or fail. Rather, the *even*-question is licensed by the (negative) evidence that has just become manifest in the context.

(13) *Context: Prof. Smith and her tutor have been grading the final exams of their students.*
Tutor: Let's see. Jack didn't solve the hardest problem. He didn't solve the intermediate problem either.
Prof. Smith: Did he even solve the EASIEST problem?

We now return to the definition of CB in (11). First, note that the contextual evidence in favor of a given answer must convey a *sufficient* degree of support.[6] That is, very weak evidence will not trigger CB. For example, a few drops on the coworker's jacket would not license the utterance in (12). Similarly, Jack's failure to solve the hardest problem, together with pending information about his performance on the intermediate one, would not suffice to license the *even*-question in (13). The role of evidence strength for shaping bias will become particularly relevant when we introduce our own proposal in Section 3.3.

Our second point about the definition of CB concerns the exclusion of contravening evidence. One case in point are alternative questions, which require that both alternatives have been raised in

[6]Compare Büring and Gunlogson (2000)'s notion of "compelling" evidence for an answer p, defined as evidence which, when considered in isolation, would allow the participants to assume the truth of p.

the preceding discourse (Bolinger 1978; Biezma and Rawlins 2012). Since this requirement implies comparable evidential support for each alternative, such questions still offer an unbiased choice and are, unsurprisingly, neutral in terms of CB. This is where the caveat of excluding a second piece of evidence supporting an alternative answer comes in. Without this caveat, the definition in (11) would incorrectly classify alternative questions as contextually biased for both alternatives simultaneously.

Finally, CB is often thought to follow a general pattern across all polar questions. Specifically, any evidence present in the context must support, or at least not contradict, the question's prejacent.[7] This makes the interesting prediction that positive polar questions, although canonical in form, may also count as contextually biased under the definition in (11). Indeed, based on examples like (14), Büring and Gunlogson (2000) establish that positive polar questions are compatible with contextual evidence for the scope proposition but incompatible with contextual evidence against said proposition.

(14) *Context: A and S have conducted a psycholingnistic experiment in which the subjects have all certified that they are right-handed. They encounter Carl, who they recognize as one of their subjects, cutting bread with his left hand.*

S: Is Carl left-handed?

S′: #Is Carl right-handed?

Notice that the same cannot be said about OB. For example, positive polar questions will not count as carrying OB in the sense of the definition in (9), since their canonical form does not require the speaker to be biased for one of the answers. This fact emphasizes an important difference between OB and CB: while the former is tightly linked to the semantics of the question, the latter is primarily a property of contexts. More generally, this suggests that OB and CB are independent of each other. That is, it is possible for a question utterance to convey just OB, just CB, or both. Examples of the second case were already cited in (12) and (13). The first case is illustrated in (15), which involves high negation questions in "suggestion" scenarios, showing that such questions obligatorily convey (positive) OB but need not involve CB as well. The third case is illustrated in (16), exploiting the fact that *really*-questions obligatorily convey (negative) OB and (positive) CB (Romero and Han 2004).[8]

(15) *Context: Kathleen and Jeff have just come from Chicago on the Greyhound bus to visit Bob in Ithaca.* (Ladd 1981: p.164)

Bob: You guys must be starving. You want to go get something to eat?

Kathleen: Yeah, isn't there a vegetarian restaurant around here—Moosewood, or something like that?

(16) *Context: The speaker thought that Mary is a teetotaler but has just observed her order a beer.*

Does Mary really drink?

[7] Analyses that try to derive this generalization include van Rooy and Šafářová (2003) and Trinh (2014) .

[8] The alignment possibilities among different bias flavors are discussed in detail in Section 3.

Having characterized the two fundamental biases, namely OB and CB, we now introduce a third bias flavor: projected bias (PB). The role of this bias flavor is to coordinate on resolving the issue raised by the question by pointing to the expected answer. While PB has not been explicitly fleshed out in previous work, it is already anticipated by accounts that appeal to future discourse developments in some shape or form (e.g., Malamud and Stephenson 2015; see Section 2.4 for further discussion). We define PB in (17).

(17) PROJECTED BIAS

A question Q uttered in a context c conveys **projected bias** for a partition cell p iff the speaker takes p to be the most likely resolution of the issue raised by Q in c.

One question type that strongly suggests the existence of PB is that of rising declaratives. As already mentioned in the paragraph immediately preceding example (12), such questions require some form of contextual support for the prejacent proposition. This support has been analyzed in terms of the commitments or beliefs of the addressee (Jeong 2018; Rudin 2022).[9] Importantly, this analysis implies that the speaker expects the issue raised by the question to be resolved in favor of the positive answer. For instance, in (18) Bertha takes it that Alvin believes that the sweater he chose is made of 100% cotton, and expects to receive a positive answer. This is so even though Bertha's prior belief leads her to doubt Alvin's judgment.

(18) *Context: Alvin and Bertha are at a department store. Bertha has asked Alvin to go pick out a 100% cotton sweater; Alvin does not know much about clothes, and is not particularly detail-oriented. He returns with a sweater.* (Rudin 2022)

Bertha: That's 100% cotton?

In Section 2.3, we propose an empirical diagnostic that targets this intuition about the expected answer, one that is based on bare polar responses.

Adding to these linguistic intuitions about the expected answer, there are also theoretical considerations that motivate the existence of PB. More specifically, the Scoreboard Model already incorporates projected states of the Common Ground. These states include the available answers for updating the current Common Ground, thereby resolving the issue placed on the Table. From here, it is only a small step to suppose that certain projected states are presented as more likely to materialize than other such states. We take this to be what questions carrying PB indeed achieve. That is, such questions raise an issue by delineating a partition while simultaneously suggesting a resolution by promoting one of the possible answers in the form of projecting a Common Ground enriched with that answer (see Section 5 for formal details).

The theoretical case for the existence of PB can be generalized in terms of the timing of bias and the corresponding evidence, as predicted by the definitions of the three bias flavors above. That is, CB represents relevant evidence present in the *current* linguistic or nonlinguistic context. In contrast, OB reflects the speaker's *prior* beliefs, formed before the speaker uttered the question or accepted any relevant contextual evidence. It therefore stands to reason that, on the basis of these two, there is a bias flavor that makes *predictions* about the answer on which the discourse participants are expected to converge. We claim that this is precisely the role of PB, namely to

[9] Specifically, Jeong (p. 342) proposes that rising declaratives have the effect of adding the positive answer (the prejacent proposition) to the addressee's projected commitment set, while Rudin (p. 351) argues that rising declaratives are felicitous only if the speaker has reason to suspect that the addressee believes the positive answer.

coordinate on the issue raised by the question while taking into consideration the two fundamental biases. The future-oriented nature of PB is helpful in explaining its role in aggregating the evidence already shared in the context (see Section 3).

2.3 Diagnosing the different bias flavors

Section 2.2 characterized and motivated the existence of three distinct notions of question bias: OB, CB, and PB. In this subsection, we derive from these definitions a set of diagnostic tests designed to tease apart and sharpen our linguistic intuitions about these bias flavors.

A standard way of diagnosing OB or CB is to specify a prior context that reports on the doxastic state of the speaker or provides a piece of relevant evidence, respectively. A convenient way to demonstrate both biases within a single example involves *really*-questions, which simultaneously convey a negative OB and a positive CB (Romero and Han 2004). One such case was already presented in (16), where both biases are established in the metalinguistic context. Another example is given in (19). Here, OB is determined by the metalinguistic context, while CB is implied by the utterance immediately preceding the target question.

(19) *Context: A and B want to go out for dinner. As far as A knows, there used to be only one vegetarian restaurant in town called Veggie Grill, but B heard that it closed down.*

A: Since you are vegetarian, let's go to a vegetarian restaurant.

B: Is there really a vegetarian restaurant in this town?

Apart from being stated in the metalinguistic context, OB can also be diagnosed by having the speaker verbalize in a follow-up utterance either their prior belief or the evidence on which that belief is based. This is illustrated in (20) and (21), respectively, where the positive OB is reinforced.

(20) Is Zelda really coming to the party? I thought she $\left\{ \begin{array}{c} \text{wasn't} \\ \text{\#was} \end{array} \right\}$.

(21) Is Zelda really coming to the party? She said she $\left\{ \begin{array}{c} \text{wouldn't} \\ \text{\#would} \end{array} \right\}$.

How can PB be probed empirically? Since this type of bias has not been fully developed conceptually, we know of no proposed empirical diagnostic that targets it. However, given the nature of the PB as pointing to a projected answer (recall definition (17)), we expect that this answer is privileged in some sense and that this has detectable empirical effects. Importantly, since biased polar questions carry proper interrogative force, both answers are available and the particular choice will depend on the epistemic state of the addressee. Nevertheless, in the presence of PB, we expect one of the answers to be more readily accessible.

With this in mind, we propose to diagnose PB by determining which of the two polar response particles, *Yes* or *No*, sounds more natural when used to answer the question without elaboration (cf. Pope 1972, ch.5; Holmberg 2013; Goodhue and Wagner 2018). We call this the "bare response particle" test. To illustrate, according to this test high negation questions (in conflicting scenarios) convey a negative PB, see (22).

(22) *Context: Paul knew that Sarah had planned to visit Europe, but just saw her showing some friends pictures of the Great Wall of China.*

Paul: Didn't you go to Europe?

Sarah: No. ≻ Yes.

Several comments are in order here. The first concerns the nature of the test as conveying a *preference* (marked as ≻), not infelicity. Importantly, we are not claiming that one of the response particles should sound completely natural while the other is ruled out. Rather, our claim is simply that—all other things being equal—there should be a contrast in flow between the two responses.

A second point is that the contrast is expected to arise with *bare* polar response particles, that is, without further elaboration. If additional information is introduced next to the response particle, the observed preference may disappear. In (22), for instance, *Yes, but I also visited China* would be perfectly fine, yet it would also convey a sense of violated expectation. Similarly, in (22) the positive answer may become fully natural (or even more natural than the negative answer) if flagged by the discourse marker *actually*, which seems to select a less expected answer, as in *Actually, yes I did*. These additional data points demonstrate once again that, while a positive answer is available, it is not projected by default.

Notice also that the two polar particles must unequivocally select a particular answer. In (22), for example, *Yes* selects 'I went to Europe' while *No* selects 'I didn't go to Europe'. When bare response particles are ambiguous between the positive and the negative answer (e.g., as when following a low negation polar question like *Did you not go to Europe?*), the above diagnostic cannot be applied without some independent means of disambiguation, such as prosodic cues.[10]

Finally, we will not elaborate on the theoretical mechanisms underlying the interpretation and use of bare polar response particles. To be sure, there are several proposals on the market, treating such particles as clauses with an elided prejacent that copies the question prejacent (Laka 1990; Kramer and Rawlins 2011; Holmberg 2013), as propositional anaphors that obey certain salience restrictions (Krifka 2013; Claus et al. 2017), or as morphological realizations of various feature bundles (Pope 1972; Roelofsen and Farkas 2015; Farkas and Roelofsen 2019). What matters for our present purposes would follow under any of these approaches as long as both bare polar particles are available, they clearly select opposing answers, and yet one of the answers is more readily accessible than the other. Foreshadowing, in Section 4 these properties will fall out from the shape of the issue on the Table and the projected state of the Common Ground.

2.4 Existing approaches to question bias

In this subsection, we briefly review three major approaches to question bias, highlighting the bias flavor they intend to capture—or, in fact, do capture. We demonstrate that none of these approaches offers a comprehensive picture on the three bias flavors identified in the previous subsections.

Epistemic operator approach One approach to question bias is based on the idea that certain question partitions express an epistemic attitude toward a possible answer (Romero and Han 2004; Repp 2012; Frana and Rawlins 2019; Silk 2020; Goodhue 2022b). Romero and Han (2004), who pioneered this approach, apply this idea to polar questions with a verum accent, high or stressed

[10]For example, Goodhue and Wagner (2018) demonstrate experimentally that, when following low negation questions or negative rising declaratives, bare response particles with a fall–rise "contradiction contour" (transcribed as %H L* L-H%) are significantly more likely to be interpreted as tagging the positive answer than bare response particles with a regular declarative fall (i.e., H* L-L%).

low negation, or epistemic *really* (examples of each of these forms are cited in (3)). What all these question forms are claimed to have in common is the presence of a metaconversational operator VERUM (cf. Höhle 1992), which encodes certainty that the question prejacent should be added to the Common Ground.

More specifically and reformulating somewhat, questions with VERUM convey bias by virtue of the two pragmatic principles in (23) and (24).

(23) PRINCIPLE OF ECONOMY

Metaconversational operators, like VERUM, are licensed only when there is an epistemic conflict about the prejacent between the speaker's belief and the contextual evidence.

(24) DOUBT THE PREJACENT

The speaker formulates their question so that they require evidence for the question prejacent, thereby conveying doubt toward that prejacent.[11]

For example, a question with a verum accent like *IS Jane coming?* conveys the speaker's doubt about the positive prejacent, hence the speaker's prior belief is that Jane is not coming. Because of the conflict requirement, the context must then support the opposite conclusion, namely that Jane is coming.[12] Reversely, a high negation question like *Isn't Jane coming* conveys doubt about the negative prejacent, amounting to the belief that Jane *is* coming. Accordingly, the contextual evidence must support the negation of that prejacent, amounting to the proposition that Jane is not coming.[13] More generally, the prediction across all VERUM-based question forms is that the contextual evidence supports the prejacent while the speaker's belief supports its negation. This prediction is borne out.

What kind of predictions does the epistemic operator approach make about the different bias flavors? As just discussed, it draws a distinction between OB and CB, deriving their opposed alignment in contexts where the question is licensed. Moreover, VERUM-based questions delineate a partition that problematizes the metalinguistic issue of which of the two ground-level propositions (p or $\neg p$) should be added to the Common Ground. This comes close to our notion of PB, as defined in (17). However, this approach makes no clear predictions about the direction of PB, nor about how PB is related to OB and CB, if at all.

Single answer approach Another approach derives the negative bias in questions with *even* by reducing the set of possible answers to a singleton as a result of a clash with its lexical presuppositions (Ladusaw 1980, ch.8; Heim 1984; Wilkinson 1996; Abels 2003; Guerzoni 2004; Chierchia 2013, ch.3). While this approach was primarily intended to capture the rhetorical reading of questions with minimizers (such as *the slightest bit* or *lift a finger*), it has also been applied to truly information-seeking questions in which *even* associates with a minimal element like *the easiest problem*. Here we focus on this latter type of questions.

[11] This second principle is less highlighted in Romero and Han (2004)'s original work. However, it plays a crucial role in their derivation of bias (see below), so we state it explicitly here.

[12] In reality, verum questions only optionally carry OB (Gutzmann et al. 2020; Goodhue 2022a). Here we set this complication aside.

[13] Here we gloss over the fact that high negation questions are also felicitous in neutral contexts (Ladd 1981), a case Romero and Han take care of by extending their ECONOMY principle with a second clause requiring compliance with Quality.

Following more closely Abels (2003), here is how to derive the negative bias of *Did Jack even solve the EASIEST problem?*. It has been independently argued that *even* triggers the following three presuppositions in (25) (cf. Karttunen and Peters 1979; Rooth 1985; Wilkinson 1996; a.o.).

(25) a. NONTRIVIALITY
The set of focal alternatives contains the prejacent and at least one other alternative.

b. SCALARITY
The prejacent of *even* is the least likely alternative.

c. TRUTH OF ALTERNATIVES
All alternatives to the prejacent are true.

Nontriviality dictates that the set of focal alternatives contains *Jack solved the easiest problem* and at least one additional element. This additional element cannot be a positive alternative like *Jack solved the intermediate problem*, since any such alternative would be less likely than the prejacent, thereby contradicting Scalarity. Consequently, the only viable option is the negative alternative *Jack solved none of the problems*, representing the negation of the disjunction of all positive alternatives. Finally, by Truth of Alternatives, this negative alternative must be true. Hence the bias for the negative answer *Jack didn't solve the easiest problem*, which is entailed by the only true focus alternative *Jack solved none of the problems*.[14]

What bias flavor is modeled by reducing the set of answers to a single negative alternative? In most of its incarnations, this approach has been intended to capture OB. The reasoning typically goes as follows: since the question carries a single possible answer, the speaker must believe this to be the true answer (e.g., Ladusaw 1980, ch.8). However, given that the nature of the explanation steers the discourse toward a single answer, what is at least partially modeled here is the notion of PB. That being said, this approach is too course-grained to distinguish between suggested and accepted answers. That is, the intuition about questions in which *even* associates with a minimal element is that they are biased yet information-seeking. Nevertheless, if taken at face value, the single answer approach implies that such questions entirely lack interrogative force.

Conventional implication approach A third approach hard-codes the bias implication into the meaning of various noncanonical questions in the form of an indirect assertion (Reese 2007), a projected commitment of a discourse participant (Krifka 2015; Malamud and Stephenson 2015); an evidence-based belief (Farkas and Roelofsen 2017), a secondary discourse topic (AnderBois 2019), or a polarity focus antecedent (Bill and Koev 2022). Here we illustrate this approach using the account of Malamud and Stephenson (2015), which adopts the Scoreboard Model and affords an easy comparison with our own approach presented in Section 3.3.

Malamud and Stephenson discuss several biased question forms, of which we mention only (postnuclear) reverse-polarity tag questions, comparing their Scoreboard effects to those of root declaratives and (neutral uses of) positive polar questions. Starting with root declaratives, these forms provide a baseline for the effects of assertion. These effects include placing the expressed

[14] According to the single answer approach, only those questions are biased where *even* associates with the *low* endpoint of a scale (Guerzoni 2004). When this is not the case, no bias need to arise (cf. *Jack solved the easy and the intermediate problems. Did he even solve the HARDEST problem?*). Moreover, questions with minimizers are given a similar treatment, assuming that such items are licensed by a covert EVEN operator (Heim 1984).

proposition on the Table as well as adding it to the speaker's discourse commitments and the projected Common Ground.[15] Positive polar questions present the opposite case: they register no commitments but place the expressed partition on the Table, creating two projected Common Grounds, one for each polar answer. Finally, reverse-polarity tag questions fall somewhere in between root declaratives and positive polar questions. Like the former, they put on the Table a single proposition and project it to the Common Ground. However, like the latter they do not incur hard commitments; they only project *future* commitments.

With regard to bias flavors, Malamud and Stephenson conceive of their account as primarily capturing OB. This is achieved via the notion of projected or tentative commitment, which they take to imply that the speaker is in possession of some "imperfect" evidence (Malamud and Stephenson 2015, p.26). At the same time, this account comes closest to modeling our notion of PB. Unlike the epistemic operator approach, it clearly delineates the idea of an expected answer; and unlike the single answer approach, it does so while keeping the interrogative force of the biased question intact. However, the account does not draw clear distinctions between different bias flavors. Consequently, no systematic link is established between PB and the two fundamental bias types, OB and CB.

Intermediate summary We have briefly discussed three major approaches to question bias, covering a wide spectrum of biased question forms. We saw that these approaches engage with the different bias flavors in various ways. Importantly, none of the approaches captures all three bias flavors, nor does any of them make predictions about the alignment possibilities of PB against OB and CB. We will therefore not adopt a specific account of question bias. Our goal here is more modest. It is to predict the existence and direction of PB, assuming that the two fundamental biases (OB and CB) have already been derived in some manner.

3 Alignment patterns

This section constitutes the meat of the paper. We first establish the alignment patterns between different bias flavors and then go on to propose how PB can be computed from OB and/or CB.

3.1 The status of projected bias

When it comes to the theoretical status of PB, there are two analytical possibilities. It could be that PB is not a real phenomenon but rather an epiphenomenon that shadows or "freerides" on one of the other fundamental biases, namely OB and CB. If this was the case, then there would not be much more to say about it other than identifying it as a secondary bias effect on the unfolding discourse. If, on the other hand, PB is a bias flavor that is truly distinct from OB and CB, then we would expect to find cases where PB aligns with OB and cases where PB aligns with CB, as well as some systematic way of choosing between these two alignments when both OB and CB are present and contradict each other. Such a pattern would warrant an explanation regarding the forces driving the alignment in each case, specifically one that links the two fundamental biases

[15] If the assertion is accepted by the addressee, the proposition will be removed from the Table and added to the Common Ground.

to the projected stages of the discourse. Below, we present evidence supporting the latter of these two possibilities, arguing for the idea that PB is a bias flavor of its own right. We then proceed to propose a basic explanation for the behavior we observe based on the notion of evidence strength.

3.2 Empirical generalizations

We now explore how PB interacts with OB and CB. Specifically, we are interested in whether PB always aligns with one of these fundamental biases, or whether we might find some variation in this regard, which would then require some further explanation.

Before we present the data, a caveat. In what follows, we do not use the same question types across all contexts to establish the three empirical generalizations we uncover. We do so for two main reasons. First, we want to ensure that the empirical claims we make are broad enough to encompass the full range of bias data, not just a subset thereof. This is as it should be, since our claims concern the alignment patterns of different biases rather than the specific properties of individual question types. The second reason is that there are certain natural limitations. That is, it would not have been possible to use the same question type throughout, since different questions impose specific restrictions on OB and CB, such that a given question type cannot appear across all contexts.[16]

To start with, we consider a context where neither OB nor CB is present. An example is shown in (26), which features a positive polar question that conveys no prior belief and is presented in a neutral context. In this case, PB is absent as well: the two responses using bare polar particles (without any further elaboration) are equally felicitous, a fact we mark with $\sim$.

(26) *Context: A has been sitting in a windowless room, having no clue about the weather conditions outside. They are wondering whether to eat their lunch inside or in the park. B arrives, having been outside.*

A: Is it raining outside?

B: Yes. $\sim$ No.

Another case where PB is absent is that of alternative questions, as in (27). As discussed in Section 2.2, such questions are used in balanced contexts where there are two pieces of contextual evidence, each supporting one of the question alternatives. Here, once again, the two bare response particles are equally felicitous.

(27) A: You said you wanted to play golf, but now you say you don't really feel like doing it. Do you want to play golf or not?

B: Yes. $\sim$ No.

As far was we can tell, it is not possible to construct a similar "balanced bias" case when it comes to OB. This is likely because the speaker cannot have believed both p and $\neg p$, as this would result in an inconsistent epistemic state. In (27), for instance, the speaker cannot be said to believe both

[16]For instance, as mentioned in Section 1, polar questions with epistemic *really* (e.g., (3d)) obligatorily convey negative belief and also require positive contextual evidence. As such, they cannot feature in the environments needed to establish the generalizations below whenever one or both of the fundamental biases have the wrong polarity or are entirely absent.

that the addressee wants and does not want to play golf. Rather, the speaker holds neither of these two beliefs.

These two data points lead us to the generalization in (28).

(28) ***Generalization 1 (No fundamental biases)***

No PB is present in the absence of the fundamental biases, and similarly when the two possible answers are equally supported by a fundamental bias.

We now turn to cases where only one of the fundamental biases is present. The outcome in such cases is fairly straightforward, with PB aligning with the relevant fundamental bias. This is clearly displayed for OB by the examples in (29)–(30), where the polarity of the OB (i.e., positive) aligns with the polarity of the more felicitous response particle.

(29) *Context: Kathleen and Jeff have just come from Chicago on the Greyhound bus to visit Bob in Ithaca.*

Bob: You guys must be starving. You want to go get something to eat?

Kathleen: Yeah, isn't there a vegetarian restaurant around here—Moosewood, or something like that?

Bob: Yes. $\succ$ No. (based on Ladd 1981, p.164)

(30) *Context: A has been sitting in a windowless room believing that it is sunny outside. Then B arrives.*

A: It's sunny outside isn't it?

B: Yes. $\succ$ No.

Similarly, in cases where only CB is present, PB aligns with it. This is illustrated with a positive polar question in (31) and a rising declarative in (32).

(31) *Context: A has been sitting in a windowless room, having no clue about the weather conditions outside. B arrives carrying a wet umbrella.*

A: Is it raining outside?

B: Yes. $\succ$ No.

(32) *Context: The speaker and her addressee made plans two days ago to get drinks tonight. They have not spoken about it since.*[17]

A: We're still on for tonight?

B: Yes. $\succ$ No. (based on Rudin 2022, p.346)

These observations lead us to posit our second generalization, stated in (33).

(33) ***Generalization 2 (Single fundamental bias)***

PB consistently aligns with the single fundamental bias present in the context.

[17]Rudin actually claims that, in the presented context, the rising declarative also conveys positive OB. However, even if A still believes that A and B will be going out for a drink, this belief is shared contextual information and would therefore count as CB under the definition in (11).

Finally, we consider cases where both OB and CB are present. When OB and CB are aligned, then PB follows suit. This is hard to demonstrate with most noncanonical questions, which typically require that OB and CB be conflicted. However, positive polar questions—while not inherently able to express OB (recall our comments in Section 2.2)—nevertheless suggest that this generalization is on the right track. This can be seen in (34).

(34) *Context: David knew that Raya wanted to visit Germany. He just saw her wear a dirndl.*

David: Did you go to Germany?

Raya: Yes. ≻ No.

More interesting are cases where both OB and CB are present but contradict each other. Crucially, in such cases PB aligns with CB and against OB. This pattern is demonstrated in both directions of misalignment for high negation questions in (35) and for questions with *really* in (36), which have flipped bias profiles.

(35) *Context: Paul was under the impression that Sarah had visited Europe but just saw her wearing a Chinese qipao.*

Paul: Didn't you go to Europe?

Sarah: No. ≻ Yes.

(36) *Context: Paul thought that Sarah would be visiting France over the summer break. However, he just saw someone looking at pictures of her on the Great Wall of China.*

Paul: Did Sarah really go to China?

Addressee: Yes. ≻ No.

These observations derive our last and most critical generalization, as stated in (37).

(37) ***Generalization 3 (Two fundamental biases)***

When both fundamental biases are aligned, PB follows their direction. However, in cases of epistemic conflict between OB and CB, PB aligns with CB.

Table 1 summarizes the different observations presented in this subsection. That is, it identifies the alignment of PB, given various alignments for OB and CB.

		CB		
		positive	*none*	*negative*
OB	*positive*	positive	positive	negative
	none	positive	none	
	negative	positive		

Table 1: Alignment pattern of PB relative to OB and CB.

3.3 Why projected bias is necessary

Before putting forward our proposal, let us point out that the alignment pattern revealed in the previous subsection argues for PB constituting a bias flavor in its own right. That is, if the bare response diagnostic indeed probes answer projection as claimed, then its results are best explained by positing a third flavor of bias. This is because there are really only three ways of reducing PB to the two fundamental biases, and all can safely be ruled out: (i) the preferred answer strictly aligns with OB, (ii) the preferred answer strictly aligns with CB, or (iii) the preferred answer may freely align with either OB or CB. Option (i) is falsified by cases where PB arises despite the absence of OB and also by cases where PB directly contradicts OB, while option (ii) is falsified by cases where PB arises despite the absence of CB. Option (iii) seems initially more plausible, but it fails to capture the robust preference of PB for CB over OB whenever the latter two are present but are in conflict. Since no other obvious reduction is available, the most natural conclusion is that PB is an additional *composite* flavor of bias based on the specific configurations of OB and CB. In short, the pattern revealed by the bare response test is simply too idiosyncratic to be explained away as keyed to one specific fundamental bias or as reflecting a random choice between the two fundamental biases.

To drive the point home, consider the following challenge, posed by an anonymous reviewer. Notice that our three generalizations about bias alignment established in the previous subsection were stated in declarative terms and employed the notion of PB. One might therefore wonder whether they can be reformulated in more procedural terms without making reference to PB at all. (38) presents one such attempt.

(38) a. In some cases no bias is present (neither OB nor CB), and so no bias is active (cf. Generalization 1).

b. In cases where only OB or only CB is present, the bias that is present is active (cf. Generalization 2).

c. In cases where both OB and CB are present, CB is active while OB is not (cf. Generalization 3).

We argue that this attempt fails to meet the challenge of reducing PB to the two fundamental biases. Specifically, we point to the wording in (38) invoking the notion of fundamental biases being "active", which begs the question, active with respect to what or in what sense? The intended interpretation seems to be that an active bias is one that has an effect on the projected development of the discourse. Thus, we are back to some version of our notion of PB.

3.4 Explaining the pattern

In the previous subsection, we identified the alignment patterns of the different bias flavors and codified our findings as the three generalizations stated in (28), (33), and (37). Taken together, these generalizations show that PB displays a strong affinity with CB: it consistently aligns with CB, even when OB is also present and misaligned. Nonetheless, PB is distinct from CB, as it also arises in the absence of CB, namely when OB is present, and aligns with it. Despite this relative autonomy, PB does depends on OB and CB in that when both fundamental biases are absent, PB does not arise either. In this subsection, we develop a proposal that predicts the attested patterns, thereby linking the two fundamental biases to the projected stages of the discourse.

Given the driving influence that OB and CB exert on PB in these generalizations, it seems reasonable to explore the idea that PB is not a semantic property of particular question forms, but is instead computed from these other fundamental biases via some kind of algorithm following general pragmatic principles. Along these lines, we propose that the key to understanding these generalizations is the notion of evidence. Specifically, we argue that the observed patterns can be derived from the interaction of three pragmatic principles. The first two are widely recognized and we take them to be quite uncontroversial; the third adopts a more substantive stance but also has a precedent in the literature on modality and evidentiality.

The first principle constitutes a requirement on the Common Ground and is presented in (39).

(39) COMMON GROUND SUPPORT

Information accepted in the Common Ground must be supported by sufficient evidence.

This principle is directly associated with Grice's (second) maxim of Quality, which instructs speakers not to assert what they lack adequate evidence for (Grice 1975). It is likewise tied to notions like acceptance (Stalnaker 1978, 1999), commitment (Brandom 1983; Krifka 2015; Geurts 2019), and knowledge (Williamson 2002), all of which come with a strong evidential requirement.

Another principle we will appeal to is presented in (40) and relates to the impact that a speaker's choice of utterance has on the interpretation.

(40) ECONOMY

Use a more complex or marked linguistic form only if a simpler or unmarked linguistic form does not have the same interpretational effects on the Common Ground.

This essentially dictates that there needs to be some payoff with regard to the information conveyed to justify a speaker's use of a more complex or marked utterance. This kind of constraint is suggested by the interplay between the maxims of Manner and Quantity, and has been widely adopted in neo-Gricean approaches to implicature computation (e.g., Horn 1984; Heim 1991; Levinson 2000; Frank and Goodman 2012).

The final principle we need is defined in (41) and states the basis for the ultimate alignment of PB.

(41) STRONGEST EVIDENCE

Discourse projections align with the strongest available evidence.

This principle essentially requires that, in situations of evidential conflict, speakers endorse the strongest evidence and ignore less reliable sources. It is compatible with qualitative theories of evidence that rank information sources and consider only the highest-ranked evidence, ignoring lower-ranked alternatives (Kratzer 1991; McCready 2015, ch.7). On the other hand, it contrasts with holistic theories, such as Bayesian updating, where a less probable piece of evidence affects the probabilities of all relevant hypotheses, including the most probable one.

We now demonstrate how, through the application of these pragmatic principles, the three generalizations of PB alignment can be derived. We start with Generalization 1 (see (28)), namely the observation that when OB and CB are absent altogether, as well as when the possible answers are equally supported by CB, then no PB is present. This pattern is captured by COMMON GROUND SUPPORT in the following way. When a polar question is asked and neither the question form nor

any other aspect of the context conveys a preference for one of the answers over the other, then there is no evidence that would justify projecting a favored answer to the Common Ground. In other words, in such a situation there would be no evidence that would justify the generation of PB, and so PB is absent. Consequently, only a neutral interpretation is available, and both answers are perceived as being equally likely to be added to the Common Ground.

Generalization 2 states that when OB or CB are the only bias present in the situation, then PB aligns with them (see (33)). We can account for this finding through an interaction between the COMMON GROUND SUPPORT and ECONOMY principles as follows. Imagine that a polar question conveys OB or CB, say a positive one (the negative case is parallel). Assuming that the two fundamental biases exert no effect on the Common Ground component of the Scoreboard (see Section 4 for formal details), it follows from ECONOMY that the question must also convey PB—otherwise, the speaker would have uttered a simpler question that conveys no OB or CB and still has no effect on the Common Ground. Moreover, the question uttered cannot convey negative PB, for which there is no supporting evidence, because this would violate COMMON GROUND SUPPORT. The only remaining option is for the question to convey positive PB.

Finally, Generalization 3 (stated in (37)) concern cases in which both OB and CB are present. If these two fundamental biases point in the same direction, then PB aligns with them, by the same reasoning that applies in the single-bias cases discussed in the previous paragraph. If, however, the biases are in epistemic conflict, then PB aligns with CB. We account for this last pattern by appealing to the ECONOMY and the STRONGEST EVIDENCE principles. As before, ECONOMY bans a neutral interpretation of the question. PB must therefore be present, and the task of the algorithm is to determine which of the two conflicting biases PB should align with. The key observation here is that OB and CB differ in discourse status, which results in a difference with regard to the evidential strength of these two biases. Specifically, OB is a *subjective* notion: it is an expression of the privately held belief of a speaker leading to a tentative discourse commitment (see Section 4 for a formal implementation). By contrast, CB is an *objective* notion, indicating the presence in the context of certain agreed-upon content recorded in the Common Ground.[18] Conventionally, the evidence standards for belief are lower than those for facts or knowledge (cf. Gettier 1963). As a result, STRONGEST EVIDENCE dictates that PB must align with CB.

4 Formal implementation

This section outlines a formal implementation of the discourse effects of questions conveying different bias flavors within the Scoreboard Model of Farkas and Bruce (2010). In doing so, we both sharpen the theoretical claims made in Section 3 and situate our proposal within a well-established pragmatic model.

We start out by defining a **context**, the formal counterpart to the scoreboard metaphor. A context c minimally consists of the following components: two sets of propositions representing the current discourse commitments of the speaker and the addressee (DC_c^A and DC_c^B, respectively), a set of propositions representing the current Common Ground (CG_c), a stack (sequence) representing what is currently at-issue or the topic of conversation (T_c, the Table), and projected future

[18] While this distinction is pragmatic in nature, it echoes the semantic distinction between "subjective" and "objective" epistemic modality (e.g., Lyons 1977: ch.17; Kratzer 1981; Portner and Rubinstein 2012; Trinh 2024; Koev 2025).

states of the first three components (DC_c^A*, DC_c^B*, and CG_c*).[19] The general form of a context is shown in (42).

(42) $c = \langle \mathsf{DC}_c^A, \mathsf{DC}_c^B, \mathsf{CG}_c, \mathsf{T}_c, \mathsf{DC}_c^A*, \mathsf{DC}_c^B*, \mathsf{CG}_c* \rangle$

There are certain relationships that must necessarily hold between the different components in order for a structure like this to qualify as a context. Here, we impose two such constraints. First, at any moment in the discourse, a Common Ground contains the propositions that are shared by the two discourse commitment sets (i.e., $\mathsf{CG}_c = \mathsf{DC}_c^A \cap \mathsf{DC}_c^B$). This means that a separate Common Ground component is not strictly needed, as it can be derived from the other components. Nevertheless, for convenience and following common practice, we will represent it explicitly.

Second, projected developments are monotonic in the sense that they inherit and extend their corresponding current counterparts (i.e., $\mathsf{DC}_c^A \subseteq \mathsf{DC}_c^A*$, $\mathsf{DC}_c^B \subseteq \mathsf{DC}_c^B*$, and $\mathsf{CG}_c \subseteq \mathsf{CG}_c*$). If a projected development fails to materialize, it must be reset to the current state of its corresponding component.

Given this structure of contexts, we propose to capture the three different bias flavors as follows. Starting with OB, this bias flavor is represented by the contrast between the current and the projected discourse commitments of the speaker (cf. Malamud and Stephenson 2015). Put in formal terms, a question Q conveys OB for p in a context c when $p \notin \mathsf{DC}_c^A$ and $p \in \mathsf{DC}_c^A*$. Since a projected discourse commitment presents a speaker as having the relevant belief, this amounts to $believe(t, A, p) \in \mathsf{DC}_c^A$, for some interval t that includes the utterance time in c. That is, OB presents the speaker as being committed to having held the belief that p, in line with the definition of OB in (9). This will hold even if p itself may never be added to the speaker's actual discourse commitments.

As for CB, we assume that it is derived from the information already contained in the current Common Ground. Building on our definition of CB in (11), a question Q conveys CB for p in a context c when the propositions in the current Common Ground provide sufficiently strong evidential support for p without also providing comparable evidential support for $\neg p$, a fact that we will express more formally by writing $p \succ_c \neg p$.[20] This includes both cases where p is plainly entailed by the Common Ground and cases where p is strongly suggested by it.

Finally, unlike the previous two bias flavors, PB directly guides the resolution of the current issue introduced by the host question (recall the definition in (17)). This means that PB is fundamentally linked to the development of the Common Ground. We therefore propose to capture the effect of this bias flavor by contrasting the current and the projected states of the Common Ground. Formally, a question Q conveys PB for p in a context c just when $p \notin \mathsf{CG}_c$ and $p \in \mathsf{CG}_c*$. More generally, PB serves to connect the various evidence sources associated with the two fundamental biases and to steer them toward the expected development of the Common Ground.

We now provide illustrations of the three main cases of PB alignment with OB and CB discussed in Section 3. The first case we discuss involves the absence of fundamental biases. It is

[19] One might wonder why there is no projected state of the Table component as well, namely T_c*. Although we know of no instance in the literature that employs this idea, a projected Table component may prove fruitful in the analysis of contrastive topics, which signal that the current issue has been answered non-exhaustively and thus predict the shape of the follow-up issue (Büring 2003; Tomioka 2010; Constant 2014; Kamali and Krifka 2020).

[20] There are many different ways of recruiting Common Ground information and making it bear on a given proposition. One straightforward strategy is to use probabilistic conditionalization on the Common Ground to gauge the likelihood of a proposition against that of its negation while also requiring it to meet some high contextual threshold. This leads to the following definition: $p \succ_c \neg p$ iff $\mathsf{Pr}(p \mid \bigcap \mathsf{CG}_c) > \mathsf{Pr}(\neg p \mid \bigcap \mathsf{CG}_c)$ and $\mathsf{Pr}(p \mid \bigcap \mathsf{CG}_c) \geq \theta_c$.

illustrated on the positive polar question in (43), where the dynamic notation $c + \phi = c'$ indicates that uttering ϕ in the input context c results in the output context c'.

(43) *Case 1: No fundamental biases*

a. $c+$ *Is it raining outside?* is defined only if $\neg rain \not\succ_c rain$.

b. If defined, $c+$ *Is it raining outside?* $= c'$, where c' is like c except that $\mathsf{T}_{c'} = \mathsf{push}(\{rain, \neg rain\}, \mathsf{T}_c)$.

Positive polar questions impose a felicity condition requiring that there be no sufficiently strong contextual evidence supporting the negative question alternative over the positive such alternative (Büring and Gunlogson 2000; Sudo 2013). Once this condition is met, the only semantic effect of such questions is to push the issue they raise on top of the Table. In (43), this is the issue of whether it is raining outside.[21] If the issue is resolved by virtue of the addressee picking out one of the two polar alternatives, it will be popped off the Table and will no longer be relevant. Most importantly, our proposal from Section 3.3 predicts no PB (Generalization 1), so no Common Ground distinct from the current one is projected.

The second case involves contexts with a single fundamental bias. It is illustrated with the high negation questions in (44), under the assumption that no negative CB is present in the context.

(44) *Case 2: Single fundamental bias*

a. $c+$ *Isn't it raining outside?* is defined only if $rain \not\succ_c \neg rain$.

b. If defined, $c+$ *Isn't it raining outside?* $= c'$, where c' is like c except that

- $\mathsf{T}_{c'} = \mathsf{push}(\{rain, \neg rain\}, \mathsf{T}_c)$,
- $\mathsf{DC}^A_{c'}* = \mathsf{DC}^A_c* \cup \{rain\}$, and
- $\mathsf{CG}_{c'}* = \mathsf{CG}_c* \cup \{rain\}$.

The felicity condition imposed by high negation questions is similar to that of positive polar questions but with the polarities reversed: it is evidence for the positive question alternative—corresponding to 'It is raining outside' in (44)—that is banned in this case. Granted that, the semantic effects of such questions are fairly rich. In addition to placing the partition they delineate on the Table, they also obligatorily convey positive OB, modeled as adding the positive question alternative to the speaker's projected discourse commitments. Given this particular bias layout (positive OB, no CB), our proposal predicts the presence of positive PB (recall Generalization 2 and especially example (15)), modeled as extending the projected Common Ground.[22] We reiterate that this last bias flavor stems not directly from the particular question form, but rather from the shape of the resulting context and the general pragmatic principles discussed in Section 3.

Our third and last case is the most interesting one and handles conflicting fundamental biases. To illustrate it, we use polar questions with epistemic *really*, as in (45).

[21] Farkas and Bruce (2010) propose that (neutral) positive polar questions additionally project both of their answers, resulting in a set of two projected Common Grounds. However, alternative questions are arguably distinguished from positive polar questions precisely by the fact that they project both answers onto the Common Ground (cf. Biezma and Rawlins 2012; Krifka 2015). We therefore assume that positive polar questions introduce no projected components.

[22] Notice that in (44) both projected sets in the output context (i.e., $\mathsf{DC}^A_{c'}*$ and $\mathsf{CG}_{c'}*$) are derived from the *projected* respective sets of the input context (i.e., DC^A_c* and CG_c*) rather than from the actual respective sets of the input context (i.e., DC^A_c and CG_c). If it is assumed that projected states are reset to their actual counterparts prior to a new utterance, the two options are equivalent. Similar remarks apply to (45) below.

(45) *Case 3: Conflicting fundamental biases*

a. $c +$ *Is it really raining outside?* is defined only if $rain \succ_c \neg rain$.

b. If defined, $c +$ *Is it really raining outside?* $= c'$, where c' is like c except that

- $\mathsf{T}_{c'} = \mathsf{push}(\{rain, \neg rain\}, \mathsf{T}_c)$,
- $\mathsf{DC}^A_{c'}* = \mathsf{DC}^A_c* \cup \{\neg rain\}$, and
- $\mathsf{CG}_{c'}* = \mathsf{CG}_c* \cup \{rain\}$.

Really-questions require that the positive question alternative be supported by the contextual evidence (Romero and Han 2004). Provided that this condition is satisfied, they push an issue on the Table and also obligatorily convey negative OB. In such cases of conflicting fundamental biases (positive CB, negative OB), our proposal predicts the presence of positive PB (see Generalization 3 and example (36)). Hence, in (45) the alternative 'It is raining outside' is added to the projected Common Ground.

By way of concluding this section, we make two general remarks about the formal implementation outlined above. Our first remark concerns the STRONGEST EVIDENCE principle, manifested as Generalization 3, which states that PB aligns with CB rather than OB in cases of conflict between the latter two. In Section 3, we proposed that this preference arises from the mutually accepted nature of the evidence associated with CB. Our point is that the technical implementation presented in this section does justice to this explanation. That is, CB is derived from Common Ground information, making it the preferred basis for PB. In contrast, OB affects the projected commitments of individual speakers, so the accompanying evidence need not be mutually accepted by all discourse participants. This is why PB relies on it only as a second choice.

Our second remark concerns how our proposal emphasizes some important features of the notion of at-issueness. Following a prominent approach to at-issueness (Simons et al. 2010; Koev 2022, ch.5), here we conceive of this notion as a *question*, namely a semantic object placed on the Table and effecting a partition on the set of live possibilities. Despite the key role of questions in determining issues, the phenomenon of question bias has rarely been directly linked to at-issueness. The new flavor of bias spelled out in this paper, namely PB, provides the missing link. It shows that the two fundamental biases (OB and CB) are merely a way to lead the conversation toward resolving the current issue. It also demonstrates that issues need not always be balanced, or introduced by neutral questions, as under the classical theory. Issues introduced by questions with PB, as in (44) and (45), can be called **unbalanced**, meaning that one of the possible outcomes is already given priority.[23] The concept of PB thus not only unifies the two fundamental biases but also sheds new light on the notion of at-issueness.

5 Conclusion

The main goal of this paper was to make a case for adding projected bias to the established question bias flavors of original and contextual bias. We pursued this goal by first proposing clear definitions of the different biases, and then, through a series of diagnostics, identifying their alignment characteristics. We were particularly interested in determining whether projected bias merely mirrored

[23] Compare Romero and Han (2004)'s "unbalanced partition" terminology, referring to questions raising metaconversational issues due to the presence of a VERUM operator. Here we repurpose their terminology to refer to regular yet skewed partitions resulting from the presence of PB.

one of the more established biases or a random choice between them, as would be expected if it were completely parasitic on them. We found that (i) when neither of the established biases were present or when the possible answers were equally supported by one of them, projected bias did not arise; (ii) when only one of the established biases were present, then projected bias was present and aligned with it; and (iii) when both of the established biases were present and in conflict, then projected bias was present and aligned with contextual bias. We interpreted this pattern of results as indicating that projected bias is a distinct flavor of question bias. That is, rather then being simply an additional effect of one of the established biases, projected bias falls out from the other two bias flavors through general pragmatic principles. We went on to articulate a proposal specifying which pragmatic principles are involved and how their application leads to the generation of projected bias. The most original part of this proposal was related to STRONGEST EVIDENCE. According to this principle, when there is a conflict between original bias and contextual bias, projected bias aligns with the latter because it is based on objective knowledge, whereas the former is based merely on subjective belief. More generally, we take this work to provide a clearer explanation for why question bias should exist in the first place. On the Question Under Discussion approach and its implementation in the Scoreboard Model, the goal of any discourse is to resolve the current issue, which in a context where a question has just been asked equates to determining which answer to this question is to be selected. The role that question bias plays in this process is made clear through the existence of projected bias, which weights the discourse resolution toward one of the possible answers.

Acknowledgments

We would like to thank Daniel Gutzmann and three anonymous reviewers at *Languages* for their very helpful feedback. This research was funded by the Israel Science Foundation (project 3464/25).

References

Abels, K. (2003). Who gives a damn about minimizers in questions? In *Proceedings of Semantics and Linguistic Theory 13*, pp. 1–18.

AnderBois, S. (2019). Negation, alternatives, and negative polar questions in American English. In K. von Heusinger, V. E. O. Gaspar, and M. Zimmermann (Eds.), *Questions in Discourse, Volume 1: Semantics*, pp. 118–171. Brill.

Arnhold, A., B. Braun, and M. Romero (2021). Aren't prosody and syntax marking bias in questions? *Language and Speech 64*(1), 141–180.

Biezma, M. and K. Rawlins (2012). Responding to alternative and polar questions. *Linguistics and Philosophy 35*, 361–406.

Biezma, M. and K. Rawlins (2017). Rhetorical questions: Severing questioning from asking. In *Proceedings of Semantics and Linguistic Theory 27*, pp. 302–322.

Bill, C. and T. Koev (2022). *Really*: Ambiguity and question bias. In *Proceedings of Sinn und Bedeutung 26*.

Bill, C. and T. Koev (2025). Bias in tag questions. In T. Trinh, A. Benz, D. Goodhue, K. Yatsushiro, and M. Krifka (Eds.), *Biased Questions: Experimental results and theoretical modelling*, pp. 145–178. Berlin: Language Science Press.

Bolinger, D. (1978). Yes–no questions are not alternative questions. In H. Hiż (Ed.), *Questions*, pp. 87–105. Dordrecht: Springer.

Brandom, R. (1983). Asserting. *Noûs 17*(4), 637–650.

Büring, D. (2003). On d-trees, beans, and b-accents. *Linguistics and Philosophy 26*(2), 511–545.

Büring, D. and C. Gunlogson (2000). Aren't positive and negative polar questions the same? Technical report, University of California, Santa Cruz.

Caponigro, I. and J. Sprouse (2007). Rhetorical questions as questions. In *Proceedings of Sinn und Bedeutung 11*, pp. 121–133.

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

Ciardelli, I., J. Groenendijk, and F. Roelofsen (2019). *Inquisitive Semantics*. Oxford: Oxford University Press.

Claus, B., A. M. Meijer, S. Repp, and M. Krifka (2017). Puzzling response particles: An experimental study on the german answering system. *Semantics & Pragmatics 10*(19), 1–56.

Constant, N. (2014). *Contrastive Topic: Meanings and Realizations*. Ph.D. dissertation, University of Massachusetts Amherst.

Farkas, D. and K. B. Bruce (2010). On reacting to assertions and polar questions. *Journal of Semantics 27*, 81–118.

Farkas, D. F. (2026). Rhetorical questions revisited. In *Proceedings of Chicago Linguistic Society 60*.

Farkas, D. F. and F. Roelofsen (2017). Division of labor in the interpretation of declaratives and interrogatives. *Journal of Semantics 34*, 237–289.

Farkas, D. F. and F. Roelofsen (2019). Polarity particles revisited. *Semantics & Pragmatics 12*(15), 1–16.

Frana, I. and K. Rawlins (2019). Attitudes in discourse: Italian polar questions and the particle *mica*. *Semantics & Pragmatics 12*(16), 1–48.

Frank, M. C. and N. D. Goodman (2012). Predicting pragmatic reasoning in language games. *Science 336*(998).

Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis 23*(6), 121–123.

Geurts, B. (2019). Communication as commitment sharing: Speech acts, implicatures, common ground. *Theoretical Linguistics 1–30; 111–125*(1-2), 45.

Goodhue, D. (2022a). All focus is contrastive: On polarity (verum) focus, answer focus, contrastive focus and givenness. *Journal of Semantics 39*, 117–158.

Goodhue, D. (2022b). Isn't there more than one way to bias a polar question? *Natural Language Semantics 30*, 379–413.

Goodhue, D. and M. Wagner (2018). Intonation, *yes* and *no*. *Glossa: A Journal of General Linguistics 3(1)*(5), 1–45.

Grice, P. (1975). Logic and conversation. In M. Ezcurdia and R. J. Stainton (Eds.), *The Semantics-Pragmatics Boundary in Philosophy*, pp. 64–75. Broadview Press.

Groenendijk, J. and M. Stokhof (1984). On the semantics of questions and the pragmatics of answers. In F. Landman and F. Veltman (Eds.), *Varieties of Formal Semantics*, pp. 143–170. Dordrecht: Foris.

Guerzoni, E. (2003). *Why Even Ask? On the Pragmatics of Questions and the Semantics of Answers*. Ph.D. dissertation, MIT.

Guerzoni, E. (2004). *Even*-NPIs in yes/no questions. *Natural Language Semantics 12*, 319–343.

Gunlogson, C. (2003). *True to Form: Rising and Falling Declaratives as Questions in English*. New York & London: Routledge.

Gutzmann, D., K. Hartmann, and L. Matthewson (2020). Verum focus is verum, not focus: Cross-linguistic evidence. *Glossa: A Journal of General Linguistics 5*(1), 51.

Hamblin, C. L. (1973). Questions in Montague English. *Foundations of Language 10*, 41–53.

Han, C.-b. (2002). Interpreting interrogatives as rhetorical questions. *Lingua 12*, 201–229.

Heim, I. (1984). A note on negative polarity and downward entailingness. In *Proceedings of the North East Linguistics Society*, Volume 14, pp. 98–107.

Heim, I. (1991). Artikel und Definitheit. In A. v. Stechow and D. Wunderlich (Eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, pp. 487–535. Berlin: de Gruyter.

Höhle, T. N. (1992). Über Verum-Fokus im Deutschen. In J. Jacobs (Ed.), *Informationsstruktur und Grammatik*, pp. 112–141. Opladen: Westdeutscher Verlag.

Holmberg, A. (2013). The syntax of answers to polar questions in english and swedish. *Lingua 128*, 31–50.

Horn, L. R. (1984). Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In D. Schiffrin (Ed.), *Meaning, Form, and Use in Context: Linguistic Applications*, pp. 11–42. Washington, DC: Georgetown University Press.

Jeong, S. (2018). Intonation and sentence type conventions: Two types of rising declaratives. *Journal of Semantics 35*, 305–356.

Jeong, S. (2021). Deriving dual dimensions of bias: Preposed negation questions with even. *Journal of Semantics 38*, 49–94.

Kamali, B. and M. Krifka (2020). Focus and contrastive topic in questions and answers, with particular reference to Turkish. *Theoretical Linguistics 46*, 1–71.

Karttunen, L. and S. Peters (1979). Conventional implicature. In C. Oh and D. A. Dinneen (Eds.), *Syntax and Semantics 11: Presupposition*, pp. 1–56. London: Academic Press.

Koev, T. (2022). *Parenthetical Meaning*. Oxford: Oxford University Press.

Koev, T. (2025). An argument for a strong force semantics of *believe*. *Linguistic Inquiry*, 1–12.

Kramer, R. and K. Rawlins (2011). Polarity particles: An ellipsis account. In *Proceedings of the North East Linguistics Society*, Volume 39, pp. 479–4927.

Kratzer, A. (1981). The notional category of modality. In H. J. Eikmeyer and H. Rieser (Eds.), *Words, Worlds, and Contexts. New Approaches in Word Semantics*, pp. 38–74. Berlin/New York: Walter de Gruyter.

Kratzer, A. (1991). Modality. In A. von Stechow and D. Wunderlich (Eds.), *Semantik/Semantics: An International Handbook of Semantics*, pp. 639–650. de Gruyter.

Krifka, M. (1995). The semantics and pragmatics of polarity items. *Linguistic Analysis 25*, 1–49.

Krifka, M. (2013). Response particles as propositional anaphors. In *Proceedings of Semantics and Linguistic Theory 23*, pp. 1–18.

Krifka, M. (2015). Bias in commitment space semantics: Declarative questions, negated questions, and question tags. In *Proceedings of Semantics and Linguistic Theory 25*, pp. 328–345.

Kripke, S. A. (1963). Semantical analysis of modal logic i: Normal modal propositional calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik 9*, 67–96.

Ladd, D. R. (1981). A first look at the semantics and pragmatics of negative questions and tag questions. In *Proceedings of Chicago Linguistic Society 17*, pp. 164–171.

Ladusaw, W. A. (1980). *Polarity Sensitivity as Inherent Scope Relations*. Ph.D. dissertation, University of Texas, Austin.

Laka, I. (1990). *Negation in Syntax: On the Nature of Functional Categories and Projections*. Ph.D. dissertation, MIT.

Levinson, S. C. (2000). *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. London: MIT Press.

Lyons, J. (1977). *Semantics, vol. 2*. Cambridge: Cambridge University Press.

Malamud, S. A. and T. Stephenson (2015). Three ways to avoid commitments: Declarative force modifiers in the conversational scoreboard. *Journal of Semantics 32*(2), 275–311.

McCready, E. (2015). *Reliability in Pragmatics*. Oxford: Oxford University Press.

Pope, E. N. (1972). *Questions and Answers in English*. Ph.D. dissertation, MIT.

Portner, P. and A. Rubinstein (2012). Mood and contextual commitment. In *Proceedings of Semantics and Linguistic Theory 22*, pp. 461–487.

Reese, B. J. (2007). *Bias in Questions*. Ph.D. dissertation, University of Texas at Austin.

Repp, S. (2012). Common ground management: Modal particles, illocutionary negation and verum. In D. Gutzmann and H.-M. Gärtner (Eds.), *Beyond Expressives: Explorations in Use-conditional Meaning*, pp. 231–274. Leiden: Brill.

Roberts, C. (2012). Information structure: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics 6*(5), 1–69.

Roelofsen, F. and D. F. Farkas (2015). Polarity particle responses as a window onto the interpretation of questions and assertions. *Language 91*(2), 359–414.

Roelofsen, F., N. Venhuizen, and G. W. Sassoon (2012). Positive and negative questions in discourse. In *Proceedings of Sinn und Bedeutung 17*, pp. 455–472.

Rohde, H. (2006). Rhetorical questions as redundant interrogatives. In *San Diego Linguistics Papers*, pp. 134–168.

Romero, M., A. Arnhold, B. Braun, and F. Domaneschi (2017). Negative polar question types in english. In *Proceedings of the North East Linguistics Society 47*, Volume 3, pp. 35–48.

Romero, M. and C.-H. Han (2004). On negative *yes/no* questions. *Linguistics and Philosophy 27*, 609–658.

Rooth, M. E. (1985). *Association with Focus*. Ph.D. dissertation, University of Massachusetts at Amherst.

Rudin, D. (2022). Intonational commitments. *Journal of Semantics 39*, 339–383.

Sailor, C. (2013, January 3). Questionable negation. Technical report, LSA Annual Meeting.

Scheffler, T. and S. A. Malamud (2023). Won't you? *Linguistics and Philosophy 46*, 1285–1327.

Silk, A. (2020). Expectation biases and context management with negative polar questions. *Journal of Philosophical Logic 49*, 51–92.

Simons, M., J. Tonhauser, D. Beaver, and C. Roberts (2010). What projects and why. In *Proceedings of Semantics and Linguistic Theory 20*, pp. 309–327.

Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and Semantics, Volume 9: Pragmatics*, pp. 315–332. New York: Academic Press.

Stalnaker, R. (1999). On the representation of context. In R. Stalnaker (Ed.), *Content and Context: Essays on Intentionality in Speech and Thought*, pp. 96–113. Oxford: Oxford University Press.

Sudo, Y. (2013). Biased polar questions in english and japanese. In D. Gutzmann and H.-M. Gaertner (Eds.), *Beyond Expressives: Explorations in Use-Conditional Meaning*, pp. 275–295. Leiden: Brill.

Tarski, A. (1956). The concept of truth in formalized languages. In A. Tarski (Ed.), *Logic, Semantics, Metamathematics*, pp. 152–278. Clarendon Press.

Tomioka, S. (2010). A scope theory of contrastive topics. *Iberia: An International Journal of Theoretical Linguistics 2*(1), 113–130.

Trinh, T. (2014). How to ask the obvious: A presuppositional account of evidential bias in english yes/no questions. In L. Crnič and U. Sauerland (Eds.), *The Art and Craft of Semantics: A Festschrift for Irene Heim*, pp. 227–249. MITWPL 71.

Trinh, T. (2024). Strengthened, and weakened, by belief. *Linguistics and Philosophy 47*, 37–76.

van Rooy, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy 26*(6), 727–763.

van Rooy, R. and M. Šafářová (2003). On polar questions. In *Proceedings of Semantics and Linguistic Theory 13*, pp. 292–309.

Wilkinson, K. (1996). The scope of *even*. *Natural Language Semantics 4*, 193–215.

Williamson, T. (2002). *Knowledge and Its Limits*. Oxford: Oxford University Press.