

Understanding focus: Pitch, placement, and coherence*

Julian J. Schlöder
ILLC / Department of Philosophy
University of Amsterdam
julian.schloeder@gmail.com

Alex Lascarides
School of Informatics
University of Edinburgh
alex@inf.ed.ac.uk

Abstract This paper presents a novel account of focal stress and pitch contour in English dialogue. We argue that one should analyse and treat focus and pitch contour jointly, since (i) some pragmatic interpretations vary with contour (e.g., whether an utterance accepts or rejects; or whether it implicates a positive or negative answer); and (ii) there are utterances with identical prosodic focus that in the same context are infelicitous with one contour, but felicitous with another. We offer an account of two distinct pitch contours that predicts the correct felicity judgements and implicatures, outclassing other models in empirical coverage or formality. Prosodic focus triggers a presupposition, where what is presupposed and how the presupposition is resolved depends on prosodic contour. If resolving the presupposition entails the proffered content, then the proffered content is *uninteresting* and hence the utterance is infelicitous. Otherwise, resolving the presupposition may lead to an implicature. We regiment this account in SDRT.

Keywords: focus, prosody, dialogue, coherence

1 Introduction

Our goals in this paper are twofold. First, we demonstrate that it is a mistake to give a semantic account of prosodic focus while ignoring the overall pitch contour of an utterance. We mean this as both a *challenge* to prior accounts that claim to model prosodic focus without considering contour and as a *methodological* point.

* This work has received funding from the European Research Council (grant no. 758540, H2020) and the European Commission's Marie Curie Actions (grant no. 607062, FP7). We are grateful to Nicholas Asher, Bob Ladd and Mark Steedman, to the audiences of the *Rutgers Semantics and Pragmatics Colloquium 2017* and the *ISSLaC 3 Workshop* at the University of Münster, and to the reviewers for *Semantics & Pragmatics* for valuable comments on earlier versions of this material. We thank the editor, Elizabeth Coppock, for feedback that greatly improved the presentation of the paper. Finally, we want to express our particular gratitude to Bob Ladd for lending his voice to our examples.

The challenge is to explain why certain utterances with identical prosodic focus (i.e., placement of the nuclear accent) are infelicitous with one pitch contour but felicitous with another. For example, if only prosodic focus is annotated (here by underlining), one is inclined to judge (1b) as felicitous and (1b′) as infelicitous (Roberts 2012: p34).

- (1) a. A: Who likes Michael?
 b. B: Nobody likes Michael.
 b.′ B: Nobody likes Michael.

However, when we consider (1b′) with different pitch contours in (@2), the judgements are more fine-grained. (Because of their familiarity, we use ToBI labels for now, but see Figure 1 for the precise contour of (@2b′) and Section 2 for discussion on how we annotate pitch contour.)

- (@2) a. Harvey: Who likes Michael?
 #b. Jessica: Nobody likes Michael.
 H* LL%
 b.′ Jessica: Nobody likes Michael.
 L*+HL- H%

(@2ab) is infelicitous but (@2ab′) is acceptable — indeed it’s natural in a context where Jessica thinks anyone liking Michael is absurd. So it is a mistake to judge (1b′) as infelicitous based on focus alone. Since contour is the only variant here, an account that ignores contour can’t model the difference between (@2b) and (@2b′). In Section 3, we square these and similar data against a number of established accounts of focus.

Our methodological point is that to empirically test an account of prosodic focus, one must consider data that is annotated with pitch contour. An account of focus that is tested against data not mentioning contour may in fact model an idiosyncratic set of interpretations that derive from assessing the data according to the different contours one unsystematically associates with prosodic focus in the absence of annotation for contour. Based on (1) and (@2), one may think that if one annotates focus without mentioning contour, one assumes the contour is H* LL%. But this is not so. In Section 3 we show that some accounts of prosodic focus make the correct predictions for one contour in some examples but correct predictions for a quite different contour in others.

Our second goal is to construct a formally precise model for focus that respects pitch contour and thereby explains (@2) and similar examples (see Section 2 for the type of data we aim to model). We reevaluate the basic data on both focus and contour in this field, leading us to reject some received concepts (namely, *question-answer congruence*, *contrasting alternatives* and *givenness*; see Section 3). But this also means that we cannot encompass here the whole extent of puzzles and problems



Figure 1 Pitch contour of (@2b'). Thick lines indicate maxima in intensity (measured in decibels). These data were extracted with Praat (Boersma & van Heuven 2001).

raised in a substantial prior literature addressing these concepts. For instance, we treat here neither focus-sensitive operators like *only* (a challenge for accounts of focus) nor prenuclear pitch accents (a challenge for accounts of pitch contours). We also make no claims about any language other than English. However, we do consider — and formalise an account that explains — some implicatures that vary with contour: *fall–rise contour* utterances appear to frequently carry *as-opposed-to* implicatures not typically associated with the *falling contour*; and the fall–rise contour appears to be more suggestive of rejection moves or negative answers (Ladd 1980, de Marneffe & Tonhauser 2019).

We find in Section 3 that extant accounts of focus predict too many utterances to be infelicitous (e.g., (@2b')). Therefore, we search for a conception of focus that is less restrictive than the received views. In this search, we hold ourselves to the standard of making use of independently motivated and tested theories whenever possible. Our account combines established theories of presupposition and coherence.

Specifically, we consider focus to be a presupposition trigger (Jackendoff 1972, Geurts & van der Sandt 2004a), but let pitch contour influence *what* is presupposed and *how* this presupposition is resolved. We let the presupposition triggered by a fall–rise contour have *underspecified modality* (whether its content is true, false or possible), but the presupposition triggered by focus with a falling contour is not similarly underspecified. Additionally, the falling contour requires that the presupposition is resolved such that the utterance's foreground content *continues* or *elaborates* on the presupposed information, whereas the fall–rise contour demands that the foreground stands in a *contrast* relation to the presupposition. While presupposition accounts of focus have been challenged (Dryer 1996, Rooth 1999, Sæbø 2016), the problematic examples have been discussed without mentioning contour — our contour-sensitive semantics accounts for them (Section 4.6).

To explain cases like (@2), we introduce a new take on the intuition that you cannot focus what is given. To wit, you cannot focus that which you also present as

presupposed. We trace this principle to Bolinger (1972, 1985), who argues persuasively that focus follows *interest* and what is obvious cannot be interesting. However, we formalise neither interest nor obviousness, instead saying that what is focal (\approx interesting) cannot be presupposed (\approx obvious). This allows us to remain entirely within independently motivated theories of presupposition and coherence. Since we let the presupposition triggered by focus vary with contour, our predictions about focus vary with contour — as is required for cases like (@2).

We think that modelling focus and contour jointly is the way to go, but acknowledge that our challenges do not conclusively rule out the following option for making them independent. One may state a semantics for focus that over-generates felicitous utterances (including (1b')) and sort out the missing infelicities in a separate model for contour (separating (@2b) from (@2b')). However, our charge against all extant accounts of focus in Section 3 is that they predict some utterances to be infelicitous that actually are felicitous (with specific contours). So adding a contour semantics to one of them is not enough; the model of focus must also be revised to permit additional felicities. Our methodological point stands either way: to determine the empirical adequacy of an account of focus, the data needs to be annotated with both focus and contour, regardless of how these are modelled.

There are prior accounts (e.g., Pierrehumbert & Hirschberg 1990, Westera 2017) that cover data marked with both focus and contour, but fail to *fully formalise* how the semantics they assign to linguistic and prosodic form supports different implicatures and/or felicity judgements in different contexts.¹ These accounts are complemented by models that *are* formal, but ignore focus (e.g., Schlöder & Lascarides 2015), ignore contour (e.g., Roberts 2012), or apply to only one pitch contour (e.g., Reese 2007, Constant 2012). Steedman (2014) takes an intermediate position in having a formal semantics, but only an informal pragmatics (we explain the problems with this in Section 3.6).

We aim to develop a formal theory that considers *both* focus and pitch contour *and* formally derives implicatures and infelicity judgements.² Our formal account is couched in Segmented Discourse Representation Theory (SDRT, Asher & Lascarides 2003), particularly its existing models for presuppositions-as-anaphora and the coherence relations *Continuation*, *Elaboration* and *Contrast*. Crucially, SDRT models the interaction between presupposition and discourse coherence (Asher &

1 Formalisation is a virtue in part because it allows one to generate testable predictions. We only discuss accounts in Section 3 about which we can compute testable predictions. While Westera (2018) is more formal than Westera (2017), he only deals with the H* L- H% contour and does not formalise the notions needed to derive predictions from the accounts of the falling and fall-rise contours in Westera (2017).

2 Some prior work goes in this direction (Büring 2003), but we argue that it is too attached to some received concepts we reject (see Section 3).

Lascarides 1998), allowing us to derive subtle implicatures and constraints on felicity from different presuppositions.

We proceed as follows. In Section 2 we clarify our methodology. In Section 3 we elaborate the empirical shortcomings of prior accounts. In Section 4, we informally motivate and describe our proposal for a semantics of focus and contour. Throughout, we discuss how and why our semantics replicates or outclasses the predictions made by prior accounts. In Section 5 we formalise our model in SDRT and compute some interpretations.

2 Data and methodology

The study of prosody is fraught with methodological issues. Some researchers aim for compositional meanings of discretised pitch accents (e.g., via ToBI categories; Silverman et al. 1992). Others provide evidence for nondecomposable *tunes* with noncompositional meanings (Ladd 1980, Bolinger 1982, Calhoun 2007). Either way, there is ambiguity in (and debate on) how one carves up the data. If pitch contours are decomposable, then one may argue about whether the accents L+H* and H* are the same (Watson et al. 2008); if they aren't decomposable, one needs to decide how and where to demarcate different tunes (Calhoun 2007).

We make no contribution to these debates. We use the term *pitch contour* (which we take to be methodologically neutral) to denote the intonational form of an utterance. The semantics we propose are for pitch contours, and we do not explore whether that semantics could be composed from the contour's parts (whatever these parts are or may be). We consider only two, easily distinguishable pitch contours. To wit, we analyse single-clause utterances in which there is a single prominent ('focal') constituent such that either

- (i) the pitch rises on this constituent, the strongest stress is on the highest part of this rise, and then it falls after (*falling contour*); or
- (ii) the pitch first falls and then steeply rises on this constituent, the strongest stress is on this rise, and the utterance ends in a final rise (*fall-rise contour*).

All our data are constructed. We annotate for contour by underlining the word on which the prominent stress is placed, adding its pitch contour type as a subscript — either *fall* or *f-r*. This notation obscures intensity: the examples should be intonated with strong and prominent stress. To aid understanding and evaluation of our claims, we provide an audio file for each data point. The audio can be obtained by clicking the @ symbol in each example.³

³ List of audio data: <http://homepages.inf.ed.ac.uk/alex/tunes/tunes.pdf>




Figure 2 Pitch contours of the falling contour of (@3a) (left), and the fall–rise contour of (@3b) (right). Thick lines indicate maxima in intensity (measured in decibel).

For example, the following utterance (@3a) is intonated with the *falling contour* and the utterance (@3b) with the *fall–rise contour* (see Figure 2).

- (@3) a. I’m a millionaire_{fall}.
 b. I’m a millionaire_{f-r}.

The falling contour corresponds roughly to what in ToBI would be annotated as $H^* LL\%$ and the fall–rise contour to $L^*+H L- H\%$. In our discussion of prior accounts in Section 3, some relevant examples satisfy this correspondence (and are annotated as such). However, we do not commit to the idea that all our examples can be reduced to these ToBI accents. Also, we do not consider utterances with multiple-focus constructions or prenuclear foci in this paper. We take such utterances to have *different contours* than the two we consider — since we remain agnostic about compositionality and the correct discretised intonational units until such time as the data and its analysis are better understood, we are forced to ignore them.

Following Steedman (2014) and others, our goal is to associate each intonational form (i.e., pitch contour) with a *single* semantic value, with its distinct implicatures in distinct contexts being derivable via independently motivated principles of pragmatics. The target data for our formal model are the felicity judgements and implicatures that one intuitively associates with a particular utterance with particular intonation in a particular context. Such intuitions, too, are sometimes vague and subject to debate; notably, there appears to be some variation between British and American English speakers (Steedman 2014: *passim*). Nevertheless, there are strong and robust intuitions regarding the felicity and meaning of some contours in some contexts: (@4) demonstrates the intuition that following a *wh*-question, focal placement should be *congruent* to the question (Halliday 1967); in (@5), the fall–rise contour leads to an *as-opposed-to* implicature (Pierrehumbert & Hirschberg 1990) (we use \rightsquigarrow to indicate implicatures and sometimes use $\not\rightsquigarrow$ to record what is *not* implicated to highlight differences in interpretations).

- (@4) a. Harvey: Who likes Michael?
b. Jessica: Rachel_{fall} likes Michael.
#b.' Jessica: Rachel likes Michael_{fall}.
- (@5) a. Louis: Is Harvey going to fire me?
b. Donna: Harvey_{f-r} is not going to fire you.
 \rightsquigarrow *but someone else is*

Even when judgements about felicity or implicatures in a specific context are clear, such judgements are always *defeasible* in the following sense. If further utterances (prior or posterior context) are present, inferences about implicatures or felicity may get revised. For example, if (@4a) is preceded by Harvey saying *I bet everyone who likes Michael also likes John*, (@4b') is acceptable. But this doesn't detract from the fact that if (@4b) is all he says, (@4b') is infelicitous. We contend that any adequate model of pragmatic interpretation should predict *both* the infelicity of (@4b') and its felicity in other contexts, as well as implicatures as in (@5).

In what follows, we construct such a model. Our target data is derived from a literature consensus of intuitions and our own judgements. We have verified our data with native English speakers, including acknowledged experts in the field (Mark Steedman and Bob Ladd).

3 Problems with prior accounts

In Section 1, we presented one example where utterances with identical prosodically marked focus (i.e., placement of the nuclear accent) are infelicitous with the falling contour but felicitous with the fall-rise contour. This troubles accounts that do not model contour alongside focus (see Sections 3.1, 3.2 and 3.3). Beyond the specific accounts we criticise, we wish to establish our methodological point: considering focus without pitch contour leads to confusion about what the modelling target is (Section 3.4). This is tentatively acknowledged (Beaver & Clark 2009: p47; Roberts 2012: p29), but we want to put pressure on the issue. Some accounts consider both focus and contour, but we demonstrate some empirical shortcomings (Sections 3.5, 3.6 and 3.7).

3.1 Question-answer congruence / Roberts 2012

Example (@4) motivates the principle of *question-answer congruence* (Halliday 1967, Büring 2003, Beaver & Clark 2009, Roberts 2012): that focus indicates the *wh*-question an assertion answers. The congruent question is the one obtained by substituting the focal constituent with a *wh*-element. So (@4b) succeeds in answering (@4a) but (@4b') does not. Now, compare (1) and (@2).

- (1) a. A: Who likes Michael?
 b. B: Nobody likes Michael.
 b.' B: Nobody likes Michael.
- (@2) a. Harvey: Who likes Michael?
 #b. Jessica: Nobody likes Michael.
 b.' Jessica: Nobody likes Michael.

H* LL%
L*+HL- H%

All congruence accounts we are aware of predict (1b') to be infelicitous — hence, falsely predict (@2b') to be infelicitous. We demonstrate this for Roberts' (2012) account. The question congruent to (@2b') is *Who is liked by nobody?*, so (@2b') is felicitous only if this question can be *accommodated* as being part of what is under discussion. (Similar appeals to accommodation are made across the congruence literature; also see Section 3.5.) There must be constraints on what questions can be accommodated in a given context, lest all focus placements be felicitous. Roberts gives the constraint (6) (our paraphrase).

- (6) A question can be accommodated if all complete answers to it partially answer the question-under-discussion.
 (Roberts 2012: pp14–15, def. 10g(iii)).

The relevant notions of answerhood are as follows. The denotation of a question is obtained by replacing all *wh*-elements in that question with free variables and computing the set of all propositions where the free variables have been instantiated with suitable referents (p10). If, say, the discourse referents are $D = \{r, k, m\}$ then the denotation of *Who likes Michael?* is (7).

- (7) $\{\text{like}(r, m), \text{like}(k, m), \text{like}(m, m)\}$.

Then, a proposition is a *complete answer* to a question if it decides (i.e., entails either truth or falsity of) *each* proposition in the question's denotation. A *partial answer* decides at least one proposition in the denotation (p11).

Now, the congruent question to (1b') is *Who is liked by nobody?*, which with respect to $D = \{r, k, m\}$ has the denotation (8).

- (8) $\{\forall x. \neg \text{like}(x, r), \forall x. \neg \text{like}(x, k), \forall x. \neg \text{like}(x, m)\}$

Some complete answers to (8) decide that $\forall x. \neg \text{like}(x, m)$ is false (i.e., that *someone likes Michael*). But those answers don't decide any propositions in (7). Thus, according to Roberts' definitions, not all complete answers to *Who is liked by nobody?* are partial answers to *Who likes Michael?*. So the former cannot be accommodated according to (6), hence (1b') is predicted to be infelicitous.

For (@2), Roberts (2012) makes the correct prediction for the falling contour, but in other cases she makes the correct prediction for the fall–rise contour. For example, both (@9b) and (@9b') are congruent to the same question (*Who does not like Michael?*). However (@9b) is felicitous while (@9b') is not.

- (@9) a. Harvey: Does Rachel like Michael?
 b. Jessica: Rachel does not like Michael.
 L*+H L- H%
 #b.' Jessica: Rachel does not like Michael.
 H* LL%

Every complete answer to *Who does not like Michael?* entails a (complete) answer to *Does Rachel like Michael?* so both (@9b,b') are predicted to be felicitous. Thus, Roberts' account makes the correct prediction for the fall–rise contour and the wrong prediction for the falling contour here. Similar objections, based on (@2) and (@9), can be made against other congruence accounts as well.

3.2 Alternative Semantics / Rooth 1992

Alternative Semantics (proposed by Rooth (1992, 2016) and developed by many others) claims that an utterance evokes a set of *alternatives*. The Roothian set of alternatives is the denotation of the congruent question. That is, the alternatives evoked by *Rachel likes Michael* are as in (10).

- (10) {Rachel likes x | $x \in D$ } for D the set of contextually available referents.

Rooth then claims that this set can relate to the prior discourse in different ways, leading to the characteristic implicatures and felicity judgements one associates with focus. The alternatives may be *congruent* in that they are the denotation of a question in the context or *contrasting* in that some alternative in the set is salient in the prior context (Rooth 1992: p85). (Rooth includes further options for focus adverbs and scalar items, but these are not relevant here.) On this account, (@2b) and (@2b') evoke the same set of alternatives that must be related to the context. Since the account does not mention contour, it has no explanation of why this succeeds for (@2b') but fails for (@2b). (The same can be said about (@9).)

An anonymous reviewer suggests a potential explanation of (@2) that doesn't require mentioning contour in the focus semantics. They point out that dialogues like (@4a,b') occur felicitously in contexts like (@11).

- (@4) a. Harvey: Who likes Michael?
 b. Jessica: Rachel_{fall} likes Michael.
 #b.' Jessica: Rachel likes Michael_{fall}.

- (@11) a. Harvey: I bet everyone who likes Michael also likes John.
 b. Harvey: Let's see: Who likes Michael?
 c. Jessica: Rachel likes Michael.
 H* LL%
 d. Jessica: But she doesn't like John.

(@11c) is felicitous because an evoked alternative — *Rachel likes John* — is salient due to (@11a), so the evoked alternatives are contrasting. The suggestion is that the pair (@4a,b') is therefore felicitous *in principle*: one *could* accommodate such information as is made explicit in (@11), but (for some reasons unrelated to focus) this is difficult in (@4). Similarly, one might claim that both responses in (@2) are in principle felicitous, but (@2b) is dispreferred to (@2b'), possibly because a fall-rise contour makes it easier to accommodate the missing information.

We don't think this argument is sound. There being *some* contexts where an utterance is felicitous does not entail felicity in *every* context. We and our informants cannot read (@2b) or (@4b') as contrasting in the way that (@11c) is and indeed cannot consider them felicitous at all. A good account of focus, we contend, must be context-sensitive enough to explain *both* the infelicity of (@4b') and (@2b) *and* the felicity of (@11c) and (@2b'). Our own account does just that (see Section 4.8)

In fact, Rooth's theory (correctly) predicts that (@4ab') is infelicitous, as the context has neither congruent nor contrasting alternatives. But the same can be said about (@2ab'), so Rooth's theory (incorrectly) predicts (@2ab') to be infelicitous. However, Rooth's theory is not necessarily mistaken. It may simply be incomplete. Possibly, it can be modified to account for (@2) by adding a relation from alternatives to context that accounts for (@2ab'). Such an amended theory still wouldn't mention contour, so it would *overgenerate* felicities (e.g., that (@2ab) is felicitous). These would need to be curtailed by an independent semantics for pitch contour. We don't know how this would be done and leave the matter open.

3.3 Givenness / Schwarzschild 1999

Another influential tradition in research on focus draws a distinction between *given* and *new* content (Selkirk 1984, Krifka 1991, Schwarzschild 1999, Büring 2006, Beaver & Clark 2009 among others). They endorse the principle GIVENNESS: "If a constituent is not F-marked, it must be given" (Schwarzschild's (1999) formulation). Our issue with these accounts is much the same as with the congruence accounts.

- (1) a. A: Who likes Michael?
 b. B: Nobody likes Michael.
 b.' B: Nobody likes Michael.

The new/given accounts we are aware of *uniformly* consider *Michael* to be given in the context after (1a) which entails that (1b') is infelicitous. We show this for Schwarzschild's (1999) account.

Schwarzschild's formal definition of givenness, when combined with GIVENNESS, entails that "an expression which is entirely new in the discourse will have to be F-marked" (p160; the formal details do not matter here). In a discourse that begins with (1a), *nobody* is a new expression in (1b) and (1b'). Thus, *nobody* must be F-marked in these utterances. This does not entail that *nobody* must be stressed, since more can be F-marked than is prosodically stressed (p155). But prosodic stress entails F-marking (*ibid.*). Hence, (1b') must F-mark (at least) *Michael* and *nobody*. But (1b) can F-mark only *nobody* (*likes Michael* is given after (1a) and hence need not be F-marked; cf. Schwarzschild 1999: sec3.2.2).

Schwarzschild uses the principle AVOIDF — "F-mark as little as possible, without violating GIVENNESS" (p156) — to predict felicity judgements. Because (1b) does not violate GIVENNESS and F-marks only one constituent, AVOIDF entails that (1b') is infelicitous for having two F-markings. So the prediction for (@2b') is also infelicity. Similar derivations of the same false prediction can be found for any other account that uses GIVENNESS+AVOIDF.

3.4 Presupposition accounts / Kratzer 1989

Many accounts of focus (e.g., Partee 1991, Geurts & van der Sandt 2004a) follow Kratzer's (1989) influential claims about (12).

- (12) a. Paula doesn't live in Paris.
 b. Paula doesn't live in Paris.

She claims that (12a) presupposes that *someone who is not Paula lives in Paris* (and this contrasts with the proffered content) whereas (12b) presupposes that *Paula lives somewhere that is not Paris* (also contrasting the proffered content). But in fact, Kratzer's claims only hold if (12) is intonated with a fall-rise contour. Observe that felicity and presupposition vary when placing (12a) in different contexts with different contours:

- (@13) a. William: Does Paula live in Paris?
 b. Edith: Paula_{f-r} doesn't live in Paris.
 ~> *someone (else) does live in Paris*
 #b.' Edith: Paula_{fall} doesn't live in Paris.

- (@14) a. William: Who does not live in Paris?
 b. Edith: Paula_{f-r} doesn't live in Paris.
 \rightsquigarrow *But is this what you wanted to know?*
 b.' Edith: Paula_{fall} doesn't live in Paris.
 $\not\rightarrow$ *someone (else) does live in Paris*

Now, on Kratzer's reading, (12a) has the truth conditions of (15a), as is indeed the case in (@13b). But (@14b') shows that (12a) can be interpreted in a way similar to (15b) too. Thus, any account that assigns to (12a) the truth-conditions of (15a) misses the data in (@14).

- (15) a. It is not Paula who lives in Paris.
 b. It is Paula who does not live in Paris.

The fact that (15a) is an anomalous response to the question (@14a) and an acceptable response to (@13a), while it is the other way round for (15b), suggests that (15a) broadly corresponds to the fall-rise contour and (15b) to the falling contour (Section 4.2 addresses how (@14b) defies this correspondence). Since (@13a) is a *prima facie* more natural context than (@14a), it seems reasonable to assume that one reads (12a) in its null context with the fall-rise contour given in (@13b); this would explain Kratzer's intuitions.

Recall our methodological point: not annotating contour confuses the modelling target. If we were to drop the annotation for contour from (@13) and (@14), we could use (@13) to criticise accounts that fail to predict the presupposition Kratzer attributes to (12a), but use (@14) against accounts that do predict it. Worse, if we (idiosyncratically) read data annotated for focus but not contour with a falling contour, we'd consider (@13b') as falsifying accounts predicting the felicity of (12a) after (@13a). But if we (idiosyncratically) read the same data with the 'most natural' contour, we may judge (12a) after (@13a) to be felicitous — drawing wholly different conclusions from the same data. Worse still, there is no guarantee that judgements about which contour is 'most natural' are intersubjective.

3.5 Büring 2003

Büring (2003) combines congruence and givenness accounts, while countenancing that contour influences how focus is interpreted in context. Specifically, this account distinguishes *focus* and *contrastive topic*; the former being indicated by a high pitched accent and the latter by a fall-rise accent.

This theory is not perfectly applicable to cases like (@2b'), since it does not model utterances with *only* a fall-rise accent (p532).

- (@2) a. Harvey: Who likes Michael?
 b. Jessica: Nobody likes Michael.
 L^*+HL- $H\%$

However, Büring (2003: p528) gives a sketch of an extension to such cases (which he embraces in Büring 2016: p75): such utterances have contrastive topics that are sets of polar questions. This means that, given a set of referents D , the contrastive topic of (@2b') would be $\{\text{does nobody like } x? \mid x \in D\}$. According to Büring's definitions (p528), a congruent (i.e., felicitous) contrastive topic contains at least two questions that are part of a strategy to answer the question under discussion. But the contrastive topic of (@2b') only contains a single question that is part of a strategy to answer (@2a): *does nobody like Michael?* So this does not work for (@2b').

3.6 Steedman 2014

Steedman (2014). provides a formally precise semantics for focus with a range of possible contours. However, it fails to predict certain implicatures. Recall (@13), here annotated using Steedman's categories for pitch.

- (@13) a. William: Does Paula live in Paris?
 b. Edith: Paula doesn't live in Paris.
 L^*+H $LH\%$
 \rightsquigarrow *someone (else) does live in Paris*

Steedman's account makes the following prediction. The L^*+H pitch expresses that Edith sees a *failure* to coordinate on what is jointly supposed and $LH\%$ attributes this failure to the hearer (i.e., William) (Steedman 2014: sec3.4.2). Paraphrased, (@13b) means *you fail to suppose that Paula does not live in Paris* (cf. Steedman's discussion of his (24)).

We would not know how to arrive at the implicature annotated in (@13b) from there. Steedman does not offer formal elaborations, but claims to derive such implicatures by *truth maintenance* (p25): the pragmatic principle that one (dynamically) constructs the interpretation of an utterance by updating the interpretation of the discourse context with the compositionally derived semantic value of the current utterance *such that the result is consistent*. By definition, only if the interpretation of an utterance is inconsistent in its context, does truth maintenance demand that certain contents in the context be modified. When there is no inconsistency, then by definition the result of update is just logical conjunction — implicatures arise via the process of restoring consistency in what would otherwise be inconsistent.

However, William's question presupposes that William doesn't know an answer, and in particular does not know whether *Paula lives in Paris*. And so it is *consistent* (with respect to classical logic, at least) with William's dialogue move that he failed

to suppose that *Paula does not live in Paris*. But if the content of (@13b) is consistent with its context, then no implicatures are generated by truth maintenance.

3.7 Fall–rise signalling nonexhaustiveness / Constant 2012

Constant (2012) offers an account that integrates focus with the fall–rise contour that comes close to predicting the implicature given in (@13). He models the fall–rise contour as expressing that the present utterance is not ‘alternative dispelling’ — meaning that the present utterance does *not* decide all propositions in the denotation of the congruent question (per the definitions in Sections 3.1). In (@13), this account would not predict the implicature we indicate, but instead that *Edith is not saying nobody (among the relevant referents) is living in Paris*.

This may be close enough — independent pragmatics might validate an inference from *not saying nobody* to *somebody* — but there is room for improvement. Our own account predicts both the strong reading we indicate in (@13) and allows for the cancellation of that to then predict Constant’s weaker reading (see Section 4.5 for details). The merits of doing so can be appreciated by considering the following example, due to Ladd (1980).

- (@16) a. Amy: Harry is the biggest liar in town.
 b. Bob: The biggest fool maybe.
 H* LL%
 b.’ Bob: The biggest fool maybe.
 L*+H L- H%

Ladd (1980) and Walker (1996a) observe that in (@16), prosody affects illocutionary force: (@16b) is interpreted to agree (*a liar and maybe also a fool*), but (@16b’) is a rejection move (*not a liar but maybe a fool*). These are *defeasible* interpretations: Bob could continue with *but not a liar*, and thus express rejection, whatever the contour. As we explained in Section 2, it is nevertheless incumbent on any model of pragmatic meaning to validate defeasible interpretations such as (@16b) vs. (@16b’) in the precise context we’ve given, and *also* validate how the interpretation of (@16b) changes when accompanied by the continuation *but not a liar*.

Contrary to its intuitive interpretation, Constant predicts that (@16b’) is *not* a rejection. If *liar* and *fool* are the only salient alternatives about Harry, then (@16b’) rejecting (@16a) means that all alternatives are dispelled — which is not the case according to Constant’s semantics for the fall–rise. The defect is that his account entails that one cannot reject using fall–rise. But one can, as in (@17).

- (@17) a. Louis: I heard you live in Cleveland now.
 b. Harvey: I live in New York City_{f-r}

Harvey’s utterance (@17b) dispells all alternatives of the form *Harvey lives in x* (given that he only lives in one place), so Constant wrongly predicts infelicity.

Constant’s account belongs to a class of analyses of the fall–rise contour that predict (in some form or other) that this contour entails the non-exhaustivity of the current utterance with respect to the current issue under discussion (e.g., Hara & van Rooij 2007, Wagner 2012, Wagner et al. 2013). These accounts differ in their formal details, but all wrongly predict that (@17b) is infelicitous, because, supposedly, (@17a) puts *Where does Harvey live?* or *Does Harvey live in Cleveland?* under discussion and (@17b) exhaustively answers either question. Our own account makes the right predictions for (@16) and (@17); see Section 5.5.

4 Our semantics for intonation

We propose a new take on the *background presupposition* semantics of focus (Jackendoff 1972, Geurts & van der Sandt 2004a), expanding the basic idea by combining it with principles of discourse coherence (see also Hobbs 1990, Reese 2007) and by making the background presupposition sensitive to pitch contour. Presuppositional accounts of focus face substantial challenges (Dryer 1996, Rooth 1999, Sæbø 2016) and we do not agree with how Geurts & van der Sandt (2004a,b) address these. Our solution to these challenges is elaborated in Section 4.6.

In this section, we develop our account piece by piece, discussing new and old data along the way. As said in Section 2, we assign semantics to *entire pitch contours* (like, e.g., Ladd 1980, Constant 2012) and remain agnostic about whether these can be decomposed. Our ambition is to demonstrate the advantages of modelling focus and pitch contour jointly, skirting the pitfalls of not doing so (see Section 3).

4.1 Background and foreground

Most accounts of focus separate *foregrounded* (focal, rhematic) content from *backgrounded* (given, thematic) content (see Féry & Ishihara 2016 for an overview). The accounts differ on how these two parts interact with the context and each other. We make the background trigger a presupposition.

(I) *Focus Semantics* (falling contour)

Focal placement separates an utterance into a *foreground* f and a *background* φ , where a variable x of the same type as f occurs freely in φ . Updating a discourse with an utterance that has a falling contour with focal constituent f proceeds as follows:

- Update with the *presupposition* φ ; that is, its free variable x must be resolved anaphorically (either bound or accommodated as $\exists x.\varphi$).

- Update the result with the *proffered* content $(\lambda x.\varphi)(f)$ (and all its presuppositions), such that the proffered content and φ are coherently connected to form a common topic (i.e., the proffered content must *elaborate* the presupposition or form a *continuation* with it (Asher & Lascarides 2003)).

We treat presuppositions as anaphora—a presupposition must be bound to an available unit in the discourse context or accommodated by existentially closing it (van der Sandt 1992, Asher & Lascarides 1998). Now, in (@4b), $f = Rachel$ and $\varphi = x \text{ likes } Michael$ (the foreground triggers a presupposition via the proper name, but by Rule (I) this updates the context *after* φ)

- (@4) a. Harvey: Who likes Michael?
 b. Jessica: Rachel_{fall} likes Michael.

After φ updates the context, the proffered content $\varphi(f)$ must attach to it with *Elaboration* (making their common topic φ) or *Continuation* (their common topic is a generalisation of their distinct but related contents). The question (@4a) presupposes *someone likes Michael*, and so the background $x \text{ likes } Michael$ binds to this, with x bound to the existential quantifier (van der Sandt 1992). The proffered content then attaches to this background with *Elaboration*; colloquially, *someone likes Michael*—*specifically, Rachel does*. The proffered content *also* attaches to (@4a) as a direct answer.

Geurts & van der Sandt (2004a) argue convincingly that standard models for how presuppositions get bound or accommodated make the right predictions for focus (but see Section 4.6). The above informal analysis of (@4ab) is an example of binding; (@18) exemplifies accommodation.

- (@18) a. Harvey: Does anybody like Michael?
 b. Jessica: Rachel_{fall} likes Michael.

Unlike (@4a), (@18a) doesn't trigger an existential presupposition, so $x \text{ likes } Michael$ is accommodated (i.e., closed with $\exists x$). The rest of the analysis then proceeds as before.

4.2 Negation and contrast

In Section 3.4 we argued that (@13b) has the truth-conditions of (15a) and (@14b') of (15b).

- (@13) a. William: Does Paula live in Paris?
 b. Edith: Paula_{f-r} doesn't live in Paris.
 \rightsquigarrow *someone (else) does live in Paris*
 #b.' Edith: Paula_{fall} doesn't live in Paris.

- (@14) a. William: Who does not live in Paris?
 b. Edith: Paula_{f-r} doesn't live in Paris.
 \rightsquigarrow *But is this what you wanted to know?*
 b.' Edith: Paula_{fall} doesn't live in Paris.
 $\not\rightsquigarrow$ *someone (else) does live in Paris*

- (15) a. It is not Paula who lives in Paris.
 b. It is Paula who doesn't live in Paris.

But in (@14b) the fall-rise contour is acceptable and its meaning is *not* that of the *it*-cleft (15a). Thus the fall-rise contour doesn't *mandatorily* result in the same presupposition as that of (15a). Furthermore, the as-opposed-to implicature associated with the fall-rise contour in (@13b) can arise in the absence of any overt negation:

- (@19) a. William: Does Paula live in Paris?
 b. Edith: Paula_{f-r} lives in Paris.
 \rightsquigarrow *someone (else) does not*

There is no negation in (@19b), so we cannot attribute its implicature to determining the relative scope of a presupposition to a linguistically-introduced negation. Rather, this as-opposed-to reading derives from *adding* a negation and determining its relative scope. We regiment this as follows: if the contour is fall-rise, then unlike Rule (I), we leave the *polarity* of the background underspecified. Moreover, to obtain the intuitive readings in (@14b) and (@19b), we specify that the proffered content is in *contrast* to the background. This semantics is expressed as follows.

(II) *Focus Semantics* (fall-rise contour, first attempt)

Updating a discourse with a fall-rise utterance with background φ and foreground f proceeds as follows:

- Update with the *presupposition* $y(\varphi)$ where y is an underspecified variable of type *polarity*; that is, $y \in \{\top, \neg\}$.
- Update with the *proffered* content $(\lambda x. \varphi)(f)$ (and all its presuppositions) such that the proffered content *contrasts with* the presupposition.

Rule (II) leaves some leeway in *what* precisely is being presupposed in a way that Rule (I) does not. This is similar to Büring's (2016) suggestion that a single fall-rise accent yields a contrastive topic of polar questions (i.e., where the polarity is left open) that has to be congruent to its context. However, our account makes no use of congruence, but instead uses general and independent mechanisms of discourse coherence to resolve in context the underspecified polarity introduced in the semantics. Hence, we understand *contrast* as in coherence theory (which differs from how Büring understands contrast). We now explain this notion in detail.

4.3 A primer on *Contrast*

Coherence relations capture the different ways in which distinct contents that are expressed in a discourse can combine to form a coherent discourse interpretation. One such way is for two contents to *contrast* one another. The relation *Contrast* corresponds to the particle *but*; that is, the contents of two clauses can contrast when they can be connected with *but*. As with any discourse relation, *Contrast* can also connect discourse segments that do not correspond to overt clauses in a discourse. For instance, a presupposition can be part of a *Contrast* (as is the case for the presuppositions triggered by fall–rise contours.)

A formal account of this goes roughly as follows. Two contents can form a contrast when their logical forms are *structurally partially isomorphic* and some isomorphic parts are *semantically dissimilar* (Asher & Lascarides 2003). The isomorphism is computed over the tree structures of the logical forms of the contrasting contents. It is partial in the following sense: one content may contain a scope bearing element (like negation) that is absent from the other. In such cases, the (partial) isomorphism ignores parts of the structure that are below such a scope bearer (Asher 1999). This is best seen by example, starting with a simple case where a complete isomorphism is possible. In *John takes Maths but Bill takes German*, the isomorphism is between $take(john, maths)$ and $take(bill, german)$, and so $take$ in the first logical form is mapped by the isomorphism to $take$ in the second, $john$ maps to $bill$ and $maths$ to $german$. Here, the requirement that there are semantically dissimilar parts in the isomorphic parts is satisfied: $maths$ is semantically dissimilar from $german$ (and arguably John is dissimilar from Bill, too). Now let's see an example of a partial isomorphism: *this is red, but it isn't scarlet*, where the logical forms for the two clauses are $red(x)$ and $\neg scarlet(x)$ (assuming the pronoun is correctly resolved). There isn't a complete isomorphism but there is a partial one: x in the first logical form maps to x in the second, and red to the semantically dissimilar $\neg scarlet$.

To judge dissimilarity, one identifies isomorphic parts from the same scale and judges their distance on the scale — farther apart means more dissimilar. This entails that some contrasts are better than others (the more dissimilar the contrasting elements, the better the contrast), so the relation *Contrast* occurs with varying *degrees of coherence* (Asher & Lascarides 2003). For example, *does/does not* and *love/hate* are (scalar) opposites, making (20a) and (20b) high quality contrasts, whereas the less opposed scalar pairs *might/definitely* and *like/love* result in acceptable, but less coherent, contrasts in (21a) and (21b).

- (20) a. Paula lives in Paris, but Jessica does not.
 b. John loves Mary, but Paula hates her.

- (21) a. Paula might live in Paris, but Jessica definitely does.
 b. John likes Mary, but Paula loves her.

There are also constraints that rule out contrasting certain contents. First, the relation *Contrast* entails the contents of the discourse segments that are so related — so out-right *contradictory* contents are not coherent, like (22a). Second, *entailment* counts as semantic similarity, which is cashed out in two ways: (i) putatively contrasting predicates cannot stand in an entailment relation, as in (22b) where \neg *bad* denotes a superset of *perfect*;⁴ (ii) the first clause of a contrast cannot entail the second, like (22c) (but the second can entail the first, like (22d)).

- (22) #a. This is scarlet, but it isn't red.
 #b. This is not bad, but it is perfect.
 #c. Katrina raced Jessica. Jessica lost, but Katrina won.
 d. Katrina and Jessica were in a race. Jessica lost, but Katrina won.

The reason for (ii) is from another aspect to contrast: two clauses may also contrast if the second defies an expectation raised by the first. For example, *Michael had a flush, but lost* is felicitous because someone having a flush is expected to win. Conversely, if the second clause *confirms* an expectation of the first, contrast is ruled out. For example, *Michael had a flush, but didn't have the lowest hand* is bad because someone *having a flush* is expected to *not have the lowest hand*. Since entailment is a form of expectation, this explains why one needs (ii). Instead of merely stipulating (ii), we would prefer a definition of contrast that unifies the ideas about expectations with those about contrasting elements. However, such a definition does not yet exist.

For the remainder of this paper, the expectation defying component of contrast is not relevant, so we do not explore it further.

4.4 Most coherent resolutions

We make use of the independently motivated principle that discourses are interpreted in a way that *maximises their coherence* (described semi-formally in Asher & Lascarides (2003), and axiomatised in Asher & Lascarides (2011)).⁵ As said, *Contrast* is a relation that varies in coherence, so preferred interpretations maximise the quality of contrasts. This interfaces with the Focus Semantics (II) as follows.

4 This constraint can be flouted for emphasis, as in *Latin is not dead, but alive and well!*. We will ignore such cases.

5 This principle has been defended extensively and via several phenomena in the wider literature on discourse coherence (e.g., Hobbs 1985, Hobbs et al. 1993). We apply this principle without rehearsing these arguments, for now relying on informal judgements about what sounds more vs. less coherent.

- (@13) a. William: Does Paula live in Paris?
 b. Edith: Paula_{f-r} doesn't live in Paris.
 \rightsquigarrow *someone (else) does live in Paris*
 #b.' Edith: Paula_{fall} doesn't live in Paris.

The utterance (@13b) presupposes $y(x \textit{ doesn't live in Paris})$. The only available referent to bind x to is *Paula*, but this is blocked by the requirement that it contrast the proffered content (if $y = \neg$, the contents contradict; if $y = \top$ there are no contrasting parts). Thus, the presupposition is accommodated as $\exists x.y(x \textit{ doesn't live in Paris})$.

Both $y = \top$ and $y = \neg$ would be permissible now, but $y = \neg$ establishes a stronger contrast to the proffered content. To see this, compare these approximate paraphrases of the two possible ways to resolve y (we eliminate the double negation in (23b)):

- (23) a.($y = \top$) *There is someone (other than Paula) who doesn't live in Paris, but Paula doesn't live in Paris.*
 b.($y = \neg$) *There is someone (other than Paula) who lives in Paris, but Paula doesn't.*

While (23a) can contrast *someone (other than Paula)* with *Paula*, the contrast in (23b) is better. The independently motivated principles that (i) stronger contrasts are more coherent and (ii) that underspecified elements in the discourse units are resolved to maximise coherence predict that the pragmatic interpretation of (@13b) can be paraphrased as (23b) rather than (23a). What is entailed by resolving an underspecification is implicated. Thus, we derive the implicature that *someone (else) lives in Paris*.

The implicature of (@19b) is computed analogously, but we will need additional machinery from Section 4.8 to explain the infelicity of (@13b'). Now recall (@14ab).

- (@14) a. William: Who does not live in Paris?
 b. Edith: Paula_{f-r} doesn't live in Paris.

The *wh*-question (@14a) presupposes *some e doesn't live in Paris*. By Rule (II), (@14b) presupposes $y(x \textit{ doesn't live in Paris})$. This presupposition can bind to the presupposition of (@14a) by making $y = \top$ and $x = e$. Such binding is impossible for $y = \neg$. Hence, as binding presuppositions is preferred, the preferred interpretation resolves y to \top .

Now, the proffered content must *contrast* with this resolution, and so binding $e = x = \textit{Paula}$ is blocked: this would result in the presupposition *Paula doesn't live in Paris* and the identical, therefore noncontrasting, proffered content *Paula doesn't live in Paris*. Thus, $e = x \neq \textit{Paula}$, resulting in the reading (23a). In other words, Edith implicates that *Paula* is an answer to William's question, but not the answer he is looking for — the desired implicature of (@14b).

Typically $y = \neg$ results in a better contrast with the proffered content and is thus often the pragmatic interpretation. In particular, in the null context $y = \neg$ is

preferred, yielding the as-opposed-to reading of fall–rise. Only in highly particular contexts, such as (@14a), does binding x to an available antecedent and resolving y to \top yield a more coherent discourse.

Finally, the negation that y resolves to can be *metalinguistic* (Horn 1989, Carston 1996; also see Beaver & Clark 2009).

- (@24) a. William: We bought po-tah-toes.
 b. Edith: We bought po-tay-toes_{f-r}.
 \rightsquigarrow not “po-tah-toes”

In (@24b) Edith is not denying the propositional content of (@24a). We can account for these cases by allowing the \neg in Rule (II) to be metalinguistic and the x in Rule (II) to resolve to prior *use* or *mention*.⁶ This also accounts for examples where fall–rise signals a speaker taking issue with the *presentation* of a proposition: While Edith gives a *positive* answer to (@25a), her intonation implicates that “in the US”, while true, mischaracterises the circumstances.

- (@25) a. William: Do you live in the US?
 b. Edith: I live in New York City_{f-r}.
 $\not\rightarrow$ not in the US.
 \rightsquigarrow not “the US”.

4.5 Uncertainty readings

Rule (II), as it stands, fails to model *uncertainty readings*.

- (@26) a. William: Did Paula eat all the cookies?
 b. Edith: Paula ate some_{f-r} of the cookies.
 \rightsquigarrow but not all of them;
 or but Edith is not sure whether it was all the cookies.

- (@27) a. William: Is Michael coming to the party?
 b. Edith: He is invited_{f-r}.
 \rightsquigarrow but he is not coming;
 or but Edith does not know whether Michael is coming.

Edith’s utterances in (@26b) and (@27b) are ambiguous: they can be interpreted as indirect *negative answers* or as indicating that Edith is *uncertain about* the answer; in the latter case, Edith is giving information that she has prosodically marked as perhaps relevant but insufficient to resolve the question.

⁶ One can formalise this as *metatalk* relations that connect the content of one utterance to the *performance* of another (Asher & Lascarides 2003: p333).

Which of the two readings is preferred seems to vary with as well as contextual knowledge of the speaker’s knowledge or intentions, but also with the intensity of the intonation and the steepness of the rises (Ward & Hirschberg 1988). Our notional categorisation of the falling and fall–rise contours underspecifies such features. Thus, to ensure that our account makes both readings *available*, we amend the Focus Semantics for fall–rise (Rule II) by adding the option that the underspecified polarity y may resolve to a modal \diamond , as in Šafářová’s (2005) semantics of final rise or Wagner’s (2012) proposal for the fall–rise contour.

(II) *Focus Semantics* (fall–rise contour, final version)

Updating a discourse with a fall–rise utterance with background φ and foreground f proceeds as follows:

- Update with the *presupposition* $y(\varphi)$ where y is an underspecified variable of type (alethic) *modality*; that is, $y \in \{\top, \diamond, \neg\}$.
- Update with the *proffered* content $(\lambda x.\varphi)(f)$ (and all its presuppositions) such that the proffered content *contrasts with* the presupposition.

As before, *maximising contrast* typically favours the \neg reading: *something isn’t P but C is P* is typically a better contrast than either *something (other than C) is P but C is P* or *something is possibly P but C is (definitely) P*. So the interpretations we outlined in Section 4.4 are replicated by this final version of the Focus Semantics. Thus our semantics favours interpreting (@26b) and (@27b) as indirect negative answers. But the reading where $y = \diamond$ is available, to interpret (@27b) (for instance) as *possibly Michael is coming, but he (definitely) is invited*, and this reading arises if the indirect answer reading is pragmatically *blocked*, for instance by the knowledge that Edith cannot know for sure whether Michael is coming.

More generally, real world knowledge can substantially affect how y resolves in context. Example (@28ab) is a case where the uncertainty reading is preferred:

- (@28) a. Amy: Does Paula like opera?
 b. Bob: She likes Wagner_{f-r}.
 \rightsquigarrow *possibly Paula likes opera.*
 b.’ Bob: She likes Wagner_{fall}.
 \rightsquigarrow *Paula likes opera.*

Axioms of rationality and cooperativity predict that responses to polar questions provide evidence for a positive answer or for a negative answer (Asher & Lascarides 2003: pp403–405); when the evidence proffered is conclusive, a particular answer is implied. Combining this expectation with the real world knowledge that *liking Wagner* is strong evidence for *liking opera (in general)* predicts that Bob has offered evidence for a *positive* answer; so the reading *Paula does not like opera (in general)*,

but she likes Wagner (i.e., $y = \neg$) is dispreferred. Further, $y = \diamond$ is preferred to $y = \top$ because $y = \diamond$ produces a better contrast: *Paula possibly likes opera (in general) but she (definitely) likes Wagner* vs. *Paula likes opera (in general) ?but she (also) likes Wagner*. Thus in this context, the fall–rise intonation conveys that Bob doesn’t quite commit to a positive answer.

This contrasts with (@28b') uttered with falling intonation, where the (same) evidence for a positive answer, provided by real world knowledge about Wagner and opera, commits Bob to a positive answer (*Paula likes opera (in general), specifically, she likes Wagner*). These differences are predicted by our semantics: the fall–rise contour demands a contrast between presupposed and proffered content, while a falling contour does not. However, such readings are a matter of degree: (@29b) (derived from Steedman 2014) is arguably ambiguous as to which answer it implicates because we cannot decide whether Bob presents *liking musicals* as positive evidence for *liking opera*, or as negative evidence.

- (@29) a. Amy: Does Paula like opera?
b. Bob: She likes musicals_{f-r}.

Implicatures towards a negative answer can arise independently of the pitch contour by purely Gricean reasoning. But our semantics makes the fall–rise contour more suggestive of a negative answer, since for a negative answer it is easier to validate a contrast than elaboration or continuation. This is in line with the empirical evidence collected by de Marneffe & Tonhauser (2019). Their results indicate that an utterance with a fall–rise contour is more likely to be read as a negative answer to a polar question than the same utterance with a falling contour. But a fall–rise contour also does not *guarantee* a negative answer interpretation. Our semantics predicts that in the cases where a fall–rise contour utterance is not a negative answer, it is preferably interpreted as expressing uncertainty, but de Marneffe & Tonhauser did not test for this.

4.6 Some challenges to a presupposition approach

Our account of focus triggers a presupposition, but such accounts have been forcefully criticised. Sæbø (2016: sec. 7.3) lists three challenges to the presupposition account of Geurts & van der Sandt (2004a), derived from the writings of Rooth (1999) and Büring (2004). The first challenge relates to Geurts & van der Sandt’s assumption that the presuppositions triggered by focus have different accessibility-constraints than other presuppositions. We do not make this assumption, so this challenge does not apply to us.

The second challenge concerns focussed quantifiers: the readings in (@30) are faulty (x ’s type appears in subscript).

- (@30) a. Nobody_{fall} likes Michael.
 ??background: x_{entity} likes Michael.
 b. Somebody_{fall} likes Michael.
 ??background: x_{entity} likes Michael.

The background in (@30a) contradicts the proffered content: x_{entity} denotes an individual in the model and so cannot be nobody. In (@30), any update with the background entails what is proffered. In both cases presupposing the background is absurd. So Geurts & van der Sandt (2004a: pp28–30) argue instead for a *polarity* focus, thereby yielding for both (@30a) and (@30b) the tautological presupposition *either nobody likes Michael or somebody does*. Sæbø (2016) and Büring (2004) argue convincingly that this strategy doesn't work, and we agree. Geurts & van der Sandt (2004a: p29) justify polarity focus for focussed quantifiers as follows.

The nonlogical part of the semantic content of words like 'somebody' and 'nobody' is so general that it is unlikely to attract the focus of a statement; 'somebody' cannot be used to mean 'some person, as opposed to some vehicle' (say).

But when we consider the fall–rise contour, we find (i) cases where *somebody* gets an 'as opposed to' reading; and (ii) cases where focus on an existential quantifier is not polar (i.e., not contrasting with 'nobody' or 'nothing'). To see (i) consider (@31), derived from a similar example by Walker (1996a).

- (@31) a. William: There is something in the garage.
 b. Edith: There is somebody_{f-r} in the garage.

The meaning of (@31) is exactly what Geurts & van der Sandt (2004a) deny: *some person, as opposed to some thing*.

Dialogue (@32) is an example of case (ii):⁷ Danny's denial move with focus on 'some' cannot be about polarity — the issue is not between *some* and *none*, but between *some* and *all*.

- (@32) a. James: [...] we're all mad, aren't we?
 b. Danny: Well, some_{f-r} of us.

We propose that in (@30–@32), the second speaker takes issue with the first speaker's choice of *specific quantifier*. So, like Constant (2012), we allow quantifiers as foregrounds, and hence as free variables in the presupposition triggered by our Focus Semantics:

⁷ This is from the British National Corpus, file HUV, lines 1468–1469. The pitch contour is constructed by us; original audio is not available.

- (@33) a. Nobody_{fall} likes Michael.
background: $x_{quantifier}z.(z \text{ likes Michael})$.
b. Somebody_{fall} likes Michael.
background: $x_{quantifier}z.(z \text{ likes Michael})$.

The presupposition in (@33) can be accommodated to form a tautology, since ‘there is a quantifier x such that $x(p)$ ’ is true of any proposition p .⁸ Thus we obtain the reading of Geurts & van der Sandt (a null presupposition) but avoid the problems with polarity focus pointed out by Büring (2004) and Sæbø (2016). In addition, this semantics is also compatible with the data (@31) and (@32) involving the fall–rise contour.

The final challenge in Sæbø’s list is due to Rooth (1999). His example (34) supposedly shows that sometimes a felicitous focus can not be replaced by an *it*-cleft triggering (on a simple background–presupposition account) the same proposition. (The *it*-cleft also triggers an exhaustivity implication, but since this is already entailed by *win*, we ignore it here.)

- (34) a. A: Did anyone win the football pool this week?
b. B: Probably not, because it’s unlikely that Mary won it,
and she’s the only person who ever wins.
b.’ B: Probably not, because it’s unlikely that it’s Mary who won it,
and she’s the only person who ever wins.

(34b’) is (supposedly) infelicitous because the presupposition of the *it*-cleft — *someone won it* — contradicts *probably not*.⁹ If focus on *Mary* triggers the same presupposition, then we should expect (34b) to be infelicitous as well (which allegedly it is not). The theoretical option to accommodate *someone won it* locally under *unlikely* is available, but there is no reason to suppose that this would be possible in (34b), but not in (34b’).

Again, we complain that this example is discussed without mentioning its potential pitch contours. For when pitch contour is taken into account, there is no counterexample. Consider (@35).

- (@35) a. A: Did anyone win the football pool this week?
??b. B: Probably not, because it’s unlikely that Mary_{fall} won it,
and she’s the only person who ever wins.
b.’ B: Probably not, because it’s unlikely that Mary_{f-r} won it,
and she’s the only person who ever wins.

⁸ But not any focus marking of a quantifier is felicitous, since there are additional constraints on focus (Section 4.8).

⁹ One reviewer reports that they find (34b’) felicitous as . . . *it’s Mary*_{f-r}, *who* . . . This may reveal something about how contours influence affect the interpretation of *it*-cleft constructions; we leave this to further work.

We consulted four native speakers of English about this example (they read A’s question and B’s full response and listened to the linked audio). All agreed that (@35b’) sounds better than (@35b): two found (@35b) outright infelicitous (as did a reviewer of this paper), while the other two were able to accommodate (@35b), albeit with some difficulty. However, those who could accommodate (@35b) were also able to accommodate (34b’).

Our semantics for the falling contour triggers the same presupposition as the *it*-cleft in (34b’). Some speakers can accommodate this presupposition — they find *both* (34b’) and (@35b) acceptable. Other speakers find *both* (34b’) and (@35b) infelicitous. Thus, our semantics for the falling contour is consistent with the data: it triggers in (@35b) the presupposition also triggered in (34b’) and speakers report identical judgements for these two utterances. Certain variations in the ability to accommodate presuppositions may have to be countenanced in the model for presupposition, but this does not affect the Focus Semantics.

In addition, our Focus Semantics for the fall–rise contour predicts that (@35b’) sounds universally acceptable. The presupposition trigger predicted by our account for (@35b’) is underspecified with regard to the following three options.

- (36) a. *someone didn’t win it*
 b. *someone possibly won it*
 c. *someone won it*

(36a) and (36b) allow for the following global accommodations; so they are preferred over any putative local accommodation of (36c).

- (37) a. *someone didn’t win, but probably no-one won . . .*
 b. *someone possibly won, but probably no-one won . . .*

(37b) is the most coherent contrast here. Indeed, (37b) appears to be the correct interpretation of (@35b’). The speaker communicates that possibly Mary won, which is unlikely, but probably no-one won. This explains why all our informants find (@35b’) felicitous.

So our proposal is entirely consistent with the judgements on Rooth’s alleged counterexample (34), once contour is taken into account.

4.7 Scalar maxima

We now briefly outline how our model accounts for some interesting fall–rise data from Constant (2012). He observes that scalar maxima like *perfect* in (@38) cannot be focussed with the fall–rise contour.

- (@38) # The food was perfect_{f-r}.

Our semantics predicts this. The three options for resolving the presupposition triggered by the fall–rise contour in (@38) can be paraphrased as follows. (Since

we interpret (@38) in a null context, the variable x (of type *property*) must be accommodated (i.e., existentially closed.)

- (39) a. ($y = \neg$) There is some $x_{\langle e,t \rangle}$ that the food did not have, but it was perfect.
 b. ($y = \diamond$) There is some $x_{\langle e,t \rangle}$ that the food possibly had, but it was perfect.
 c. ($y = \top$) There is some $x_{\langle e,t \rangle}$ that the food had, but it was perfect.

All three options sound odd, since one cannot think of an appropriate $x_{\langle e,t \rangle}$. If $x_{\langle e,t \rangle}$ is a positive property, then (39a) is contradictory, since if the food did not have this property, it wasn't perfect. Similarly, if $x_{\langle e,t \rangle}$ is positive, then (39b) and (39c) fail to be contrasting (since positive properties are entailed by *perfect*). If, on the other hand, $x_{\langle e,t \rangle}$ is a negative property, then (39a) fails to be a contrast (because perfection entails no negative properties), and (39b) and (39c) are contradictory. If $x_{\langle e,t \rangle}$ is a neutral property, it fails to establish a contrast in all three cases (it would invite you to draw a comparison between, say, *red* and perfection, with no information on why they contrast each other). Thus, the presupposition cannot be accommodated at all and, thus, (@38) sounds infelicitous.

By the same reasoning, we can account for Constant's observation that fall-rise intonation can disambiguate quantifier scope (also see Wagner 2012).

(@40) $\underline{\text{All}}_{f-r}$ my friends didn't come.

Were it not for the fall-rise contour, the surface form of (@40) would be ambiguous between a linear scope reading (*all didn't*) and a scope inversion (*not all*). The contour disambiguates to the scope inversion (Constant 2012: p408).

On our account, the linear reading is unavailable because it would require to resolve the background $y(x_{\text{quantifier of my friends didn't come}}$) in a way that establishes a contrast with what is proffered (*all my friends didn't come*). Any resolution of y and x either entails that (*maybe*) *some of my friends came* or that (*maybe*) *none of my friends came*. The latter is inconsistent with the proffer (*all my friends didn't come*) and the former fails to establish a contrast. Either way, we cannot validate the requisite *Contrast* relation.

Thus, the scope inversion reading must be selected. Its background presupposition $y(\neg x_{\text{quantifier of my friends came}})$ can be resolved as $y = \neg$, $x = \exists$ to yield an interpretation paraphrasable as *some of my friends came, but not all of my friends came*.

4.8 Given and interest

Our account does not yet explain why in many cases like (@4), prosodic focus appears to follow the structure of an antecedent question.

- (@4) a. Harvey: Who likes Michael?
 b. Jessica: Rachel_{fall} likes Michael.
 #b.' Jessica: Rachel likes Michael_{fall}.

Intuitively, Bolinger's (1972) slogan *focus follows interest* offers an explanation. One way to be interesting is to be important (Bolinger 1985). It seems reasonable to assume that *Rachel* is the part of Jessica's response that is important, since Harvey would understand her even if she mumbled the rest. But *Michael* does not seem to be important or otherwise interesting. Thus, Jessica should focus *Rachel*.

We cannot define what it means for a speaker to find something interesting. But we can identify and formalise conditions that are necessary for something to be *presented as* interesting. Bolinger (1972) provides one such condition: what is obvious cannot be interesting. We find this suggestion plausible and, based on the good predictions our formal theory will make, fruitful. We formally approximate Bolinger's suggestion by saying that to be (presentable as) interesting it is necessary (but not sufficient) to be not (also presented as) given. Because focus marks the speaker's interest, we take the backgrounded content of the current utterance to contribute to what is considered given. Further, given information is restricted to the (salient) part of the discourse context which is being addressed by the current utterance, and thus Rule (IV) draws on the definition (III):

(III) *Relevant Segment*

The *relevant segment* for a foreground–background pair $\langle f, \varphi \rangle$ is the segment of the prior discourse that the *proffered* content most coherently relates to.¹⁰

(IV) *Givenness*

The given information is everything entailed by the coherent update of the relevant segment of the prior discourse with the (presupposed) background of the current utterance.

This definition of Givenness (IV) departs substantially from prior accounts (cf. Section 3.3). We do not say that what *can* be background is *constrained* by what is given, but instead *use* the background to *compute* what is given.

Now, if the most coherent way to update the discourse context with the background content results in a meaning that entails the proffered content, then by Rule (IV) what is proffered is *given*. Rule (V) makes this anomalous:

¹⁰ In all our examples, the relevant segment is the last part of the prior discourse, but in general it can be any prior discourse unit (Hobbs et al. 1993).

(V) *Necessary Condition for Interest*

A foreground–background pair $\langle f, \varphi \rangle$ is *not interesting* if the proffered content $\varphi(f)$ is given.

We will sometimes say *f is not interesting* for $\langle f, \varphi \rangle$ is *not interesting* if it is clear from context what φ is.

Now, recall that according to the Focus Semantics for the falling contour (I), a foreground–background pair $\langle f, \varphi \rangle$ triggers a presupposition of φ (with a free variable x for the focal constituent f) and proffers the content $\varphi(f)$. And according to the Focus Semantics for the fall–rise contour (II), $\langle f, \varphi \rangle$ triggers a presupposition of $y(\varphi)$ (with a free variable x for f and y underspecifying modality) and also proffers the content $\varphi(f)$. Hence, according to (V) a falling contour utterance is *not interesting* if it is given that x resolves to f ; for the fall–rise contour additionally that y resolves to \top .¹¹

Only the *relevant* segment is given because attention in discourse is limited and dynamic (Walker 1996b), and so a fall–rise contour can be used to repeat prior information.

- (@41) a. Harvey: Who likes Michael?
 b. Jessica: Rachel_{fall} likes Michael.
 ⋮
 y. Harvey: Who likes Michael again?
 z. Jessica: Rachel_{f-r} likes Michael.

As demonstrated in Section 4.4, the Focus Semantics predicts an implicature for (@41z) that can roughly be paraphrased *but Rachel is not who I take your question to be about*. Depending on the circumstances, this could express confusion or exasperation.

¹¹ Elizabeth Coppock points out to us that, *prima facie*, discourses like (a) *Larry brought flowers*. (b) *And not only Larry_{fall} brought flowers!* (c) *Sean_{fall} did too!* seem to contradict Rule (V). If the focus presupposition of (b) is *not only x brought flowers*, then in the context of (a), x is resolved to *Larry* — so the proffered content of (b) is given by Rule IV. (Note that (b) arguably sounds more natural with a fall–rise contour. But with a fall–rise contour, the example is harmless, since in the context of (a), the underspecified modality y is not resolved to \top .)

So far, we have computed what is background simply by abstracting a variable in the focal constituent, but this is a simplification. Properly, this turns on a compositional semantics that in particular respects focus-associating operators like *only*. We do not elaborate an account of *only* here, but are sympathetic to Geurts & van der Sandt’s (2004a) suggestion that *only*-clauses trigger the focus presuppositions of the content below *only* and add to the proffer that the focus presupposition is instantiated by at most one referent (thus, for *not only* it must be instantiated by more than one). This is compatible with Rule (V) and explains Coppock’s example: (b) presupposes *x brought flowers*, which in the context of (a) is resolved to *Larry brought flowers*. Since this does not entail that a second person brought flowers, the proffered content of (b) is not entailed. We leave to further work whether our account can adopt this semantics for *only* without modification.

Now note that (III), (IV) and (V) predict that (@4b') sounds odd. In the context of (@4a), the most coherent way to resolve the presupposition *Rachel likes x* triggered by (@4b') is to bind x to *Michael* to form an answer to (@4a). Thus the proffered content is given, violating Rule (V).¹² So we capture the basic data motivating question-answer congruence. But other than the accounts we criticised in Section 3, Rule (V) allows that (@2b') is acceptable.

- (@2) a. Harvey: Who likes Michael?
 b.' Jessica: Nobody likes Michael_{f-r}.

In (@2), the presupposition triggered by *nobody likes Michael*_{f-r} is $y(\textit{nobody likes } x)$. This can be resolved to $y = \neg$, $x = \textit{Michael}$ to bind to the presupposition triggered by the *wh*-element of (@2a): *someone likes Michael*. Since such binding is preferred over any other option to deal with a presupposition, $y = \neg$ is given, and hence *nobody likes Michael*_{f-r} passes the interest test (V).¹³ (We formalise this reasoning in Section 5.5.)

Just like presuppositions generally, the coherent interpretation of the background presupposition needn't be unique; in such cases, the given information is the information all equally coherent interpretations agree on. Also, it is possible that once the proffered content is considered, the initial interpretation of the given information must be revised. Dialogue (@13) is a case where a highly salient coherent interpretation of given information gets overridden by proffered content:

- (@13) a. William: Does Paula live in Paris?
 b. Edith: Paula_{f-r} doesn't live in Paris.

The presupposed content of (@13b) is $y(x \textit{ does not live in Paris})$. The most coherent update of (@13a) with the (underspecified) presupposition *on its own* yields $x = \textit{Paula}$ but doesn't resolve y uniquely: $y = \top$ and $y = \neg$ both supply (full) answers to the question, while $y = \diamond$ provides a (dispreferred) partial answer. Either way, the presupposition is coherent (though ambiguous), with $x = \textit{Paula}$. However, when updating this with the proffered information, the Focus Semantics for fall-rise demands a *Contrast*. This cannot be reconciled with $x = \textit{Paula}$, so this initial defeasible inference is overridden: x gets existentially bound and y resolves to \neg (resulting in *someone else lives in Paris*). This (still) coherently attaches to the question (as a commentary rather than an answer) but it also contrasts the proffered content (which in turn attaches to (@13a) as an answer), as demanded by the Focus Semantics.

12 The same would be true if (@4a) had the fall-rise contour, since $y(\textit{Rachel likes } x)$ is preferably resolved to $y = \top$ and $x = \textit{Michael}$, since this forms an answer to (@4a), whereas other options would result in less coherent partial answers.

13 In computing what is given, the proffered content is not considered; so it doesn't matter that the resolution *someone likes Michael* cannot contrast the proffered content in (@2b').

Since the contours trigger different presuppositions, the focal placement alone doesn't determine whether the foreground can be interesting. Rather, interest is governed by the focal placement and contour in combination. Unlike (@13b), (@13b') is predicted to be infelicitous.

- (@13) a. William: Does Paula live in Paris?
 #b.' Edith: Paula_{fall} doesn't live in Paris.

The presupposition of (@13b') is *x does not live in Paris*; in the context of (@13a), the most coherent resolution is *x=Paula*. So by Rule (V), (@13b') is not interesting and thus incoherent. In contrast, (@13b) also initially resolves *x=Paula*, but does not resolve the polarity of the answer. *Mutatis mutandis* this also explains why (@42b) is coherent (where the focus is on the polarity).

- (@42) a. William: Does Paula live in Paris?
 b. Edith: Paula does not_{fall} live in Paris.

Also, Rule (V) generalises Rooth's (1992) notion of a contrasting alternative. For Rooth, a focus is contrasting if there is a contextually salient alternative (some instance of the background φ) that is not the alternative selected in the proffered content. In such cases, the necessary condition for interest is satisfied: if an instance of φ (that is not $\varphi(f)$) is contextually available, then it is not given that $\varphi(f)$ (unless other information would lead one to bind *x* to *f* over the contrasting alternative).

4.9 Beyond interest

Our necessary condition is at best a small component of what *interest* is. Fundamentally, interest is paralinguistic and subjective. Bolinger (1985) suggests the Boston Strangler might utter (@43a) while a sane individual prefers (@43b).

- (@43) a. I'm looking for a girl_{fall} to strangle.
 b. He's looking for a girl to strangle_{fall}.

For the Boston Strangler, it's clear that he will strangle someone, so who exactly is his matter of interest. Without what Bolinger rightly calls mind-reading, we cannot account for such variation.

Also, Bresnan (1971) and Bolinger (1972) note that it is difficult to focus constituents with apparent low semantic content; for example, one cannot focus *someone* in (@44b), but the more informative *policeman* in (@44b'').

- (@44) a. William: What did John do?
 #b. Edith: John killed some_{fall}one.
 b.' Edith: John killed_{fall} someone.
 b.'' Edith: John killed a policeman_{fall}.
 #b. ''' Edith: John killed_{fall} a policeman.

Rule (V) can explain this: In (@44b) the presupposed background is *John killed x* for a quantifier *x*. Since *someone* requires a domain of restriction of *people* (as opposed to, say, animals), the quantifier *x* is of type *quantifier(people)*.¹⁴ Since *John killed nobody* is a dispreferred response to (@44a), it is given that *x* resolves to a nonempty quantifier (i.e., one that entails *someone*). *John killed someone* is therefore given and so by Rule (V), (@44b) is anomalous and (@44b') is the correct way to mark interest. But *John killed a policeman* is not given, so (@44b'') is fine. However, explaining why (@44b''') is infelicitous requires principles going beyond interest.

Clearly, there are linguistic constraints other than Rule (V) on focus. Notably, if there are multiple parts of a clause that could in principle be focal — according to some mechanism like interest, givenness, or congruence — there is a tendency to place focus on the right-most of them (Selkirk 1984). This entails that (@44b'') is preferred over (@44b'''). Thus, unlike Bolinger, we do not say that ‘policeman’ in (@44b'') is *the* or *the most* interesting constituent, since *prima facie* ‘killed’ has an equally good claim to this distinction. Rather, it is *one* of the constituents that *can* be marked as interesting, and other principles (possibly not related to interest at all) govern the placement among these.¹⁵

4.10 Coherence

Coherence relations play another useful role: if one cannot find a relevant segment, or the background does not cohere with it, then the dialogue is incoherent.

(@45) *Context: Jessica and Katrina are both job hunting.*

- a. Harvey: Did you get a job?
- b. Jessica: Katrina_{f-r} got a job.
- #b.' Jessica: Katrina got a job_{f-r} .
- #b.'' Jessica: Katrina got a job_{fall} .

No coherent interpretation can be computed of the presupposition *Katrina got an x* that's triggered by (@45b',b'') in the context (@45a), as there is no value for *x* that makes it a coherent response (in particular, *x=job* doesn't work, because *Katrina got a job* and *Jessica got a job* are logically independent, and so this resolution isn't an indirect answer to (@45a)). But if a presupposition cannot be coherently related to

14 There seems to be a difference between focussing *someone* and *someone*. We would locate this difference within the type of *x*: the former having type *quantifier(people)* and the latter having type *quantifier*, but cannot go into detail here.

15 Based on similar observations, Wagner (2006) suggests that givenness might be gradient. Analogously, interest might be gradient. We are sympathetic to this, but leave the matter open.

its context, the dialogue is infelicitous (Asher & Lascarides 1998). A similar defect explains the infelicity of (@46b').¹⁶

- (@46) a. Julian: Is Nicholas coming to my talk?
 b. Alex: Ernie_{f-r} is coming to your talk.
 #b.' Alex: Ernie_{f-r} isn't coming to your talk.

In (@46b) the proffered content attaches to the question (@46a) as a *Plan-Correction* (Asher & Lascarides 2003: p320), meaning that Alex cooperatively amending Julian's conversational strategy according to what she infers about their *goals*. It is accommodatable in (@46) that Julian's goal is, say, to know how challenging the question session will be; by answering *Ernie is coming* Alex is giving relevant, possibly *better*, information towards this goal. But it is much harder to accommodate that *Ernie is not coming* attaches as a Plan-Correction to (@46a) — or indeed with any other relation. Hence, (@46b') is infelicitous because, unrelated to intonation, it is a *non sequitur*.

4.11 Summary: Intonated discourse update

In sum, our proposed analysis of intonation is as follows:

- i. The grammar produces a foreground–background pair $\langle f, \varphi \rangle$, where φ features a free variable x of the same semantic type as f , where f is the (unique) focal constituent (recall that our analysis is currently restricted to pitch contours with only one such constituent).
- ii. From f and φ , compute the *proffered* and (underspecified) *presupposed* content according to the Focus Semantics (I, II).
- iii. Compute the *relevant segment* according to general principles for computing how the proffered content updates the discourse (Rule III).
- iv. Compute whether and how the *presupposed* content attaches to the relevant segment, again via general principles for presupposition resolution. The result is the *given* information (Rule IV).
- v. Check whether the given information and proffered content satisfy the necessary condition on interest (Rule V).
- vi. If all is well, update the discourse with both the presupposed and proffered content — again via reasoning about discourse coherence and respecting the coherence relations entailed by the Focus Semantics (Rules I and II).

¹⁶ We thank the audience of the 2017 Rutgers Semantics & Pragmatics Colloquium, in particular Sam Carter, for pressing us on (@46).

Steps (i–ii) serve to define the meaning representation of the utterance given its intonational form (though we forego *deriving* these within the grammar): for present purposes, φ and f are simply computed by λ -abstracting the focal constituent (but this is a simplification). Steps (iii–v) jointly are a *check* on the felicity conditions of those proposed (underspecified) meaning representations. This check makes use of notions related to coherent discourse update, but does not amount to an *actual* update to the current context; this is executed in Step (vi) only if the felicity condition in (v) is satisfied.

5 Formalised account

We now formally regiment the above analysis within Segmented Discourse Representation Theory (SDRT, Asher & Lascarides (2003)). We use SDRT because it has already been used extensively to model the interaction between discourse coherence and presuppositions and provides an axiomatisation for interpreting discourse so that coherence is *maximised* (Asher & Lascarides 2011). We start by giving a brief description of SDRT, to then use it to formalise principles (I–V).

5.1 Discourse structure

SDRT models discourse structure by connecting the contents of utterances with *coherence relations* like *Narration*, *Elaboration* and *Correction*. Logical forms in SDRT consist of a set of *labels* π_1, π_2, \dots that each represent a unit of discourse, and an *assignment function* \mathcal{F} that associates each label π with a formula φ , representing the unit’s interpretation. We write $\mathcal{F}(\pi) = \varphi$ or $\pi : \varphi$ to express this mapping. The content φ can consist of coherence relations among labels, so \mathcal{F} induces a partial order: π_1 *outscores* π_2 if π_2 occurs in $\mathcal{F}(\pi_1)$. A *coherent* logical form — known as an SDRS — has a unique root under this partial order.

Cue phrases (e.g., *then*, *therefore*, *but*) can entail coherence relations, but frequently they are inferred via commonsense reasoning with linguistic and nonlinguistic information. But ambiguity can persist: simplifying somewhat (ignoring presuppositions, tense and so on), (47) could express that because the meeting was cancelled, Nicholas stayed at home (48) or that the meeting was cancelled because Nicholas stayed at home (49).

(47) The meeting was cancelled. Nicholas stayed at home.

$$\begin{aligned}\pi_1 &: (\mathbf{1}x)(meeting(x) \wedge cancel(e_1, x)) \\ \pi_2 &: (\mathbf{1}y)(stay(e_2, n) \wedge home(y) \wedge at(e_2, y))\end{aligned}$$

(48) $\pi_0 : Result(\pi_1, \pi_2)$

$$(49) \quad \pi_0 : \textit{Explanation}(\pi_1, \pi_2)$$

SDRSs are assigned a dynamic semantics, where one starts to unpack its content from its root label. The semantics of a coherence relation $R(\pi_1, \pi_2)$ is defined in terms of its arguments' contents (i.e., $\mathcal{F}(\pi_1)$ and $\mathcal{F}(\pi_2)$). For instance, the general rubric for *veridical* relations like *Explanation* and *Narration* is given in (50), where C and C' are the contexts of interpretation (typically, sets of world-assignment pairs), \wedge corresponds to *dynamic conjunction* (i.e., $\llbracket \varphi \wedge \psi \rrbracket = \llbracket \varphi \rrbracket \circ \llbracket \psi \rrbracket$), and $\varphi_{R(\pi_1, \pi_2)}$ is content that is specific to the coherence relation R and is specified in terms of $\mathcal{F}(\pi_1)$ and $\mathcal{F}(\pi_2)$:

$$(50) \quad C \llbracket R(\pi_1, \pi_2) \rrbracket C' \textit{ iff } C \llbracket \mathcal{F}(\pi_1) \wedge \mathcal{F}(\pi_2) \wedge \varphi_{R(\pi_1, \pi_2)} \rrbracket C'.$$

For example, *Background* is a veridical relation, where $\varphi_{\textit{Background}(\pi_1, \pi_2)}$ is equivalent to the condition that the event e_1 described by π_1 spatio-temporally overlaps the event e_2 described by π_2 . Other relevant relations for present purposes are *Elaboration* (which entails that the second-part is a more specific description of the event described in the first part), *Continuation* (which entails that the first-part and second-part share a topic), and *Contrast* (which we described in Section 4.3). See the appendix of Asher & Lascarides (2003) for a glossary of coherence relations and their meanings.

5.2 Construction of logical form and maximising discourse coherence

These dynamic semantics capture how to *evaluate* an interpretation of a discourse. But *constructing* which logical form is the intended interpretation is a different task, carried out in a separate logic (Asher & Lascarides (2003) provide detailed motivation for this separation).

This *glue logic* consists of *default axioms* that model how commonsense reasoning with both linguistic and nonlinguistic information validates defeasible inferences about which available unit(s) in the context the current unit connects to, which coherence relations connect them, and how other underspecified elements are resolved (e.g., anaphora and the relative semantic scope of presuppositions). The glue logic reasons over *underspecified logical forms* (ULFs), which in turn express *partial descriptions* of fully specific logical forms (SDRSs).

ULFs are computed from surface forms. Pronouns introduce a condition $x = ?$, which means that x must be co-referent with an available antecedent, but exactly which antecedent isn't known. More generally, a ULF uses a variable $?$ of an appropriate sort whenever a specific value of some construction isn't known. For instance, $\lambda : ?(\alpha, \beta)$ means that β is connected to α , forming part of the discourse segment λ , but the coherence relation isn't known.

The glue logic contains axioms expressing (defeasible) pragmatic preferences and uses them to enrich the information in an ULF. These axioms use a defeasible conditional: $\varphi > \psi$ means *If φ , normally ψ* . For example, IQAP is a glue-logic default axiom which stipulates that normally, a response to a question is an indirect answer (IQAP stands for *Indirect Question Answer Pair* and its semantics entails that direct answers are also indirect answers):

$$(IQAP) (\lambda : ?(\alpha, \beta) \wedge \text{interrogative}(\alpha) \wedge \text{spk}(\alpha) \neq \text{spk}(\beta)) > \lambda : IQAP(\alpha, \beta)$$

In words, if β is connected to α but we don't (yet) know with what coherence relation, α is an interrogative and α and β are said by different people, then normally, β answers α .

In addition to such axioms, SDRT formalises that one (always) interprets discourses in a way that *Maximises Discourse Coherence*. As we have mentioned throughout, discourse coherence is not a yes/no matter; it can vary in quality. SDRT's principle MDC defines factors that affect that quality. Roughly put, they are as follows (formal details are in Asher & Lascarides (2003: p233)):

Maximise Discourse Coherence (MDC).

Given competing interpretations of a discourse, select for \leq^c -maximality. For SDRS K', K , it is the case that $K' \leq^c K$ iff all of the following hold:

- i. If K' is consistent, then so is K .
- ii. *Prefer rich structure*: K has at least as many coherence relations as K' .
- iii. *Prefer flat structure*: K has at most as many labels as K' unless K' has a *semantic clash* and K does not.¹⁷
- iv. *Prefer better relations*: Each rhetorical connection in K is at least as coherent as those in K' . Recall that the *Contrast* connection varies in quality (see Section 4.3); similarly, a *Continuation* is better the more specific the common topic of the contents it relates are.

¹⁷ A *semantic clash* occurs if a segment α of an SDRS appears both veridically and nonveridically (i.e., one segment outscoping α entails α 's content, but another segment outscoping α does not). The SDRS (51c) is an example of a semantic clash, because the content of π_2 occurs veridically in *Parallel* but nonveridically in the *If(-then)* relation; (51b) is thus more coherent (despite having more labels).

- (51) a. π_1 : If a shepherd goes to the mountains,
 π_2 : he normally brings his dog.
 π_3 : He brings a good walking stick too.
 b. π_0 : *If*(π_1, π)
 π : *Parallel*(π_2, π_3)
 c. π_0 : *If*(π_1, π_2) \wedge *Parallel*(π_2, π_3)

- v. *Prefer resolution*: K resolves (as computed by dynamic update through the coherence relations) at least as many underspecifications as K' does.

There isn't always a *unique* maximally coherent interpretation. For instance, MDC does not distinguish between the two alternative SDRSs (48) and (49) of (47).

We are ready to define discourse update. The definition uses the nonmonotonic proof theory \vdash_g that validates intuitively compelling patterns of the conditional $>$, such as Defeasible Modus Ponens ($\varphi, \varphi > \psi \vdash_g \psi$ but $\varphi, \neg\psi, \varphi > \psi \not\vdash_g \psi$) (Asher & Lascarides 2003: ch5).

Definition 1 (Update). Let Γ be a ULF for the discourse context and $\pi : \mathcal{K}$ be a ULF representing new information. Then $update(\Gamma, \pi : \mathcal{K})$ is the set of all (and only) those SDRSs that satisfy the glue logic consequences of attaching π to some available segment α in Γ . More formally: $K \in update(\Gamma, \pi : \mathcal{K})$ iff K is an SDRS and there is an available segment α in Γ such that for all glue-formulae φ

$$\text{If } \Gamma, \pi : \mathcal{K}, \lambda : ?(\alpha, \pi) \vdash_g \varphi, \text{ then } K \models \varphi,$$

where $K \models \varphi$ means that the fully specific SDRS K includes the (underspecified, partial) information in φ .

To obtain the pragmatically preferred interpretation, one then selects the \leq^c -maximal SDRSs from $update(\Gamma, \mathcal{K})$.¹⁸ If $update(\Gamma, \pi : \mathcal{K}) = \emptyset$, the discourse is incoherent.

5.3 Presuppositions

Following van der Sandt (1992), SDRT assumes that the linguistic grammar derives a ULF in which proffered content is separated from presupposed content and their relative semantic scope is underspecified. For instance, in both theories the presupposition trigger *regret* yields the logical form for (52a) given in (52b):

- (52) a. A man didn't regret smoking.
 b. proffered: $\pi_1 : \exists x(man(x) \wedge \neg regret(e, x, \wedge smoke(e', x))$
 presupposed: $\pi_2 : smoke(e', x)$

How the presupposed and proffered contents coherently relate to their context is resolved in the glue logic. First update with the presuppositions. Glue logic axioms (defeasibly) entail that presuppositions *bind* to a prior unit; when binding isn't coherent, another axiom entails that the presupposition is *accommodated* by attaching it to the outscoping-maximal unit where such attachment is coherent

¹⁸ It is possible to axiomatise MDC within a *dynamic* glue logic and thereby including it in the definition of *update* (Asher & Lascarides 2011), but the added formalism would distract here.

(i.e., local accommodation is dispreferred). These defaults don't apply to proffered content.

For example, simplifying somewhat (only the ULF labels that contribute to anaphoric or semantic scope ambiguities are shown), the grammar generates for the two sentences in (53a) the ULFs (53b) and (53c) (where presupposed content is marked with ∂).

- (53) a. A man had a health scare. But he didn't regret smoking.
 b. $\pi_1 : \exists x \exists y (man(x) \wedge health-scare(y) \wedge have(e_1, y, x))$
 c. $\pi_2 : Contrast(\pi, \pi_3), \pi = ?$
 $\pi_3 : \neg regret(e_2, z, \wedge smoke(e_3, z))$
 $\pi_4 : \partial smoke(e_3, z), z = ?$
 d. $\pi_2 : Background(\pi_1, \pi_4) \wedge Contrast(\pi_1, \pi_3)$

Sentence-initial *but* introduces a *Contrast* relation whose first argument is anaphoric ($\pi = ?$). Given number and gender constraints (omitted here), the only candidate for resolving $z = ?$ is $z = x$. So x must be made available, which means that the presupposition π_4 must connect to π_1 . Both these discourse units describe *states*, and so the glue logic axioms validate a (defeasible) inference that they connect as *Background*(π_1, π_4). MDC then predicts that $\pi = ?$ resolves to $\pi = \pi_1$: this forms a better quality *Contrast* (and a flatter structure) than the alternative (i.e., *Contrast*(π, π_3) where $\pi : Background(\pi_1, \pi_4)$). So the final SDRS is as shown in (53d): this entails that smoking occurred, even though *smoking* was syntactically outscoped by *not*.

5.4 Intonated discourse update

We now formalise our model from Section 4. Definition 2 formalises the Focus Semantics (I, II).

Definition 2 (Focus Semantics). Let $\langle f, \varphi \rangle$ be the foreground–background pair of the current utterance. The discourse update associated with $\langle f, \varphi \rangle$ is an update with $\pi^b : \partial \mathcal{K}_{\pi^b}$, $\pi^f : \mathcal{K}_{\pi^f}$ and $\pi : R^{ctr}(\pi^b, \pi^f)$, where:

- \mathcal{K}_{π^f} is the ULF corresponding to $\varphi(f)$.
- If the contour is falling, then $R^{ctr} = ?_{topic}$ (which can be resolved to *Continuation* or *Elaboration*) and \mathcal{K}_{π^b} is the ULF corresponding to φ , where x occurs free in φ and of the same type as f , and the semantic index (an eventuality term) of φ is syntactically distinct in \mathcal{K}_{π^b} and \mathcal{K}_{π^f} (although they can denote the same eventuality).

- If the contour is fall–rise, then $R^{cntr} = Contrast$ and \mathcal{K}_{π^b} is the ULF corresponding to $?_{mod}(\varphi)$, where x occurs free in φ and of the same type as f , and $?_{mod}$ underspecifies *modality* — that is, it can resolve to \top , \diamond or \neg . Again, the semantic index of φ is syntactically distinct in \mathcal{K}_{π^b} and \mathcal{K}_{π^f} .

Note how focal placement *in combination* with contour produces distinct contents with which to update the discourse context. Specifically, the presupposition that’s triggered from prosodic form — i.e., the underspecified logical form \mathcal{K}_{π^b} — depends on both focal placement and contour: the focal placement determines φ ; the contour determines whether its polarity is underspecified and how it relates to the proffered content.

To compute what is given (Rule IV), compute the maximally coherent interpretation of the relevant segment, updated with the background content.

Definition 3 (Givenness). The relevant segment α^r is the segment of the prior discourse where \mathcal{K}_{π^f} most coherently attaches. If there are multiple such segments let α^r be the minimal segment (w.r.t. the order induced by outscoping) that outscopes all most coherent attachment points. Write $\Gamma \upharpoonright \alpha^r$ for the subset of Γ containing α^r and all segments outscoped by α^r . Then, ϕ is *given* if and only if for all maximally coherent $K \in update(\Gamma \upharpoonright \alpha^r, \pi^b : \mathcal{K}_{\pi^b})$, $K \models \phi$.

Definition 4 (Intonated Discourse Update). Let Γ be the prior context and $\langle f, \varphi \rangle$ be the foreground–background pair of the current utterance.

- i. Compute the ULFs \mathcal{K}_{π^f} and \mathcal{K}_{π^b} as in Definition 2.
- ii. If $update(\Gamma, \pi^f : \mathcal{K}_{\pi^f}) = \emptyset$, break.
- iii. If $update(\Gamma \upharpoonright \alpha^r, \pi^b : \mathcal{K}_{\pi^b}) = \emptyset$, break.
- iv. If \mathcal{K}_{π^f} is given, break.
- v. Do discourse update on Γ with $\pi^b : \partial K_{\pi^b}$, $\pi^f : K_{\pi^f}$ and $R^{cntr}(\pi^b, \pi^f)$.

One can express (ii–iv) as monotonic axioms in the glue logic to reduce Definition 4 to standard SDRT-update (Definition 1). For simplicity, we won’t do this here.

5.5 Some formal analyses

As mentioned in Section 3, intonation can affect the illocutionary force of an utterance: (@16b) is interpreted as agreement, but (@16b’) is a rejection move.

- (@16) a. Amy: Harry is the biggest liar in town.
 b. Bob: The biggest fool_{fall} maybe.
 b.’ Bob: The biggest fool_{f-r} maybe.

Simplifying somewhat (treating “maybe” as equivalent to \diamond , omitting tense, treating subject pro-drop like a pronoun and over-simplifying the semantics of “biggest”) the ULF of (@16a) is (54a) and the Focus Semantics (Definition 2) yields the ULFs (54b) and (54b’) for (@16b) and (@16b’):

- (54) a. $\alpha^r : liar(e_{\alpha^r}, h) \wedge biggest(e_{\alpha^r})$ *Harry is the biggest liar.*
 b. $\pi : ?_{topic}(\pi^b, \pi^f)$.
 $\pi^b : \partial \diamond (P(e'_\beta, x) \wedge biggest(e'_\beta)) \wedge P = ?$ *he is maybe the biggest P.*
 $\pi^f : \diamond (fool(e_\beta, x) \wedge biggest(e_\beta))$ *he is maybe the biggest fool.*
 b'. $\pi' : Contrast(\pi'^b, \pi'^f)$.
 $\pi'^b : \partial ?_{mod} \diamond (P(e'_\beta, x) \wedge biggest(e'_\beta)) \wedge P = ?$ *y(he is maybe the biggest P).*
 $\pi'^f : \diamond (fool(e_\beta, x) \wedge biggest(e_\beta))$ *he is maybe the biggest fool.*

In (54ab), one must first ensure that π^b can coherently update the context α^r to yield given information that doesn’t entail π^f . MDC prefers binding P to the available antecedent *liar* (and x to the available antecedent Harry) rather than resolving P via existential quantification. Thus the proffered content isn’t given and discourse update can proceed. First, resolving P to *liar* validates the relation *Accept* between α^r and π^b . Then, the underspecified relation $?_{topic}$ resolves to *Continuation*: the glue logic axioms don’t validate inferring *Elaboration* (conventionally, “liar” is not more specific than “fool”, nor “fool” more specific than “liar”). Since flat structures are preferred by MDC, the final discourse structure is: $\pi_0 : Accept(\alpha^r, \pi^b) \wedge Continuation(\pi^b, \pi^f)$. Its dynamic semantics entail that Bob is committed to Harry being the biggest liar and also maybe the biggest fool.

Now consider (@16ab’), where the ULFs are (54a) and (54b’). First compute how just the presupposed background content π^b would update α^r . As with (54ab), binding P to the available antecedent *liar* and binding x to Harry is preferred (via MDC). The underspecified modality $?_{mod}$ can resolve to \top (in which case Bob’s response can be paraphrased as *Maybe he is*), \diamond (assuming that $\diamond \diamond \phi \models \diamond \phi$, this is also paraphrased as *Maybe he is*) or \neg (assuming $\neg \diamond \phi$ is equivalent to $\Box \neg \phi$, this is paraphrased as the correction *He’s definitely not*). All of these yield coherent updates, and render the proffered content π'^f not given. So the example passes the interest test.

However, when updating the context with both π'^b and π'^f a clearly preferred resolution of the modality emerges via MDC. Specifically, $?_{mod} = \neg$ maximises contrast. Compare:

- (55) a. *Harry is definitely not the biggest liar, but he is maybe the biggest fool*
 b. *Harry is maybe the biggest liar, ?but (also) maybe the biggest fool.*

Saying what Harry is *vs.* what he is *not* yields a better contrast than contrasting two things that he is. So by MDC, $?_{mod}$ resolves to \neg . Thus, the contents of π^b and α^r are contrary, so the glue logic infers a *Correction* relation, yielding (56).

$$(56) \quad \begin{array}{l} \pi_0 : \text{Correction}(\alpha^r, \pi^b) \wedge \text{Contrast}(\pi^b, \pi^f). \\ \alpha^r : \text{liar}(e_{\alpha^r}, h) \wedge \text{biggest}(e_{\alpha^r}) \quad \text{Harry is the biggest liar.} \\ \pi^b : \neg \diamond (\text{liar}(e_{\alpha^r}, h) \wedge \text{biggest}(e_{\alpha^r})) \quad \text{Harry is not the biggest liar.} \\ \pi^f : \diamond (\text{fool}(e_{\beta}, h) \wedge \text{biggest}(e_{\beta})) \quad \text{Harry is maybe the biggest fool.} \end{array}$$

Now we give the formal account of how our semantics predicts felicity judgements and implicatures in the context of *wh*-questions.

- (@2) a. Harvey: Who likes Michael?
b.' Jessica: Nobody likes Michael_{f-r}.

Definition 2 yields the ULF of (@2b') given by π , π^b and π^f in (57), and these must update Harvey's move α^r (where we have already resolved the presupposition triggered by the *wh*-question). Here, α^f and α^b are equally coherent attachment points for the proffered content π^f ; it attaches to α^f as an answer and to α^b as a *Correction*. So α^r is their supersegment.

$$(57) \quad \begin{array}{l} \alpha^r : \text{Background}(\alpha^b, \alpha^f). \\ \alpha^f : ?\lambda x. \text{like}(x, m) \quad \text{Who likes Michael?} \\ \alpha^b : \exists y. \text{like}(y, m) \quad \text{someone likes Michael} \\ \pi : \text{Contrast}(\pi^b, \pi^f). \\ \pi^b : \partial ?_{mod} (\neg \exists z. \text{like}(z, x)) \wedge x = ? y (\text{nobody likes } x) \\ \pi^f : \neg \exists z. \text{like}(z, m) \quad \text{nobody likes Michael} \end{array}$$

First check that proffered content isn't given. The glue logic validates that π^b binds to α^b , thereby resolving $?_{mod} = \neg$ and $x = m$. So interest is satisfied — π^f isn't entailed by this result. But $x = m$ *cannot* be a part of the *final* update with π^f , whatever the resolution of $?_{mod}$:

- (58) a. $?_{mod} = \top \rightsquigarrow \pi^b$: nobody likes Michael.
b. $?_{mod} = \diamond \rightsquigarrow \pi^b$: possibly, nobody likes Michael.
c. $?_{mod} = \neg \rightsquigarrow \pi^b$: somebody likes Michael.

(58c) is inconsistent with π^f and the resolutions in (58a) and (58b) each fail to establish a contrast.

Thus, in the update with both the background and proffered content (step (v) in Definition 4), x is accommodated via an existential quantifier (rather than binding it to m). For the usual reasons, resolving $?_{mod}$ to \neg maximises contrast with π^f

(someone likes someone (other than Michael), but nobody likes Michael), which yields (59):

$$(59) \quad \pi_0 : \text{Background}(\alpha^f, \pi^b) \wedge \text{Contrast}(\pi^b, \pi^f) \wedge \\ \text{Correction}(\alpha^b, \pi^f) \wedge \text{QAP}(\alpha^f, \pi^f)$$

$\alpha^f : ?\lambda z. \text{like}(z, m)$	<i>Who likes Michael?</i>
$\alpha^b : \exists z. \text{like}(x, m)$	<i>someone likes Michael.</i>
$\pi^b : \exists x. \exists z. \text{like}(z, x)$	<i>there is someone that somebody likes.</i>
$\pi^f : \neg \exists z. \text{like}(z, m)$	<i>nobody likes Michael.</i>

Here, Jessica answers Harvey's question but corrects its presupposition. Due to the resolution of π^b Jessica tacitly acknowledges that there was a liking, but she denies that anyone likes Michael. While this sounds odd for *liking*, the relevance of this acknowledgement becomes more apparent when considering a telic verb like *vote*. Replacing 'like' by 'vote' in (@2) yields the interpretation that Jessica acknowledges that *there was a vote* but proffers that *nobody voted for Michael*.

If (@2b') were uttered with a falling pitch contour, then the background ULF is $\partial \neg \exists y. \text{like}(y, x)$, and the most coherent way to interpret this sets $x = m$, so that it corrects α^b . But this entails the proffered content, and so by Definition 3 it is anomalous. Thus with stress on *Michael*, the fall-rise contour successfully voices the denial of the question's presupposition but the falling contour doesn't.

Now, we analyse the infelicity of (@4).

$$(@4) \quad \text{a. Harvey: Who likes Michael?} \\ \text{\#b.' Jessica: Rachel likes } \underline{\text{Michael}}_{\text{fall}}.$$

As before, the question (@4a) yields the SDRS rooted at α^r in (60); (@4b') yields the ULFs π , π^b and π^f :

$$(60) \quad \alpha^r : \text{Background}(\alpha^b, \alpha^f) \\ \alpha^b : \exists x. \text{like}(e_\alpha, x, m) \\ \alpha^f : ?\lambda y. \text{like}(e_\alpha, y, m) \\ \pi : ?_{\text{topic}}(\pi^b, \pi^f) \\ \pi^b : \partial \text{like}(e'_\pi, r, z) \wedge z = ? \\ \pi^f : \text{like}(e_\pi, r, m)$$

For the interest test, update α^r with π^b . The most coherent update resolves z to m , as this can attach as an answer to α^f . As then π^b entails π^f , the test fails — (@4a,b') is infelicitous.

If (@4b') is uttered with a fall-rise contour, it remains infelicitous: If Jessica's response has a fall-rise contour, then the ULFs are as above except that π^b features the underspecified modality $?_{\text{mod}}$ and *Contrast* replaces $?_{\text{topic}}$. The most coherent update of α^r with π^b still resolves z to m and $?_{\text{mod}}$ to \top because this results in

an answer to (@4a), whereas $?_{mod} = \neg$ or $?_{mod} = \diamond$ provide only *partial* answers (SDRT makes complete answers more coherent). So this is also infelicitous.

All semi-formal analyses from Section 4 can be analogously formalised. For instance, consider (@17). The presupposition of (@17b) — $y(\textit{Harvey lives in } x)$ — is initially resolved to $x = \textit{Cleveland}$, so the interest test is passed. Computing maximal contrast (resolving $?_{mod} = \neg$) yields a logical form where (@17b) can be paraphrased as *I don't live in Cleveland, but in NYC*.

(@17) a. Louis: I heard you live in Cleveland now.

b. Harvey: I live in New York City_{f-r}

Thus, in conclusion, our account solves the challenges we raised in Section 3 and retains the good predictions of prior models. We account for data like (@4) that motivate congruence and also for (@2) and predict the implicatures missed by Steedman (2014) (Section 3.6). Also, we replicate Constant's (2012) good predictions about the fall-rise contour (Section 4.7), but not the bad ones (Section 3.7).

6 Conclusion

We have proposed a formal semantic analysis of two pitch contours in English discourse. We model focus and contour jointly: both the *placement* of the stress and the overall *contour* determine which (underspecified) presupposition gets triggered. The most coherent interpretation of this presupposition in context then determines whether the focal element (determined by where the stress is placed) can be *interesting* (and so felicitous), or not. We do not provide a model theory of interest, nor even a definition. Rather, we impose a necessary condition — interesting content must be coherently related to the given information, but not be entailed by it.

While there may be ambiguity in mapping a raw acoustic signal to a specific intonation contour (Calhoun 2007), we postulated no ambiguity in the mapping from such a specific contour to its meaning representation. Distinct pragmatic interpretations in distinct contexts are then determined entirely by how discourse coherence interacts with linguistic and nonlinguistic content, according to principles also used in analysing other linguistic phenomena such as anaphora, elided constructions and presuppositions (Hobbs 1985, Kehler 2002, Asher & Lascarides 1998). The fall-rise contour is often a vehicle for conveying content indirectly, but what exactly is implicated varies radically from one context to another. We capture that variation by underspecifying certain semantic elements, and then capturing how those elements are resolved to specific values via reasoning about discourse coherence and its interaction with compositional and lexical semantics, and world knowledge.

Our model extends the empirical coverage of prior accounts, both in terms of felicity judgements and predicted implicatures. However, we have not modelled certain perlocutionary effects that the prior literature attests: that the fall-rise can

sometimes add an implicature paraphrasable as *I thought you knew* (Ladd 1980, Steedman 2014). We ignore them here, but have formalised them elsewhere (Schlöder & Lascarides 2015). Further, we have here treated only two pitch contours — so we offer just a first step towards achieving a coverage as broad as Steedman (2014). But ours achieves formally precise *pragmatic* derivations, not just semantic ones.

We hope to have demonstrated the pitfalls of ignoring contour in the study of focus, but also to have shown a promising avenue to address them by combining principles of discourse coherence with underspecified presuppositions and a notion of what can be interesting.

References

- Asher, Nicholas. 1999. Discourse and the focus/background distinction. In Peter Bosch & Rob van der Sandt (eds.), *Focus: Linguistic, cognitive, and computational perspectives*, 247–267. Cambridge: Cambridge University Press.
- Asher, Nicholas & Alex Lascarides. 1998. The semantics and pragmatics of presupposition. *Journal of Semantics* 15(3). 239–300. <http://dx.doi.org/10.1093/jos/15.3.239>.
- Asher, Nicholas & Alex Lascarides. 2003. *Logics of conversation*. Cambridge: Cambridge University Press.
- Asher, Nicholas & Alex Lascarides. 2011. Reasoning dynamically about what one says. *Synthese* 183(5). 5–31. <http://dx.doi.org/10.1007/s11229-011-0016-4>.
- Beaver, David I & Brady Z Clark. 2009. *Sense and sensitivity: How focus determines meaning*. Oxford: Wiley–Blackwell.
- Boersma, Paul & Vincent van Heuven. 2001. Praat, a system for doing phonetics by computer. *Glott International* 5(9/10). 341–345. <http://dx.doi.org/11245/1.200596>.
- Bolinger, Dwight. 1972. Accent is predictable (if you're a mind-reader). *Language* 48(3). 633–644. <http://dx.doi.org/10.2307/412039>.
- Bolinger, Dwight. 1982. Intonation and its parts. *Language* 58(3). 505–533. <http://dx.doi.org/10.2307/413847>.
- Bolinger, Dwight. 1985. Two views of accent. *Journal of Linguistics* 21(1). 79–123. <http://dx.doi.org/10.1017/S0022226700010033>.
- Bresnan, Joan W. 1971. Sentence stress and syntactic transformations. *Language* 47(2). 257–281. <http://dx.doi.org/10.2307/412081>.
- Büring, Daniel. 2003. On D-trees, beans, and B-accent. *Linguistics and Philosophy* 26(5). 511–545. <http://dx.doi.org/10.1023/A:1025887707652>.
- Büring, Daniel. 2004. Focus suppositions. *Theoretical Linguistics* 30(1). 65–76. <http://dx.doi.org/10.1515/thli.2004.007>.

- Büring, Daniel. 2006. Focus projection and default prominence. In Valéria Molnár & Susanne Winkler (eds.), *The architecture of focus*, 321–346. Berlin: de Gruyter.
- Büring, Daniel. 2016. (Contrastive) topic. In Caroline Féry & Shinichiro Ishihara (eds.), *The oxford handbook of information structure*, 64–85. Oxford: Oxford University Press. <http://dx.doi.org/10.1093/oxfordhb/9780199642670.013.002>.
- Calhoun, Sasha. 2007. *Information structure and the prosodic structure of english: A probabilistic relationship*: University of Edinburgh dissertation. <http://hdl.handle.net/1842/8120>.
- Carston, Robin. 1996. Metalinguistic negation and echoic use. *Journal of Pragmatics* 25(3). 309–330. [http://dx.doi.org/10.1016/0378-2166\(94\)00109-X](http://dx.doi.org/10.1016/0378-2166(94)00109-X).
- Constant, Noah. 2012. English rise-fall-rise: a study in the semantics and pragmatics of intonation. *Linguistics and Philosophy* 35(5). 407–442. <http://dx.doi.org/10.1007/s10988-012-9121-1>.
- Dryer, Matthew S. 1996. Focus, pragmatic presupposition, and activated propositions. *Journal of Pragmatics* 26(4). 475–523. [http://dx.doi.org/10.1016/0378-2166\(95\)00059-3](http://dx.doi.org/10.1016/0378-2166(95)00059-3).
- Féry, Caroline & Shinichiro Ishihara (eds.). 2016. *The oxford handbook of information structure*. Oxford: Oxford University Press. <http://dx.doi.org/10.1093/oxfordhb/9780199642670.001.0001>.
- Geurts, Bart & Rob van der Sandt. 2004a. Interpreting focus. *Theoretical Linguistics* 30(1). 1–44. <http://dx.doi.org/10.1515/thli.2004.005>.
- Geurts, Bart & Rob van der Sandt. 2004b. Interpreting focus again. *Theoretical Linguistics* 30(2–3). 149–161. <http://dx.doi.org/10.1515/thli.2004.30.2-3.149>.
- Halliday, Michael A K. 1967. Notes on transitivity and theme in english: Part 2. *Journal of Linguistics* 3(2). 199–244. <http://dx.doi.org/10.1017/S0022226700016613>.
- Hara, Yurie & Robert van Rooij. 2007. Contrastive topics revisited: A simpler set of topic-alternatives. *North East Linguistic Society (NELS)* 38.
- Hobbs, Jerry. 1985. On the coherence and structure of discourse. Tech. Rep. CSLI-85-37 Center for the Study of Language and Information, Stanford University.
- Hobbs, Jerry. 1990. The Pierrehumbert-Hirschberg theory of intonational meaning made simple: Comments on Pierrehumbert and Hirschberg. In Jerry Morgan Philip R Cohen & Martha E Pollack (eds.), *Intentions in communication*, 313–323. Cambridge, MA: MIT Press.
- Hobbs, Jerry R, Mark E Stickel, Douglas E Appelt & Paul Martin. 1993. Interpretation as abduction. *Artificial Intelligence* 63(1–2). 69–142. [http://dx.doi.org/10.1016/0004-3702\(93\)90015-4](http://dx.doi.org/10.1016/0004-3702(93)90015-4).
- Horn, Laurence. 1989. *A natural history of negation*. Chicago, IL: University of Chicago Press.

- Jackendoff, Ray. 1972. *Semantic interpretation in generative grammar*. Cambridge, MA: MIT Press.
- Kehler, Andrew. 2002. *Coherence, reference and the theory of grammar*. Stanford, CA: CSLI Publications.
- Kratzer, Angelika. 1989. An investigation of the lumps of thought. *Linguistics and philosophy* 12(5). 607–653. <http://dx.doi.org/10.1007/BF00627775>.
- Krifka, Manfred. 1991. A compositional semantics for multiple focus constructions. In Steven K Moore & Zachary Wyner Adam (eds.), *Proceedings of the 1st Semantics and Linguistic Theory Conference (SALT 1)*, 127–158. <http://dx.doi.org/10.3765/salt.v1i0.2492>.
- Ladd, D Robert. 1980. *The structure of intonational meaning: Evidence from english*. Bloomington & London: Indiana University Press.
- de Marneffe, Marie-Catherine & Judith Tonhauser. 2019. Inferring meaning from indirect answers to polar questions: The contribution of the rise-fall-rise contour. In Malte Zimmermann, Klaus von Stechow & Edgar Onea Gaspar (eds.), *Questions in discourse*, Leiden: Brill Publishing. http://dx.doi.org/https://doi.org/10.1163/9789004378322_006.
- Partee, Barbara. 1991. Topic, focus and quantification. In Steven K Moore & Adam Zachary Wyner (eds.), *Proceedings of the 1st Semantics and Linguistic Theory Conference (SALT 1)*, 159–188. <http://dx.doi.org/10.3765/salt.v1i0.2918>.
- Pierrehumbert, Janet & Julia Hirschberg. 1990. The meaning of intonational contours in the interpretation of discourse. In J. Morgan P. R. Cohen & M. E. Pollack (eds.), *Intentions in Communication*, 271–311. Cambridge, MA: MIT Press.
- Reese, Brian. 2007. *Bias in questions*: University of Texas at Austin dissertation. <http://hdl.handle.net/2152/3280>.
- Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5(6). 1–69. <http://dx.doi.org/10.3765/sp.5.6>.
- Rooth, Mats. 1992. A theory of focus interpretation. *Natural Language Semantics* 1(1). 75–116. <http://dx.doi.org/10.1007/BF02342617>.
- Rooth, Mats. 1999. Association with focus or association with presupposition? In Peter Bosch & Rob van der Sandt (eds.), *Focus: Linguistic, cognitive, and computational perspectives*, 232–244. Cambridge: Cambridge University Press.
- Rooth, Mats. 2016. Alternative semantics. In Caroline Féry & Shinichiro Ishihara (eds.), *The oxford handbook of information structure*, 19–40. Oxford: Oxford University Press. <http://dx.doi.org/10.1093/oxfordhb/9780199642670.013.19>.
- Sæbø, Kjell Johan. 2016. Information structure and presupposition. In Caroline Féry & Shinichiro Ishihara (eds.), *The oxford handbook of information structure*, 128–146. Oxford: Oxford University Press. <http://dx.doi.org/10.1093/oxfordhb/9780199642670.013.012>.

- Šafářová, Marie. 2005. The semantics of rising intonation and interrogatives. In Emar Maier, Corien Bary & Janneke Huitink (eds.), *Sinn und Bedeutung* 9, 355–369. <https://ojs.ub.uni-konstanz.de/sub/index.php/sub/article/view/774/665>.
- van der Sandt, Rob. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9(4). 333–377. <http://dx.doi.org/10.1093/jos/9.4.333>.
- Schlöder, Julian J & Alex Lascarides. 2015. Interpreting english pitch contours in context. In Christine Howes & Staffan Larsson (eds.), *Workshop on the Semantics and Pragmatics of Dialogue (SemDial)* 19, 131–139. <http://semdial.org/anthology/papers/Z/Z15/Z15-3018/>.
- Schwarzschild, Roger. 1999. GIVENness, AvoidF and other constraints on the placement of accent. *Natural language semantics* 7(2). 141–177. <http://dx.doi.org/10.1023/A:1008370902407>.
- Selkirk, Elisabeth O. 1984. *Phonology and syntax*. Cambridge, MA: MIT Press.
- Silverman, Kim, Mary Beckman, John Pitrelli, Mori Ostendorf, Colin Wightman, Patti Price, Janet Pierrehumbert & Julia Hirschberg. 1992. ToBI: a standard for labeling english prosody. In *International conference on spoken language processing 2*, 867–870.
- Steedman, Mark. 2014. The surface-compositional semantics of english intonation. *Language* 90(1). 2–57. <http://dx.doi.org/10.1353/lan.2014.0010>.
- Wagner, Michael. 2006. Givenness and locality. In Masayuki Gibson & Jonathan Howell (eds.), *Semantics and Linguistic Theory (SALT)* 16, 295–312. <http://dx.doi.org/10.3765/salt.v16i0.2938>.
- Wagner, Michael. 2012. Contrastive topics decomposed. *Semantics and Pragmatics* 5(8). 1–54. <http://dx.doi.org/10.3765/sp.5.8>.
- Wagner, Michael, Elise McClay & Lauren Mak. 2013. Incomplete answers and the rise-fall-rise contour. In Raquel Fernández & Amy Isard (eds.), *Workshop on the Semantics and Pragmatics of Dialogue (SemDial)* 17, 140–149. <http://semdial.org/anthology/papers/Z/Z13/Z13-3018/>.
- Walker, Marilyn A. 1996a. Inferring acceptance and rejection in dialogue by default rules of inference. *Language and Speech* 39(2–3). 265–304. <http://dx.doi.org/10.1177/002383099603900306>.
- Walker, Marilyn A. 1996b. Limited attention and discourse structure. *Computational Linguistics* 22(2). 255–264.
- Ward, Gregory & Julia Hirschberg. 1988. Intonation and propositional attitude: the pragmatics of L*+ H LH%. In *5th eastern states conference on linguistics*, 512–522.
- Watson, Duane G, Michael K Tanenhaus & Christine A Gunlogson. 2008. Interpreting pitch accents in online comprehension: H* vs. L+H*. *Cognitive Science* 32(7). 1232–1244. <http://dx.doi.org/10.1080/03640210802138755>.

- Westera, Matthijs. 2017. *Exhaustivity and intonation: A unified theory*: University of Amsterdam dissertation. <https://hdl.handle.net/11245.1/25ab109d-1e52-4ba9-bdd8-4719c17c5a3b>.
- Westera, Matthijs. 2018. Rising declaratives of the quality-suspending kind. *Glossa* 1(3). 1–32. <http://dx.doi.org/10.5334/gjgl.415>.

Julian J. Schlöder
ILLC, University of Amsterdam
P.O. Box 94242
1090 GE Amsterdam
The Netherlands
julian.schloeder@gmail.com

Alex Lascarides
School of Informatics,
University of Edinburgh
10 Crichton Street
Edinburgh, EH8 9AB
Scotland
alex@inf.ed.ac.uk