Focus constraints on ellipsis — An Unalternatives account [1]

Muriel ASSMANN — *University of Vienna*
Daniel BÜRING — *University of Vienna*
Izabela JORDANOSKA — *University of Vienna*
Max PRÜLLER — *University of Vienna*

Abstract This paper presents a new account of the generalization that focused elements cannot be elided, framed within Unalternative Semantics, a framework that does away with syntactic F-marking. We propose the mirror image of the generalization: what is elided cannot introduce alternatives. We implement this as a focus restriction in UAS and then go on to show how to account for MAXELIDE effects using the same technique, without making reference to any transderivational constraints.

Keywords: focus, ellipsis, alternative semantics, unalternative semantics, MaxElide, secondary focus.

1. Introduction

In this paper we discuss the interaction of focus and ellipsis in English in the Unalternative Semantics framework (UAS, see Büring 2015, 2016a, b). Consider examples (1) and (2), which illustrate the Focus–Ellipsis Generalization in (3).

(1) (Who was Kim going to kiss?) —
  a. ALEX.
  b. Kim was going to kiss ALEX.

(2) (What was Kim going to do?)
  a. Kim was going to kiss ALEX.
  b. Kiss ALEX.
  c. #ALEX.

(3) The FOCUS–ELLIPSIS GENERALIZATION (FEG): Focal elements cannot be elided.

FEG at first sounds like a truism: If ellipsis is the most radical form of deaccenting, it seems trivial that a focus—the accent bearer *par excellence*—could not be elided. But (2) already shows that things aren't that simple: even though *kiss* need not bear a pitch accent in a VP focus answer like (2a)—and can in principle be elided, as in (1a)—such elision is impossible when the focus is a VP, as in (2c). Instead, only the non-focal subject and auxiliaries can be elided, as in (2b).

The straightforward move would seem to allude to a syntactic marking of focus. If VP were F-

[1]We would like to thank James Gray, Nina Haslinger, the reviewers of SuB 22 and the people who came to our poster presentation. This work was supported by the Austrian Science Fund (FWF), project grant P29180-G23, 'Unalternative Constraints Crosslinguistically'.

 In: Uli Sauerland and Stephanie Solt (eds.),
Proceedings of Sinn und Bedeutung 22, vol. 1, ZASPiL 60, pp. 109–126. ZAS, Berlin.

marked in (2), but the object alone in (1), the ellipsis pattern would follow by equating 'focal' in (3) with 'F-marked'. Note that on this view, the FEG provides an argument for *syntactic* focus marking: without F-markers or something like them, there is nothing in the structure for FEG (or the principles that account for it) to tell the difference between (1a) and (2c).

In this paper we argue against this. We show, in Section 2, how focusing and the FEG can be modelled without F-markers, based on the idea that ellipsis itself contributes to identifying (non-)focal material: what is elided cannot have (non-trivial) focus alternatives, as schematized in (4a) (where ε is a marker for elision of its sister à la Merchant 2004).

(4) a. weaker hypothesis b. stronger hypothesis

Starting in Section 3 we explore a stronger version of this hypothesis, namely that ellipsis furthermore marks its remnant as focal, (4b). This stronger hypothesis turns out to give us a direct implementation of some MAXELIDE effects.

In Section 4 we examine some apparent problems regarding MAXELIDE and extraction and propose a solution to it that invokes SECONDARY FOCUS (SF), while Sections 5 and 6 elaborate on the further predictions we make regarding SF. Section 7 concludes. Data are either taken from previous literature, the Corpus of Contemporary American English (COCA, Davies 2008) or constructed by ourselves and judged by a native speaker of English.

2. Background

In this section we present the Unalternative Semantics framework. We show how it relates to more standard versions of focus semantics, and how the FEG can be captured in it without recourse to syntactic F-markers.

2.1. Unalternative Semantics

Unalternative Semantics (UAS) takes a syntactic tree annotated with metrical weights and directly derives the set of focus alternatives for each node. As such it rolls into one what is done by the rules for stress/accent-to-F, F-projection, and F-interpretation in frameworks that base Roothian alternative semantics on syntactic F-marking.

Crucially, UAS restricts focal alternatives at branching nodes only, in one of two ways. If the metrical weights among sisters are *reversed from the default*—where the default, for the time being, would be*weak–strong*—STRONG RESTRICTION applies.

(5)  STRONG RESTRICTION (whenever s/w is *reversed* from the default): short: $x^{\backslash B}\ C$

A only allows alternatives that differ in B (=strong), but are the same in C (=weak)

Case (5) corresponds closely to the 'traditional'  : B gets to have non-trivial alternatives, those each get combined with the literal meaning of C to form A's alternatives;[2] we write this as $x^{\backslash B}\ C$ ('combine the ordinary meaning of C with any alternative to B, (except the ordinary meaning of B)').

In case B and C show the default weak–strong pattern, a WEAK RESTRICTION is imposed.

(6)  WEAK RESTRICTION (default *w–s* pattern): short: $\overline{x^{\backslash B}\ C}$

A allows alternatives *except* those that differ in B (=weak), but are the same in C (=strong)

This case has no corresponding configuration in an F-marking framework: it contains all alternatives one would get from (7a), plus A's literal meaning ('alternative' to (7b)), *plus those alternatives to A that are not in the alternatives to (7c).*

(7) a.  b.  c. 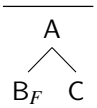

We write this as $\overline{x^{\backslash B}\ C}$, to be read as: any alternative of type A, except those that replace B but not C ('if the weak daughter is replaced, the strong one must be as well').

Finally, restrictions from lower nodes *propagate*, so that for example A in (8a) (default *weak–strong* twice) allows for all alternatives except i) those that 'replace' *B* but not *C*, *and* ii) except those that replace *D* but not *E* (regardless of whether they replace B). Technically, the (weak) restriction introduced on C, $\overline{x^{\backslash D}\ E}$, propagates to A as $\overline{y\ x^{\backslash D}\ E}$ ('no alternatives that replace D and keep E, regardless of whether they replace B'[3]). A itself introduces the weak restriction $\overline{y^{\backslash B}\ C}$, so that the sum total restriction on A is 'does not contain E, except if combined with B and D'. (8b) gives a parallel derivation involving strong restriction.[4]

(8) a. 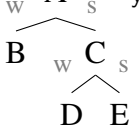 $\overline{y^{\backslash B}\ C}\ \&\ \overline{y\ x^{\backslash D}\ E} \equiv \overline{xy^{\backslash BD}\ E}$ $\overline{x^{\backslash D}\ E}$ b. 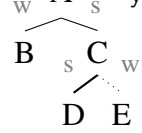 $\overline{y^{\backslash B}\ C}\ \&\ y\ x^{\backslash D}\ E \equiv y\ x^{\backslash D}\ E$ $x^{\backslash D}\ E$

This much background should suffice to understand our proposal regarding ellipsis (see the appendix for more details). In fact, even though UAS knows *four* possible states with respect to

[2]The only difference is that B's literal meaning is not allowed to be used, a fact not relevant to the present paper.

[3]Note that *y* here is not restricted, so it could be the literal meaning of B, or an alternative to it.

[4]To aid reading, strong daughters are set in bold, and weak daughters that have undergone reversal are dotted.

introducing alternatives—1. must be replaced by alternatives (strong sister in reversed structure, SR) 2. must be unchanged in all alternatives (weak sister in reversed structure, SR), 3. may or may not be changed in the alternatives (strong sister in default structure, WR), and 4. may not be changed, unless its sister is changed, too (weak sister in default, WR)—as opposed to standard alternative semantics' two (F or not), our final proposal for ellipsis merely requires 1. and 2., the classical 'focal' and 'non-focal'. This means that the gist of our proposal should be understandable even without the details of UAS. The crucial theorems of UAS on which the proposal relies, though, require those details and are explicated in the appendix.

2.2. The Focus-Ellipsis Generalization within UAS

To make things more perspicuous, we assume a syntactic element ε which marks the deletion/non-spell out of its sister constituent. ε is borrowed from Merchant (2001), and assumed to be subject to contextual restrictions, roughly that the denotation of its sister must be contextually given ('ellipsis under identity'), the exact formulation being irrelevant here.

Crucially, we put an additional restriction on ε, to the effect that its sister must not contain focal material (i.e. constituents that introduce non-trivial alternatives). In standard alternative semantics this would amount to requiring that in *[ε B]*, B only has the trivial alternative, its literal meaning; indirectly this ensures that B does not (bear or) dominate any F markers. We get the same effect in UAS requiring that the only alternative allowed for *[ε B]* is (the literal meaning of) B, written as B. Take the term answer *ALEX* from (1a) and (2) above, which we assume to be represented as in (9).

(9) ALEX ε ~~Kim was going to kiss t_{Alex}~~

By virtue of ε, *Kim was going to kiss* is marked as non-focal, so all focus alternatives at the sentence level will be built around that property, i.e. $\mathsf{Kim\ was\ going\ to\ kiss}\ x$. This makes (9) a good answer to the question 'Who was Kim going to kiss?', but not 'What was Kim going to do?'. So the FEG is turned around: We do not prohibit deleting something focal, but rather mark something that might otherwise contain focal material as *non*-focal in the process of ellipsis (i.e. as a condition on the presence of ε).[5]

It bears pointing out that the problem with the term answer *ALEX* to a VP-question as in (2c) is

[5]There is a complication here in that the question–answer condition (QAC) used with UAS in Büring (2015)—that A can answer Q if at least some answers in $[\![Q]\!]$ are permitted as alternatives of A—actually fails to rule out an answer with only 'Kim was going to kiss *x*' alternatives as an answer to a VP question like 'What was Kim going to do?'. The reason is that even an answer like *Kim was going to kiss ALEX* ('VP focus') does not have all propositions of the form 'Kim was going to *Q*' as possible focus alternatives; it lacks those in 'Kim was going to $\mathsf{Q}^{\backslash kiss}$ Alex' (for good reasons); therefore Büring (2015) relaxed the QAC so as to be content as long as *some* answers to the question are also permitted alternative of the answer, which, alas, is *also* the case if the permitted focus alternatives are just 'Kim was going to kiss *x*' or just 'Kim was going to *R* Alex'. The correct version of QAC should be one that does not mind if, say, an answer to 'What did Kim do?' lacks *some* of the 'Kim *Q*' propositions as alternatives, but *does* mind if it *only* allows alternatives that are of the sort 'Kim kiss *x*' or 'Kim *R* Alex'. But since QAC does not know about the 'form' of propositions, it is unclear to us at this point how to best state such a condition, so we leave this for another occasion.

not in general accounted for by demanding that the complement of ε is given (as is standardly assumed, e.g. in Merchant 2001). To see that, consider (10):

(10) (Kim was going to call me, but then my phone went dead. What is Kim going to do now?) — *(call) ALEX

In (10) 'Kim was going to call' is made contextually salient, so the elided part of (9), ~~*Kim is going to call*~~, is given. Yet, (10) is no better than (2c). We conclude that the oddness of (2c), like that of (10), is not attributable to eliding something non-given (*Kim called* in (2c)); rather, we submit, the problem is that in both cases *call* is marked as non-focal (by virtue of being elided), although it is part of the focus. The standard treatment of ε alone thus does not derive the FEG.

On the treatment just proposed, the restriction imposed by ε is different from both Weak and Strong Restriction; it is just the weak daughter condition of Strong Restriction ('can't be focal'), with out its relational counterpart ('must be focal'). While this is certainly possible to do in UAS, it may be worth exploring the idea that ε does in fact impose a Strong Restriction: not only must the elided part be non-focal, its sister (or, counting ε itself: its aunt) must be focal, as schematized in (4b) above.

3. The remnant must be focal

If ε imposes a SR, this means that the immediate remnant of an ellipsis must be focal. This seems generally on the right track, and in particular, it derives a number of so-called MAXELIDE effects, exemplified in (11).

(11) a. *John saw something, but we don't know WHAT he did.
b. John saw something, but we don't know WHAT.

The definition of MaxElide, as stated originally in Merchant (2008) is given in (12).

(12) MAXELIDE: Let XP be an elided constituent containing an A'-trace. Let YP be a possible target for ellipsis. YP must not properly contain XP. (Merchant 2008: p.141)

To a first approximation, MAXELIDE means that a smaller ellipsis like VP in (11a) is ungrammatical in a context where a bigger ellipsis, like TP in (11b), is possible. Using ε with a Strong Restricton, as suggested above, the two competing structures are given in (13).[6]

(13) a.

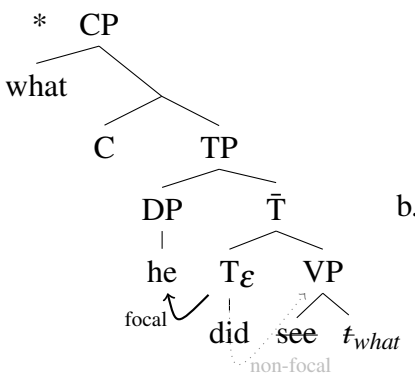

b.

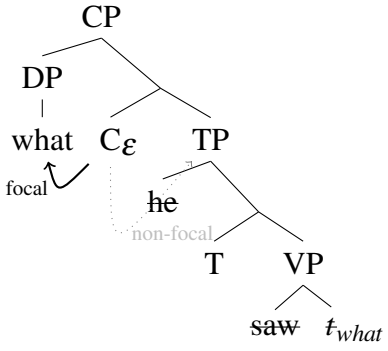

In (13a) *he* needs to be focal, by ε, and there should be a contrastive target of the form 'somebody else saw something', which there isn't. Furthermore, *what* fails to be marked as focal in (13a), though it clearly is the element contrasting with *something*.

In (13b), on the other hand, *what* is marked as focal, and everything else as background, which exactly matches the context in (11). So using a strong restriction with ε not just makes sure that nothing focal is elided (which is the case in both (13a) and (13b)), but also that the final non-elided remnant is itself focal. This captures the basic MaxElide effect. What is more, it does so without invoking a transderivational constraint, which most other accounts of MaxElide effects we know of do.[7]

4. Apparent Problem: MaxElide and Extraction

What our story so far does not capture is that MaxElide effects appear to be restricted to cases in which the ellipsis site contains a trace. Thus in (14), which does not involve extraction, both small (lower VP) and maximal (higher VP) ellipsis are equally acceptable, apparently ignoring MaxElide.

[6]Note that unlike in our earlier trees, ε in (13) is sitting on preterminal nodes—T and C, respectively—rather than being adjoined; as a consequence, the SR does no longer regard sisters, but aunt–niece pairs (here: SpecT and VP/SpecC and TP). We believe that this is at least an option (though it wouldn't make a difference in (13), since there ε could also just adjoin to $\bar{C}/\bar{T}$), as English in some cases allows for a non-focal head between the ellipsis site and the closest focus. This is shown in (i), where an ε adjoined right above the ellipsis site, as in (ib), would wrongly force *did* to be focal.

(i) They left before BO did ~~leave~~ a.

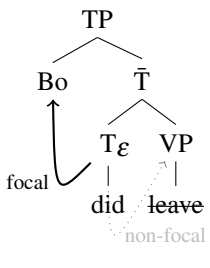

b. *

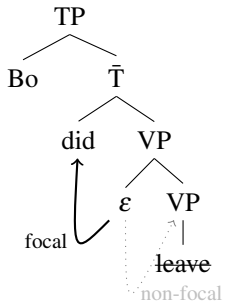

We are not concerned with what categories can and cannot be elided in English in general. See e.g. Miller and Pullum (2013) for further information.

[7]See Messick and Thoms (2016) and Griffiths (2017) for different accounts of deriving MaxElide effects.

(14) John said Mary likes Peter. a. No, BILL said she **does**. b. No, BILL did.

According to what we said so far, (14a) should only be possible if *does* (or *she*) were focal. But is it? Clearly it does not bear the nuclear pitch accent, which in both cases is on *Bill* (as one would expect). We submit, however, that *does* is a SECONDARY FOCUS (SF), i.e. a focus that is contained in the background of the main focus (*BILL*). As in (14a), we indicate a SF by boldface, reflecting that it is marked by stress, but not accented (which would be indicated by capitals). The structures for (14a) and (14b) are then as in (15a) and (15b), respectively.

(15) a.

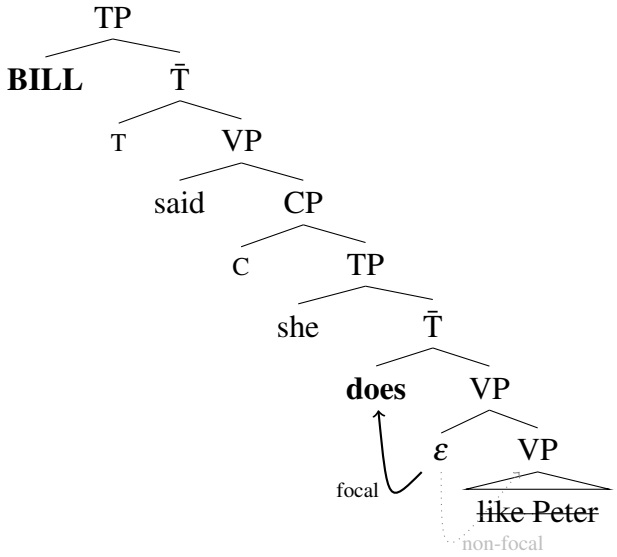

b.

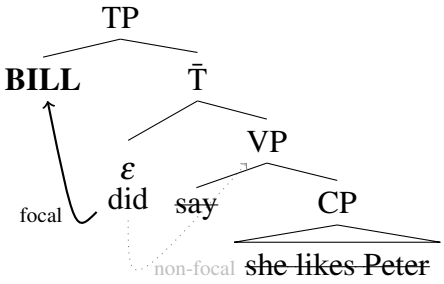

Section 6 below provides independent evidence that, indeed, the final remnant before an ellipsis is always focused, even when not accented.

But for now we want to make sure that the introduction of SFi as a general possibility does not throw out the proverbial baby with the bathwater: If a SF can generally obscure the effect of MAXELIDE in this way, why are there unacceptable cases at all? That is to say, why can we not claim that (13a) also contains a SF on *did* or *he*, as in (16)?

(16) John saw something, but we don't know...

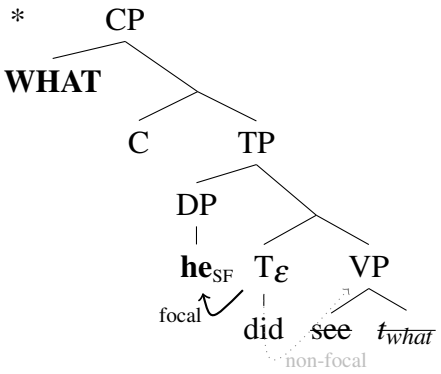

Our answer to this question closely follows the one given in Takahashi and Fox (2005), who argue that if an ellipsis site, say VP, contains an unbound variable, its antecedent cannot be

just a VP; rather one has to find an antecedent matching a PARALLELISM DOMAIN, which includes, in addition to the ellipsis site, the binders for the variable. For example, while *[$_{\text{VP}}$likes Peter]* in (14a)/(15a) can directly be licensed by *[$_{\text{VP}}$likes Peter]* in the first sentence (their semantic identity can be established at the VP level), *[$_{\text{VP}}$see t_{what}]* in (16) cannot, because it contains an unbound trace (roughly, we do not know its denotation, and hence cannot establish synonymy with any antecedent). Instead a parallelism domain that includes the antecedent, *what*, is required, i.e. the entire CP. In (15), on the other hand, any constituent containing the ellipsis site (and even the ellipsis site itself) is a potential PD (as there are no variables waiting to be bound).

Takahashi and Fox (2005) then cash out MAXELIDE as 'Delete the biggest deletable constituent within a PD'. For example, since the minimal PD in (16) is CP, ellipsis must delete TP, not just VP. The minimal PD in (15a), on the other hand, is the elided VP itself (since there are no unbound traces within it), within which, trivially, that VP is the biggest deletable constituent. If one picks the matrix VP or TP instead, the biggest deletable constituent within *that* is the matrix VP. So still, for any given PD, the ellipsis site is maximal, it's just that there are various choices of PD. Where there is a trace involved, however, choices are effectively restricted to a domain containing at least the antecedent; (17) below illustrates again what rules out smaller PDs for such a case, here (16).

(17) John saw something, but we don't know

a.	*what he did [~~see t_{what}~~]$_{\text{PD}}$.	PD contains trace, but not antecedent
b.	*[what he did ~~see t_{what}~~]$_{\text{PD}}$.	PD ok, ellipsis not maximal in PD
c.	[what ~~he did see t_{what}~~]$_{\text{PD}}$.	PD contains antecedent, ellipsis maximal

In our proposal, the role of Takahashi & Fox's PD is roughly played by the DOMAIN OF THE SECONDARY FOCUS, i.e. the domain that consists of the secondary focus and its background. Like Takahashi & Fox's PD, the domain of a focus is in principle free (so long as it contains the focus, of course), provided it does not contain unbound variables. As a consequence, if the SF c-commands a trace, its domain must include the antecedent of that trace, again just like Takahashi & Fox's PD.

Crucially, and unlike in the case of Takahashi & Fox's PD, there is also an upper limit on the choice of the domain for a SF: it cannot include the main focus. This is a consequence of the UAS mechanism, as we demonstrate in the appendix. For now, we simply state that the choice of domain for the SF is limited: big enough to contain the antecedent of a trace (if there is one), but not including the main focus.

Importantly, these two conditions cannot possibly both be met in case the extractee is itself the main focus of a sentence. For example, in (16), repeated in (18a), the domain of the SF *he* has to include *what* so as to have the VP internal trace bound, and at the same time must *not* include *what*, since that is the main focus. This dilemma will present itself whenever the extractee is the main focus; in other words, there can be no SF in such cases, and hence no non-maximal ellipsis.

(18) a.

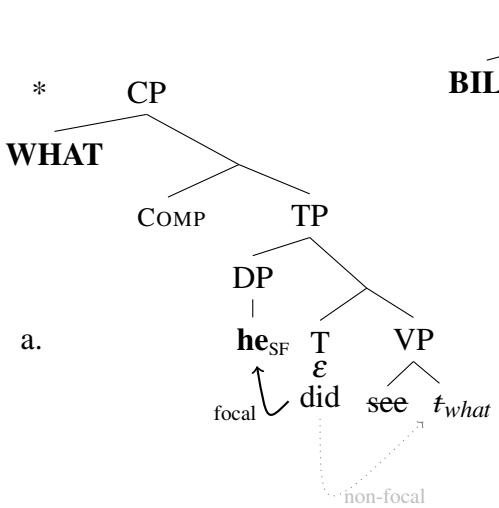

b.

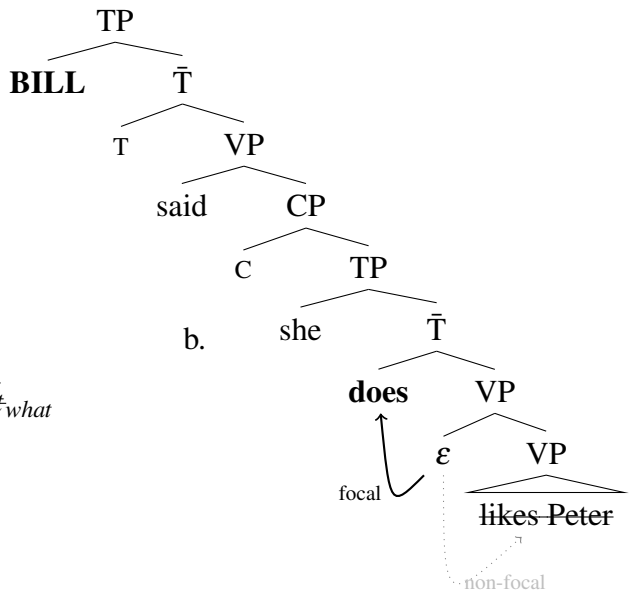

In a non-extraction case like (15a), repeated in (18b), no such problem arises: Since there is no trace waiting to be bound, the domain of the SF on *does* can be as small as (embedded) TP, CP or (higher) VP, which all exclude the main focus *Bill*. We thus predict that—other constraints on the choice of the focus domain notwithstanding—ellipsis size is flexible in these cases.

This concludes our derivation of MAXELIDE effects and their (apparent, see Section 6) absence in non-extraction contexts. In the next Section we will discuss in more detail the interaction between SF and extraction. But first, let us take stock: the basic MAXELIDE effect is written into the semantics of ellipsis itself, i.e. once ε/ellipsis is used, not only does the elided part need to be non-focal, but the nearest remnant must be focal. No transderivational constraints are needed (cf. 'biggest deletable'), it's just that any smaller ellipsis would wrongly mark something as focal that ought to be in the background. Furthermore there is now a reason why the size of what Takahashi and Fox (2005) call P(arallelim) D(omain) relates to the minimum ellipsis size: The minimal PD is the domain in which no S(econdary) F(oci) can occur. Lastly, our proposal directly answers the question why, unlike similar principles like Maximize Presuppositions (Heim 1991), MAXELIDE does not penalize the complete absence of ellipsis even where possible, that is, why the urge to elide as much as possible is only activated once *some* ellipsis has taken place: the alleged 'principle' MAXELIDE is just a consequence of the conditions that come with ε; no ε—no MAXELIDE effect.

5. More Complex Interactions Between Ellipsis and Extraction

So far, we have paid attention only to configurations in which the phrase extracted from the elided VP was itself the primary focus. However, the generalization we have derived is a different one: that MAXELIDE effects will be observed *unless* the primary focus is higher than the extractee. This section will look at the two sides of this prediction that were not discussed so far: cases in which the primary focus is higher than the extractee, and cases in which it is lower.

5.1. Focus Below the Extractee

These cases have actually been discussed in the literature a lot: A focus below the extractee *forces* ellipsis to be smaller.[8]

(19) a. I think YOU should ride the TALLEST camel, but I don't know which one PHIL should.
b. I don't know which puppy you should agree to adopt, but I know which one you should NOT.

In existing accounts this is because a focus renders any constituent containing it undeletable. On our account, basically the same holds: As soon as a constituent is not given, it cannot be below ε. Another prediction we make for these cases is that ellipsis *does* have to be maximal below that focus. We share this prediction with Takahashi and Fox (2005), who contrast (19b) with the unacceptable (20), in which the ellipsis ends unnecessarily far below the focused *not*.

(20) *I don't know which puppy you should agree to adopt, but I know which one you should NOT agree to.

In contradistinction to that, Griffiths (2017) claims that a focus underneath the extractee basically neutralizes MaxElide, providing examples such as those in (21) (his (13)).

(21) a. I know who MARY thinks he'll kiss and also who SUE thinks he will.
b. I know who BILL hopes to kiss and also who BOB hopes to.

Our account clearly predicts these to be ungrammatical, so to the extent that Griffiths' judgements are shared[9], more needs to be said here. On the other hand, we correctly predict the contrast between (19b) and (20), which is surprising given Griffiths' account.

5.2. Focus above the Extractee

We assume that no MaxElide effects will be observed if there can be a (variable-free) focus domain below the main focus. So far we looked at cases in which the main focus was an element extracted from the ellipsis site, so that the *de facto* there could be no MaxElide obviations in structures with extraction.

But in principle, an extractee need not be focused, in which case we predict that MaxElide obviations are possible. A pertinent example is given in (22) ((33) from Merchant 2008):

(22) ??Ben knows who she invited, but Charlie DOESN'T know who she did.

[8]Examples (48) from Schuyler (2001) and (32) from Takahashi and Fox (2005).
[9]Not all speakers accept these examples, so there seems to be a fair amount of inter-speaker variation. Furthermore, some speakers find (21b) worse than (21a), we leave this open for now.

According to our proposal, if *did* in (22) is a SF, it could take either the *who* clause or the VP headed by *know* (i.e. anything below the sentential focus *doesn't* and above the extractee) as its focus domain (FD). In either case, there shouldn't be a problem and the resulting sentence is predicted to be grammatical. (23) illustrates the case where the *who* clause is the FD of *did*:

(23) DOESN'T know who she **did** ~~invited t_{who}~~ domain of SF *did*

The literature partly bears out this prediction in that Merchant (2008) judges (22) better than 'regular' MAXELIDE violations—i.e. ones in which the extractee is focused—such as (24) (his (30)).

(24) *I know we invited SOMEone, but I cant remember WHO we did.

But (22) is still judged as degraded, which is not predicted by what we said so far. We are not sure what the cause of this degradation is; there are reasons to believe, though, that it is unrelated to the MAXELIDE effect. Observe that Merchant judges (22) on a par with (25), which, notably, involves extraction of an adjunct, *when*.

(25) ??Abby knew when he had quit, but Beth DIDN'T know when he had.

Crucially, adjunct extraction generally does not lead to MAXELIDE effects at all (as they do not involve VP *internal* traces, cf. again (12)), as has been observed in the literature, (26) (ex. (16a) from Schuyler 2001), and indeed we found examples of this kind in the COCA corpus (Davies 2008), e.g. (27).[10]

(26) I think you should adopt one of these puppies, but I don't know WHEN you should.

(27) a. Sean: And the airline was not willing or forthcoming today, General McInerney, with any information about why he stopped his training back in 2009. That's somewhat of a puzzle, right?
b. General McInerney: It's very puzzling, Sean. **We have to do a deep dive into that to see why he did.** Did he suffer depression?

Likewise (25)'s counter-part in (28) seems impeccable.

(28) Abby knew that he had quit, but not WHEN he had.

Based on this, it seems justified to assume that (22) and (25) are degraded for the same reason, and that, in the light of (25), that reason should crucially *not* be related to the explanation for MAXELIDE effects, but something independent. While of course we would like to know what that something is, we will have to leave that question for another occasion.[11] For now, our account

[10] The context is provided in (27a), while the relevant sentence is marked in bold in (27b).

[11] The crucial difference between (28) and (25) appears to be that the main focus has 'moved' from the *wh*-word to something further to the left, so that the *wh*-word itself becomes part of the background, as schematized in (i).

(i) MAIN FOCUS... wh-phrase... aux ~~...t_{wh}...~~/ t_{wh}~~...~~

predicts—correctly, as we just argued—that (22), like (25), does *not* show a MAXELIDE effect, and that both are of equal acceptability, which is significantly higher than 'classic' MAXELIDE violations.

6. Effects of SF

In this Section we motivate the assumption that cases of non-maximal ellipsis indeed involve a S(econdary) F(ocus). Or put differently, that ellipsis is really always maximal, once we consider SF, and as predicted by our claim that ellipsis itself marks its closest remnant as focal.

For this, we need to look at cases that do not involve extraction from the ellipsis site, such as those in (29) ((35) from Merchant 2008: with indication of accents added).

(29) a. Ben knows that she invited Klaus, but her father DOESN'T
b. Ben knows that she invited Klaus, but her father DOESN'T know that she did.

According to existing accounts of MAXELIDE effects, the choice between (29a) and (29b) is optional, as the choice of ellipsis size is in general where no extraction from within VP is involved. On the present proposal, (29b) must involve a SF on *did*, that is: it emphasizes that her father assumes that she *didn't* invite Klaus. While this seems consistent with our intuition, it turns out difficult to really pinpoint these aspects of meanings. In particular, we would like to see cases in which a non-maximal ellipsis is unacceptable because its remnant cannot be focal for independent reasons.

To do so, we will concentrate on *associated* foci to demonstrate the connection between ellipsis size and focus interpretation. For an illustration consider (30).

(30) A: Bob will only tell me WHERE he's going (not when).
B: (i) I wonder WHY.
(ii) *I wonder WHY he will only tell you.
(iii) I wonder WHY he will only tell you where.

By inserting *only* below the remnant of the higher sluice, we can quite reliably control the position of the SF in the small ellipsis cases, here on *where*. (30-ii) and (30-iii) set our baseline in that they show that, once *only* is overt, its associate must be, too, where both follow the sentential focus and are therefore deaccented, resulting in a rather small ellipsis.[12]

Note that the extraction in (30) is from the non-elided VP headed by *(only) tell*, so a proponent of the 'classical' MAXELIDE approaches could claim that (30-iii) is unsurprising, given that in non-extraction cases the choice of ellipsis size is free anyway. But our, stronger, claim is that in fact, the small ellipsis in (30-iii) is *only* possible if *where* is a SF. To bring home that point,

We suspect that the resulting configuration, then, involves two secondary foci, one on *had*, one on *when*, both embedded under the main focus *didn't*, which might simply be too hard to contextualize.

[12]This generalization—that (secondary) foci may be deleted only if the element they associate with is deleted, too (see also Han and Romero 2004: note 15, p.199, and Büring, 2015:note 23)—is actually derived by our account: the focus sensitive element 'retrieves' the focus, i.e. allows the alternatives to be reset.

we need to look at (30)'s minimal cousin (31).

(31) A: Bob will only tell ME where he's going.
B: (i) I wonder WHY.
(ii) I wonder WHY he will only tell you.
(iii) *I wonder WHY he will only tell you where.

First off, note that (31-ii) is acceptable in this context (unlike (30-ii) in (30)), because *only* associates with the overt (and by hypothesis SF marked) *you*. More importantly, (31-iii) is *un*acceptable here (contrasting with (30-iii)); its only reading is one in which *only* associates with *where* ('where, but not when', just as in (30)), which is infelicitous in this context as it differs from the main focus in A's utterance in (31). This is exactly the evidence we are after: for the ellipsis to be non-maximal, the final remnant *must* be (secondarily) focal.

Additionally, (30) and (31) between them show that previous accounts of MaxElide-like effects are incomplete in several ways. First, (31-iii) should not fall under the purview of MaxElide to begin with, since it does not involve extraction from the ellipsis site at all. On our account, *all* ellipsis is subject to MaxElide—or more precisely: marks the final remnant as focal—, so effects like in (30)/(31) are predicted.

Second, this effect cannot be due to competition between ellipses of different sizes. The only *eliptical* competitor in (30) and (31) alike is (B-i); (31-ii) (like (30-ii)) does not involve ellipsis, but null-complement anaphora.[13] But if (30-i) were to block (31-iii), it should do so in (30-iii) as well. On our account, (30-iii) is grammatical because *where* can be a SF in the context of (30), and (31-iii) is ungrammatical because it cannot—the SF must be *you*, not *where*. There is no competition involved.

In fact, (31-iii) actually *is* the maximal *ellipsis* structure that contains *only* and its associated focus *you*. So a version of MaxElide that claims (30-iii) to be grammatical because it is the biggest ellipsis excluding *where*—the focus associated with *only*—would have to do the same for (31-iii); the next bigger ellipsis, the VP headed by *tell*, would fatally include *you*.[14] Again, the present account has no trouble with this since it does not invoke competition (of ellipses or otherwise) at all.

[13]Even if you think that null-complement anaphora is a form of ellipsis, it still should not compete, lest (ia) be blocked by (ib).
(i) a. Mary has a hunch why Bob bailed, and Sue even KNOWS why.
b. Mary has a hunch why Bob bailed, and Sue even KNOWS.

[14]Throughout this Section we have ignored the option of deleting the middle VP, i.e. a reply like *I wonder why he will* in (30)/(31). These sound bad in either context, as predicted. However, it seems generally hard to elide a VP with initial *only*; that is to say, (i) seems to mean that Steve will tell me where he (Steve) is going, not that he will *only* tell me.
(i) Bob will only tell ME where he's going, and Steve will, too.
Whatever the reason for this, it may explain independently why *I wonder WHY he will* sounds odd in our examples, so we did not consider it in the paradigms.

7. Conclusion

In this paper we have put forward an account of the generalization that focus cannot be elided, without making reference to syntactic F-markers. We propose that ellipsis itself imposes a restriction on the available (non-trivial) focus alternatives: Whenever ellipsis applies, the elided part cannot introduce alternatives. In a second step we proposed that ellipsis imposes an additional focus restriction to the effect that the final remnant is itself focal. This was shown to capture so-called MAXELIDE effects, and to do so without using transderivational constraints.

Crucially our proposal assumes that this marking occurs with any instance of ellipsis, not just those containing traces. Cases in which smaller ellipses are permitted are analyzed as involving a secondary focus on the final remnant as well. This assumption was motivated using examples with associated second occurrence foci in non-maximal ellipses. On the resulting picture there is no MAXELIDE principle, rather the effect follows, like the ban on deleting focal material, from our general modelling of ellipsis, without F-markers or any kind of competition-based principles.

Appendix: Secondary Focus in UAS

A SF results when a focal constituent, though locally strong, ends up within a weak branch higher up in the tree. To make this more perspicuous we follow Rooth (1992) and mark the domain of a focus by a SQUIGGLE OPERATOR $\sim$ adjoined to that domain; some examples are given in (32).

(32)

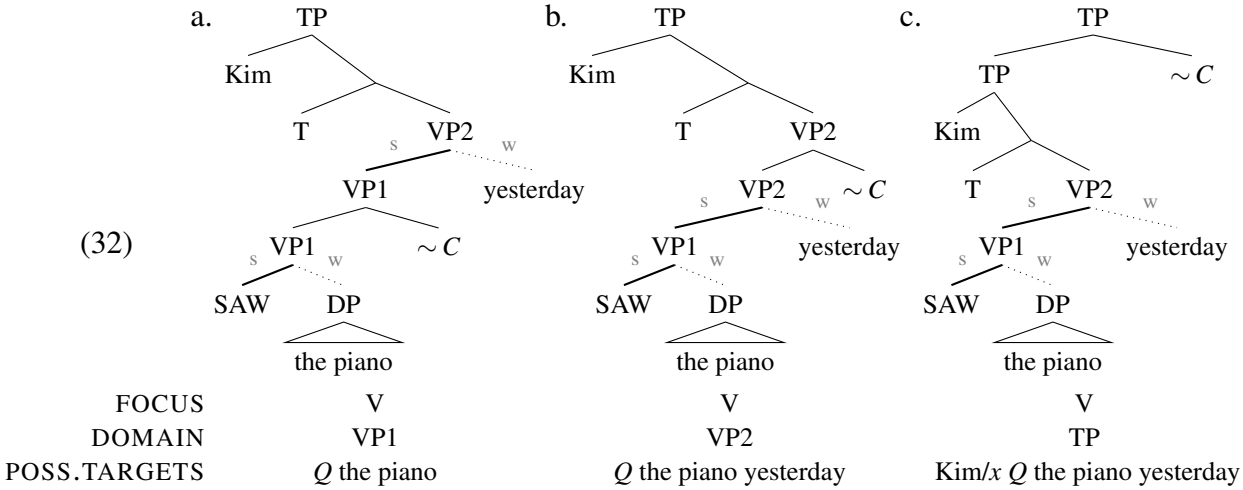

The squiggle operator RETRIEVES the focus, which—just like in Rooth (1992)—involves two things: First, it checks that the value of *C* (a covert pronoun, the focal target) is allowed as a focus alternative to the domain, i.e. compatible with the restrictions accumulated so far (otherwise the structure is undefined). If so, it, second, optionally RESETS the focus, i.e. sets the only possible alternative to the domain to be its literal meaning.

As detailed in Büring (2015), the resetting is crucial for the treatment of secondary foci. A typical SF configuration is SECOND OCCURRENCE FOCUS, where the domain of one focus

(the secondary) is included in the background of, and follows, the main, focus, as in (33a).[15]

(33) a. (The kids only skimmed the book.) Even **JOHN** only **skimmed** the book.

b. 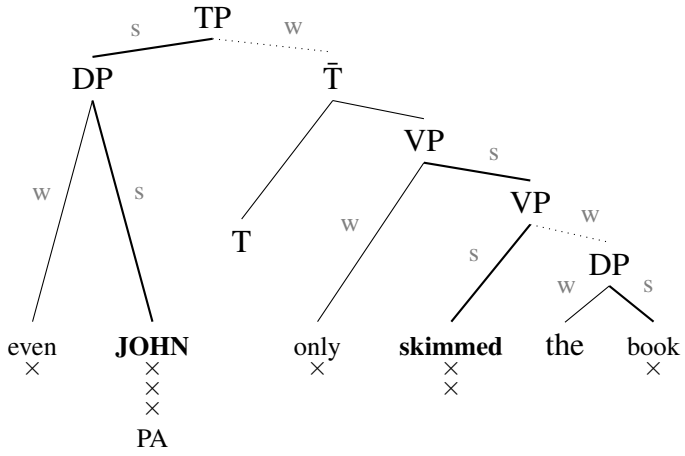

(33b) gives the full representation of the second clause, including a metrical grid compatible with the weights in the tree, and the resulting accent placement. *skimmed* is focal, as its mother node VP has undergone prosodic reversal, and the entire $\bar{T}$ is in the background of the focal *even John*, whose mother TP likewise is reversed.

The crucial generalization that follows in such a configuration is that the focus on the SF *skimmed* has to be retrieved *below* the higher focus; put differently, the domain of the focal *skimmed* may be at most as big as $\bar{T}$. As Büring (2015), following Büring (2015), shows, this generalization is born out by the facts around second occurrence focus.

The ellipsis configurations we argued for in this paper are structurally parallel to SOF, see for example (34): *does* is a SF, marked in this case not by prosodic reversal, but by ellipsis/ε, whose domain is in the background of the main focus on *Bill* (marked by prosodic reversal on TP).

[15]Every Second Occurrence Focus is a Secondary Focus, but not necessarily *vice versa*. The domain of a SOF, as the name suggests, needs to have a more or less *verbatim* antecedent, such as the first clause in (33a). A Secondary Focus is simply a focus whose domain is in the background of another focus. If there can be non-anaphoric focus domains, as we assume there can be, a SF need not have an antecedent.

(34)

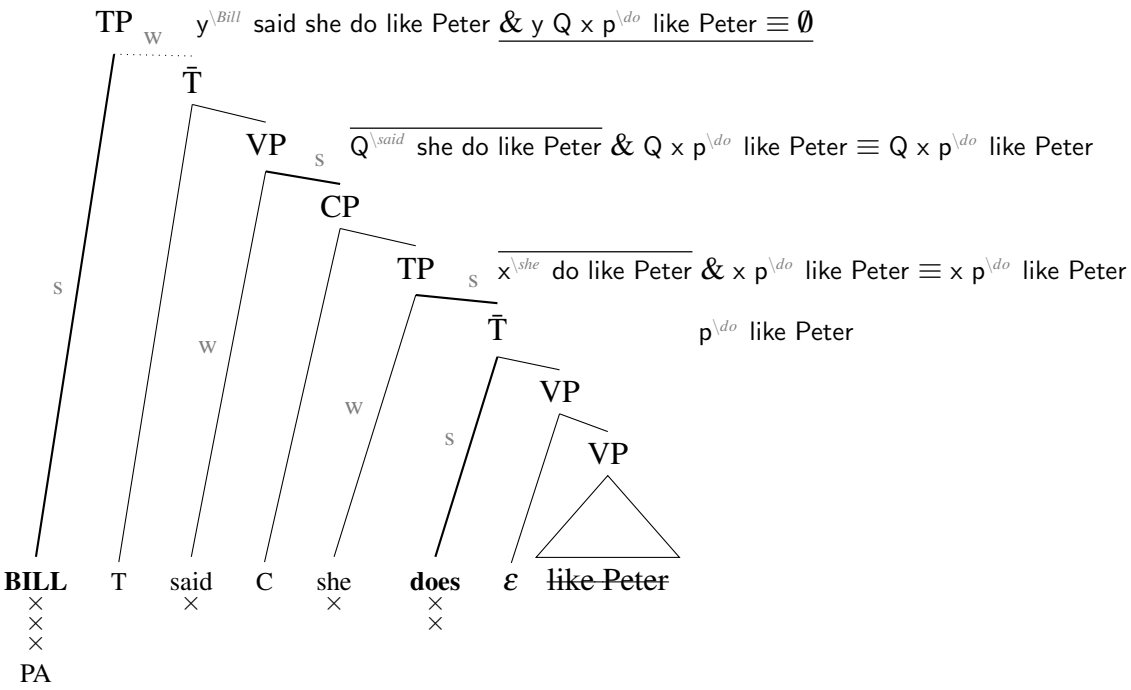

In (34) we also indicated the restrictions on focus alternatives imposed on the various nodes. Using those, we can now show that the restriction that SF needs to be retrieved and reset 'before' the main focus, indeed just follows from the general UAS system.

The lower $\bar{T}$ marks *does* as focal (by ellipsis), a restriction that is propagated up. All nodes above $\bar{T}$ and below the root have default weak–strong patterns and add rather trivial weak restrictions: they may or may not be focal. The only alternative to *do(es)*, we assume, is 'does not', so $\bar{T}$, TP and higher VP all require the focus alternatives to be about not liking Peter.

Fatally, the restriction originating with ε—that alternatives must involve *not* liking Peter—clashes with the strong restriction imposed by the prosodic reversal of the root TP, which requires 'said she likes Peter' to be background, i.e. be constant in any alternative. As indicated in the underlined part on the top line, these two are incompatible (one wants 'does' one wants 'doesn't').

In such a configuration the SF needs to be retrieved before the higher focus comes along. This is shown in (35), where ® marks the node at which the permitted focus alternatives are reset to the literal meaning.

(35)

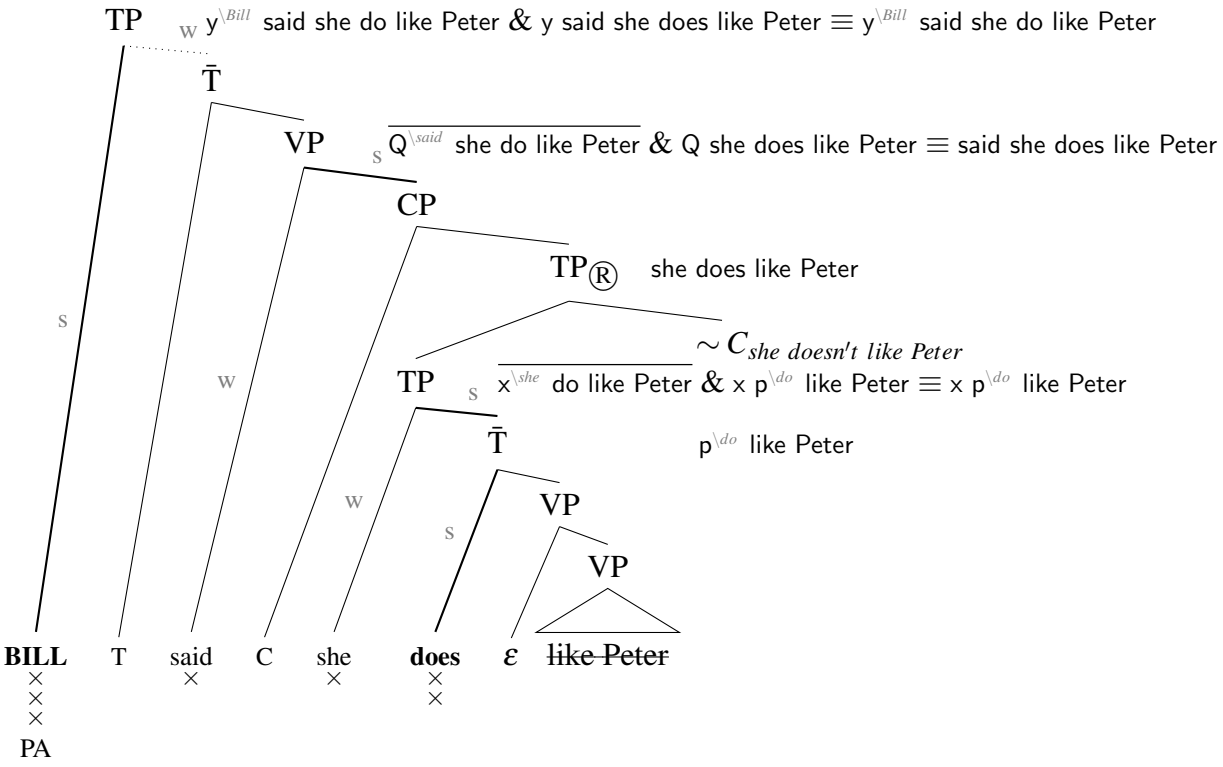

When the focus is retrieved by $\sim C$ (by juxtaposing 'she does like Peter' with the value of C 'she doesn't like Peter'), it is reset, so TP only has the trivial alternative. That in turn is propagated in the usual way until it meets the SR at the root level. This time the local strong restriction at the root and the propagated y said she does like Peter are compatible (in fact, the former simply subsumes the latter).

Finally, and crucially, a focus cannot be retrieved in a domain that contains unbound variables, (36a) and, as we just saw, it *has* to be retrieved before the 'next' focus. So if the binder to a variable in the focus domain *is* the next focus up, retrieval is impossible, and the structure crashes, (36b). As a consequence, there can be no SF in such configurations and maximal ellipsis up to the main focus is obligatory, (36c).

(36)

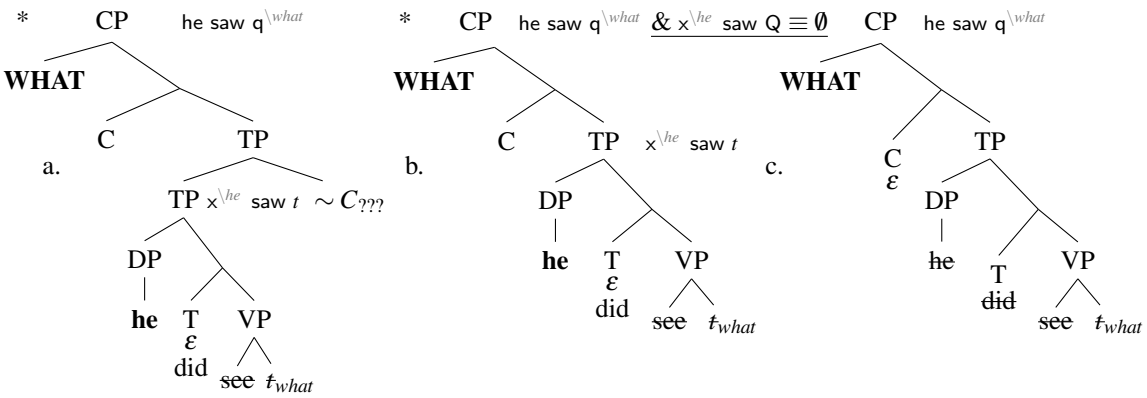

References

Büring, D. (2013/2015). A theory of second occurrence focus. *Language qs a Cognitive Process/Language, Cognition and Neuroscience 30*(1–2), 73–87.

Büring, D. (2015). Unalternative semantics. In S. D'Antonio, M. Moroney, and C. R. Little (Eds.), *Proceedings of SALT 25*, pp. 550–575. Linguistic Society of America.

Büring, D. (2016a). *A Beginner's Guide To Unalternative Semantics*. manuscript University of Vienna. http://homepage.univie.ac.at/daniel.buring/phpsite/index.php?page=UAS.

Büring, D. (2016b). Discontinuous foci and unalternative semantics. Volume 56 of *Current Trends in Generative Linguistics*, pp. 67–82. Ljubljana: Filozofska Fakulteta Ljubljana.

Davies, M. (2008). The corpus of contemporary American English (COCA): 560 million words, 1990–present. *Available online at https://corpus.byu.edu/coca/*.

Griffiths, J. (2017). Beyond MaxElide: A'-extraction from ellipsis sites. Paper presented at the *Generative Linguistics in the Old World (GLOW) 40*.

Han, C.-H. and M. Romero (2004). Disjunction, focus, and scope. *Linguistic Inquiry 35*(2), 179–217.

Heim, I. (1991). Artikel und Definitheit. Volume 6 of *Handbücher zur Sprach- und Kommunikationswissenschaft*, pp. 487–534. Berlin: Walter De Gruyter.

Merchant, J. (2001). *The Syntax of Silence: Sluicing, Islands, and the Theory of Ellipsis*. Oxford: Oxford University Press.

Merchant, J. (2004). Fragments and ellipsis. *Linguistics and Philosophy 27*(6), 661–738.

Merchant, J. (2008). Variable island repair under ellipsis. In K. Johnson (Ed.), *Topics in Ellipsis*, pp. 132–153. Cambridge: Cambridge University Press.

Messick, T. and G. Thoms (2016). Ellipsis, economy, and the (non) uniformity of traces. *Linguistic Inquiry 47*(2), 661–738.

Miller, P. and G. K. Pullum (2013). Exophoric VP ellipsis. In P. Hofmeister and E. Norcliffe (Eds.), *The Core and the Periphery: Data-driven Perspectives on Syntax Inspired by Ivan A. Sag*, pp. 5–32. Stanford CA: CSLI Publications.

Rooth, M. (1992). A theory of focus interpretation. *Natural Language Semantics 1*, 75–116.

Schuyler, T. (2001). *Wh*-movement out of the site of VP-ellipsis. Master's thesis, University of California, Santa Cruz.

Takahashi, S. and D. Fox (2005). MaxElide and the re-binding problem. In *Proceedings of SALT XV*.