EISS 13

Empirical Issues *in* Syntax *and* Semantics 13

Edited by

Christopher Pinon

and

Laurent Roussarie

2020
CSSP
Paris
http://www.cssp.cnrs.fr/eiss13/
ISSN 1769-7158

Contents

Preface

This is the thirteenth volume of the series *Empirical Issues in Syntax and Semantics* (EISS), which, like the preceding twelve volumes of the series, is closely related to the conference series *Colloque de Syntaxe et Sémantique à Paris* (CSSP). The six papers included in the present volume are based on presentations given at CSSP 2019, which took place on 02–04 October 2019 at UPS Pouchet CNRS, hosted by Université Paris 8.[1] CSSP 2019 had a small thematic session entitled *Computational semantics and pragmatics*, but since the number of papers from the thematic session submitted to the volume was low, they are not grouped separately.

We would like to take this opportunity to thank the reviewers, whose comments aided the authors in revising the first drafts of their papers, sometimes substantially. With their permission, the reviewers were (in alphabetical order by column):

Heather Burnett	Bert Cappelle	Alessandro Lenci
Patricia Cabredo Hofherr	Caroline Heycock	Philip Miller

Finally, we would also like to thank both the scientific committee and the organizing committee of CSSP 2019[2] for their efforts in planning and organizing a memorable conference.

Christopher Pinon
Laurent Roussarie
December 2020

[1] http://www.cssp.cnrs.fr/cssp2019/index_en.html
[2] http://www.cssp.cnrs.fr/cssp2019/contact/index_en.html

Agreement and information structure in Surati Gujarati

Kinjal Joshi

Abstract In Surati Gujarati and in Standard Gujarati, verbal elements can agree with a morphologically overt accusative case marked argument. This agreement pattern is unique because in other Indo-Aryan languages like Hindi and Marathi, presence of morphologically overt case on an argument blocks agreement with that argument, and the verb shows default agreement as a result. In this article, I show a novel empirical fact that in the causative constructions of Surati Gujarati, the verb can either agree with an accusative case marked object or show default agreement. I argue that this difference in agreement (agreement with accusative marked object vs. default agreement) is due to the use of two homophonous but different case markings on the direct object, namely, differentially marked, [+SPECIFIC] Accusative case vs. [+/−SPECIFIC] ordinary Dative case. To account for dative case assignment, I propose that dative case is a dependent (structural) case in Surati Gujarati. Further, I establish a relationship between information structure and agreement. To account for both case alternation and information structure-agreement relationship in Surati Gujarati, I propose a focus-driven object shift analysis. Under the Minimalist Program (MP) framework, focus movements are considered A-bar movements while movement related to case assignment is A movement. Here, the same movement serves both purposes. Thus, a movement of this type raises a larger theoretical question on the status of A vs. A-bar movement.

Keywords information structure · object shift · agreement · causative · case alternation

K. Joshi, University of Oslo, kinjal.joshi@iln.uio.no

In Christopher Pinon & Laurent Roussarie (eds.), Empirical Issues in Syntax and Semantics 13, 1–24. Paris: CSSP.
http://www.cssp.cnrs.fr/eiss13/

1 Introduction

Surati Gujarati is a dialect of Gujarati mainly spoken in the city of Surat and in areas surrounding Surat in India. The data in this article is from the

field work that I conducted in the city of Surat to collect data on Surati Gujarati. The field study was conducted on 25 language consultants, it consisted of 14 female and 11 male speakers. Their age groups ranged between 20–70 years.[1,2] Cardona (1965), Mistry (1998), and Suthar (2005) claim that in Standard Gujarati, when both the subject and the object are overtly case marked, the predicate must agree with the case marked object as seen in (1), and, if the predicate shows the default neuter agreement, it renders the sentence ungrammatical, as in (2). The agreement facts described in (1) and (2) for Standard Gujarati also hold for Surati Gujarati.[3]

(1) *ram-e* ***bakri-ne***$_{Acc}$ ***khaadhi*** *(Standard Gujarati)*
Ram.M.SG-ERG goat.F.SG-ACC eat.PFV.F.SG
'Ram ate the goat.'

(2) **ram-e* *bakri-ne*$_{Acc}$ *khaadhu* *(Standard Gujarati)*
Ram.M.SG-ERG goat.F.SG-ACC eat.PFV.N.SG
'Ram ate the goat.'

Now switching to the Surati Gujarati data, in examples (3) and (4), the use of a causativized predicate allows default agreement on the main verb in Surati Gujarati. This, I argue is due to case alternation. In a nutshell, I propose that Surati Gujarati has two distinct *-ne* markers, namely, -ne$_{Acc}$ vs. -ne$_{Dat}$. Although the case marker *-ne* is homophonous for Accusative and Dative case, the position of the accusative and dative case marked objects make the difference in the case markers explicit. I illustrate this in sentence (3), where the causativised predicate *dawdaav* ('cause to run') can show default neuter agreement morphology as Dative case does not control agreement in Surati Gujarati. A crucial point here is that the dative

[1]The Gujarati spoken in the cities of Vadodra and Ahemdabad of the Gujarat state in India is considered the Standard variety of Gujarati.

[2]I have excluded the field work methodology adopted to collect the data for reasons of space.

[3]The following abbreviations are used: ACC = Accusative, CAUS = Causative, DAT = Dative, ERG = Ergative, F = Feminine, *F* = Focus, M = Masculine, NOM = Nominative, NEG = Negation, N = Neuter, PLURAL = Plural, PFV = Perfective, SG = Singular

Note that nominative case in Hindi, Marathi, and Surati Gujarati is morphologically non-overt. Thus, referring to non-overt as zero case marking is merely a terminological choice and should not be taken as an analytical commitment.

marked object *ram* must always appear in the in-situ position to the right of the adverb *varamvaar* ('frequently').

(3) *Priyanka-e* ***lagbhag*** ***varamvaar*** $_{\text{VP}}$[*ram-ne*$_{Dat}$
Priyanka.F.SG-ERG probably frequently Ram.M.SG-DAT
vhel-lu *dawdaivu*]
early-N.SG run.CAUS.PFV.N.SG
'Priyanka probably frequently made Ram run early.'

Conversely, in sentence (4), the causativised predicate *dawdaav* ('cause to run'), can also agree with the direct object *ram* for masculine singular features. Note, that the accusative case marked object *ram* must always appear in a structurally higher position than the dative marked object in (3). The accusative object invariably appears to the left of the adverb *varamvaar* ('frequently') as seen in example (4).

(4) *Priyanka-e* ***lagbhag*** *ram-ne*$_{i\,Acc}$ $_{\text{VP}}$[***varamvaar***
Priyanka.F.SG-ERG probably Ram.M.SG-ACC frequently
t_i ***vhel-lo*** *dawdaivo*]
early-M.SG run.CAUS.PFV.M.SG
'Priyanka probably frequently made Ram run early.'

The main goal of the paper is to explore this puzzling nature of agreement with causative predicates, illustrated in examples (3) and (4). I describe the puzzle in detail in §2 of the paper. The paper is structured as follows; in §2, I provide a detailed description of the puzzle. In §3, I discuss object case in Surati Gujarati. §4 provides supporting evidence for case alternation in Surati Gujarati, namely, Differential Object Marking Accusative vs. Ordinary Dative case. In §5, I provide data and diagnostics to show that information structure plays a vital role in accusative case assignment in the language. In §6, I account for dative case assignment in Surati Gujarati. And in §7, I argue for an object shift type of analysis.

2 The puzzle

Looking at these agreement patterns in both Standard and Surati Gujarati the use of causative predicates with overtly case marked arguments, presents an intriguing puzzle. We obtain an unexpected behaviour of the agree-

ment patterns with causative predicates (both causativized transitive and intransitive predicates). The core puzzle here is that both object agreement and default neuter agreement on the predicate is possible with the causative predicates, as seen in sentences (5) and (6), respectively. Conversely, default agreement was not possible with an ordinary transitive predicate *kha* ('eat') as seen in (1) and (2). The *-ne*$_{Acc}$ on the direct object *ghoda* ('horse') in example (5), is a case of differential object marking (henceforth DOM) (Aissen 2003). I will argue that this agreement alternation on the causative predicates seen in examples (5) and (6) is a reflex of the DOM phenomenon (i.e. differential object *-ne*$_{Acc}$ vs ordinary *-ne*$_{Dat}$).

(5) ***priyanka-e*** ***ghoda-ne***$_{Acc}$ ***dawdaivo***
Priyanka.F.SG-ERG horse.M.SG-ACC run.CAUS.PFV.M.SG
'Priyanka made the horse run.'

(6) *priyanka-e* *ghoda-ne*$_{Dat}$ *dawdaivu*
Priyanka.F.SG-ERG horse.M.SG-DAT run.CAUS.PFV.N.SG
'Priyanka made a/the horse run.'

For the sake of simplicity, I use a causativized unergative predicate *dawdaav* ('cause to run') to demonstrate the puzzle in sentences (5) and (6).[4] It must be noted that the *-ne*$_{Acc}$ object in (5), is definite whereas the *-ne*$_{Dat}$ marked object in sentence (6) is ambiguous between definite and indefinite. Further, the agreement pattern seen in examples (5) and (6) is also available for +HUMAN proper names. I illustrate this in (7) and (8). In example (7), the causative predicate *dawdaav* ('cause to run'), agrees with

[4] Similar patterns are observed with causativized transitives, as in (i) and (ii) with an ambiguity in (i).

(i) *ram-e* ***bakri-ne***$_{Acc}$ *khawdaavi*
Ram.M.SG-ERG goat.F.SG-ACC eat.CAUS.PFV.F.SG
Reading A: 'Ram made the (specific) goat eat.'
Reading B: 'The (specific) goat was fed by Ram to someone.'
(Free Translation: 'Ram fed the goat')

(ii) *ram-e* *bakri-ne*$_{Dat}$ *khawdaivu*
Ram.M.SG-ERG goat.F.SG-DAT eat.CAUS.PFV.N.SG
Reading A (Only): 'Ram made a/the goat eat.'

the *-ne*$_{Acc}$ marked direct object *ram*.

(7) *priyanka-e* ***ram-ne***$_{Acc}$ ***dawdaivo***
Priyanka.F.SG-ERG ram.M.SG-ACC run.CAUS.PFV.M.SG
'Priyanka made Ram run.'

Conversely, in example (8), the causative predicate *dawdaav* ('cause to run'), shows default neuter singular agreement.

(8) *priyanka-e* *ram-ne*$_{Dat}$ *dawdaivu*
Priyanka.F.SG-ERG ram.M.SG-DAT run.CAUS.PFV.N.SG
'Priyanka made Ram run.'

My proposal for solving this puzzle is to argue that this distinct agreement pattern is due to alternation in case, as indicated by the subscripted 'Acc' and 'Dat', thereby proposing two distinct *-ne* forms, namely, differential object *-ne*$_{Acc}$ vs. ordinary *-ne*$_{Dat}$. The phenomenon of case alternation has been attested for many languages like Russian, Icelandic, Spanish and Dutch (see Demonte 1995; Pineda 2013; Sigurðsson & Wood 2012).

DOM in Surati Gujarati is only available for animates. It is obligatory for [+HUMAN] proper names but the DOM marker can be dropped with a [−HUMAN, +ANIMATE] direct object, in which case they receive an indefinite interpretation as seen in example (9). The direct object *ghodo* ('horse') in (9) receives an indefinite interpretation without the DOM marker and the verb in this case must obligatorily agree with the direct object as seen in (9). This is in sharp contrast to example (5) where the object marked with the overt DOM marker *-ne*$_{Acc}$, receives a definite interpretation.

(9) *priyanka-e* ***ghodo*** ***dawdaivo***/**dawdaivu*
Priyanka.F.SG-ERG horse.M.SG run.CAUS.PFV.M.SG/*N.SG
'Priyanka made a horse run.'

In the following section, I discuss object case in greater detail.

3 Basics of object case in Surati Gujarati

In this section, I discuss the properties and the impact of the *-ne* marker, a homophonous marker for accusative and dative case on the interpretation of the sentence (for Standard Gujarati, see Mistry 1998). In Table 1, I

Table 1 Licensing conditions for DOM in Surati Gujarati

	DIRECT OBJECT	DOM (*-ne*$_{Acc}$)	Example
I.	[+HUMAN]/Proper names	Obligatory	(10)
II.	[+ANIMATE,−HUMAN]	Optional	(11)
III.	[−ANIMATE]	Not licensed	(12)

summarize the licensing conditions of DOM in Surati Gujarati.

The overt accusative case marking on direct objects in Surati Gujarati is sensitive to animacy and definiteness. Further, the overt case marker *-ne* is obligatory for direct objects that are proper names as seen in (10). The case marker *-ne* is optional for objects with [+ANIMATE] features, as in (11).

(10) *raj-e* ***ram*(-ne***$_{Acc}$***)*** ***joyo***
Raj.M.SG-ERG Ram.M.SG-ACC see.PFV.M.SG
'Raj saw Ram.'

(11) *raj-e* ***bakri(-ne***$_{Acc}$***)*** ***joyi***
Raj.M.SG-ERG goat.F.SG-ACC see.PFV.F.SG
'Raj saw a (the) goat.'

[−ANIMATE] objects cannot license the overt case marker in Surati Gujarati as it is sensitive to animacy, as seen in sentence (12), where the inanimate direct object *gaadi* (car) cannot be used with the *-ne* marker. [5]

(12) *raj-e* ***gaadi****(*-ne*$_{Acc}$*)* ***joyi***
Raj.M.SG-ERG car.F.SG(*-ACC) see.PFV.F.SG
'Raj saw a (*the) car.'

The hierarchy for licensing an overt accusative case marker on the object is illustrated in (13).

(13) **Animacy scale**

[5]The rules for licensing of an overt case marker on the object seem to be different for Standard Gujarati. Compare Mistry (1998:429) where he claims that even an inanimate definite object like *kaagal-ne* ('the letter') can license an overt accusative case marker.

Human > Animate Definite > Animate Indefinite > Inanimate

DOM is optional on [+ANIMATE] direct objects and its absence yields an indefinite interpretation of these direct objects as seen in (14), where *ghodo* ('horse') has an indefinite reading. In (14) both the adverb *vhel-* ('early') and the predicate *joyo* ('see') agree with the direct object *ghodo* ('horse') for masculine and singular features.

(14) *priyanka-e* *(varamvaar)* ***ghodo*** *(??varamvaar)*
Priyanka.F.SG-ERG frequently horse.M.SG frequently
vhel-lo ***joyo***
early-M.SG see.PFV.M.SG
'Priyanka frequently saw a horse early.'

The presence of DOM imparts a definite interpretation to the object as seen in (15). The position of the DOM marked object *ghoda-ne*$_{Acc}$ ('horse') in (15) must be noted as it is higher than the indefinite *ghodo* ('horse') in (14). I use adverbs *varamvaar* ('frequently') and *vhel-* ('early') to indicate the difference in the positions of the two objects in (14) and (15).

(15) *priyanka-e* *(??varamvaar)* ***ghoda-ne***$_{Acc}$ *(varamvaar)*
Priyanka.F.SG-ERG frequently horse.M.SG-ACC frequently
vhel-lo ***joyo***
early-M.SG see.PFV.M.SG
'Priyanka frequently saw the horse early.'

By contrast, with ditransitives the *-ne*$_{Dat}$ marked indirect object *ghodo* ('horse') in (16) does not trigger agreement and can never do so. The adverb *vhel-* ('early') and the predicate *aap* ('give') both agree with the unmarked object *chana* ('chickpeas') for number. Another important observation is the in-situ position of the *-ne*$_{Dat}$ marked object *ghodo* ('horse'), which is similar to the position of the unmarked object in (14). Note also that the indirect object is ambiguous between definite and indefinite in (16), parallel to what we saw in (6).

(16) *priyanka-e* *(varamvaar)* *ghoda-ne*$_{Dat}$ *(??varamvaar)*
Priyanka.F.SG-ERG frequently horse.M.SG-DAT frequently

Table 2 Evidence for Case Alternation

	Argument	Example
I.	The use of -ne_{Acc} is optional but -ne_{Dat} is obligatory	(18), (20), (22)–(23)
II.	-ne_{Acc} object is structurally higher than -ne_{Dat} object	(24)–(26)
III	Agreement with Negative particle blocked with -ne_{Dat} object	(30)–(31)

chana ***vhel-la*** ***aapya***
chickpeas.PL early-PL give.PFV.PL
'Priyanka frequently gave chickpeas to a/the horse early.'

4 Evidence for case alternation in Surati Gujarati

As already indicated, I propose the following explanation for the distinct agreement pattern in Surati Gujarati: when the case on the object is accusative case, the verb must agree with the case marked object, and when it is dative, the verb exhibits default neuter agreement (see Mistry 1998 for Standard Gujarati). In this section, I will provide arguments to make the difference between differential object -ne_{Acc} vs. ordinary -ne_{Dat} more explicit.

4.1 Diagnostics for case alternation

The hypothesis must be empirically substantiated, as the two *-ne* morphemes are homophonous and this difference is visible in syntax (also see Mistry 1998). I summarize the arguments I provide in this section in Table 2.

A piece of evidence to show the presence of accusative-dative case alternation is seen in sentences (18) and (20) for contexts (17) and (19) respectively (see §5 on the role of information structure in case marking); here, the -ne_{Acc} marker is the accusative case marker (DOM) and it can be dropped. The predicate *dawdaav* ('cause to run') and the adverb *vhel-*

('early') in both these examples agree with the object.

(17) **Context:** Priyanka is a shepherd and has three horses Y1, Y2, Y3. Frequently, Y1 entered her kitchen early.

(18) *Priyanka-e* ***ghoda-ne***$_{Acc}$ *varamvaar* ***vhel-lo***
Priyanka.F.SG-ERG horse.M.SG-ACC frequently early-M.SG
dawdaivo
run.CAUS.PFV.M.SG
'Priyanka frequently made the horse run early.'

(19) **Context:** Some horse or other from Priyanka's village frequently entered her house early.

(20) *Priyanka-e* *varamvaar* ***ghodo*** ***vhel-lo***
Priyanka.F.SG-ERG frequently horse.M.SG early-M.SG
dawdaivo
run.CAUS.PFV.M.SG
'Priyanka frequently made a horse run early.'

By contrast, the *-ne* case marker is always obligatory when there is default agreement as seen in sentences (22) and (23) for context (21). I argue that this follows from analyzing it as a dative (non-DOM) case marker here.

(21) **Context:** Priyanka's mother asked: what did Priyanka do?

(22) *Priyanka-e* *varamvaar ghoda-ne*$_{Dat}$ *vhel-lu*
Priyanka.F.SG-ERG frequently horse.M.SG-DAT early-N.SG
dawdaivu
run.CAUS.PFV.N.SG
'Priyanka frequently made a/the horse run early.'

(23) **Priyanka-e* *varmavaar ghodo* *vhel-lu*
Priyanka.F.SG-ERG frequently horse.M.SG early-N.SG
dawdaivu
run.CAUS.PFV.N.SG
'Priyanka frequently made a/the horse run early.'

It must be noted that (18), (20) and (22) are all grammatical in their respective contexts. However, there is no context in which (23) is grammat-

ical. Thus, it seems that the distribution of accusative and dative case on the object is contextually determined. Another piece of evidence for case alternation in Surati Gujarati is presented by double object constructions. It seems very clear from Cardona & Suthar (2003) that SUBJ-IO-DO-VERB can be considered the base word order for double object constructions in Standard Gujarati as seen in (24) and this also holds for Surati Gujarati. The double object constructions are essential as dative case is considered the case of the indirect objects (see Dryer 1986; Haspelmath 2005) as shown in (24). In (24), the beneficiary is unambiguously dative marked. And the verb never agrees with the beneficiary. A causativized transitive will also end up looking, as in (25).[6,7]

(24) *ram-e* *raj-ne*$_{Dat}$ ***bakri*** ***aapi***
Ram.M.SG-ERG Raj.M.SG-DAT goat.F.SG give.PFV.F.SG
'Ram gave Raj a goat.'

(25) *ram-e* *varamvaar raj-ne*$_{Dat}$ VP[***bakri(*-ne***$_{Acc}$***)***
Ram.M.SG-ERG frequently Raj.M.SG-DAT goat.F.SG(*-ACC)
vhel-li ***khawdavi***]
early-F.SG eat.CAUS.PFV.F.SG
'Ram frequently fed (a/*the) goat to Raj early.'

Crucially, the language cannot license accusative case on the direct object in sentence (25) as long as it is in the in-situ position. But when the direct object *bakri* ('goat') moves out of the VP then it can be *-ne* marked, as in (26). Comparing (25) with (26), I show that that the accusative case marked causee moves out of the VP domain and it is structurally higher as seen in (26) than the dative argument, as in example (25).

(26) *ram-e* ***bakri-ne***$_{i\,Acc}$ *varamvaar raj-ne*$_{Dat}$ VP[$\boldsymbol{t}_i$
Ram.M.SG-ERG goat.F.SG-ACC frequently Raj.M.SG-DAT

[6]The presence of dative subjects and dative objects in Standard Gujarati has been documented in previous studies (see Mistry 2004; 1998).

[7]The important point that the paper tries to make is to illustrate that Surati Gujarati has two distinct *-ne* markers in causative constructions. Whether these markers are truly dative or accusative, as discussed by Manzini & Franco (2016), is beyond the scope of the paper.

vhel-li khawdavi]
early-F.SG eat.CAUS.PFV.F.SG
'Ram frequently fed the goat to Raj early.'

A dative-marked argument can only occur to the right of the adverb *varamvaar* ('frequently') as seen in sentences (22), (25) and (26), whereas the direct object can only be accusative-marked if it occurs to the left of *varamvaar* ('frequently') as seen in (18) and (26) as opposed to (25). Another piece of evidence comes from negation in Surati Gujarati. Surati Gujarati has negated auxiliaries which can either agree with the subject or with the object if both are available i.e., if the subject is unmarked for case as seen in (27). Here, the main verb *laav-* ('bring') agrees with the subject *Ram* and the negation can either agree with the subject *Ram* or with the object *gaadi* ('car') (For further discussion, see Joshi 2019).

(27) *ram gaadi nho-ti/nho-to lav-yo*
Ram.M.SG car.F.SG NEG-F.SG/M.SG bring-PFV.M.SG
'Ram did not bring a car.'

If the subject is unavailable, i.e. if the subject is ergative case marked, both the main verb and the pre-verbal *nho-* agree with the object as seen in (28) and (29).

(28) *ram-e* ***gaadi nho-ti chalaavi***
Ram.M.SG-ERG car.F.SG NEG-F.SG drive.PFV.F.SG
'Ram had not driven a car.'

(29) *ram-e* ***bakri-ne***$_{Acc}$ ***nho-ti khaa-dhi***
Ram.M.SG-ERG goat.F.SG-ACC NEG-F.SG eat-PFV.F.SG
'Ram had not eaten the goat.'

The observation crucial to the argument is that negation can agree independently of the main verb as seen in example (27), where the negation can either agree with the subject or with the object. Based on the above paradigm, I predict that if both subject and object are unavailable for agreement, then the main verb and *nho-* both must show default agreement. This is exactly what we observe in example (30) where both the main verb *dawdaav* ('made to run') and the negation show default agree-

ment morphology. The sentence in (30) supports an analysis where *-ne* is a dative marker when the verb exhibits default agreement since, whenever the dative object is unavailable to the main verb for agreement, it also remains unavailable to the negation for agreement.

(30) *Priyanka-e ghoda-ne$_{Dat}$ nho-tu dawdaivu*
Priyanka.F.SG-ERG horse.M.SG-DAT NEG-N.SG run.CAUS.PFV.N.SG
'Priyanka did not make a/the horse run.'

If the object in example (30) was accusative, i.e., available for agreement then it should have allowed the sentence in (31) to be grammatical, considering the empirical fact illustrated in (27) that negation can agree independently of the main verb. However, example (31) is ungrammatical.

(31) **Priyanka-e ghoda-ne$_{Dat}$ nho-to dawdaivu*
Priyanka.F.SG-ERG horse.M.SG-DAT NEG-M.SG run.CAUS.PFV.N.SG
'Priyanka did not make a/the horse run.'

As a result, it confirms the analysis that the object in (30) with default agreement on the main verb is dative and not accusative. Thus, based on the empirical evidence presented above I claim that the difference in the agreement patterns of Surati Gujarati is due to accusative-dative case alternation.

5 Role of information structure in case assignment

In this section, I argue that information structure plays a vital role in case alternation. To explain the role of information structure in case assignment, I propose the following: In the Surati Gujarati examples under discussion, the case-marked objects of the causative predicates are marked accusative case if the focus is narrow focus on the object. By contrast, the object is marked with dative case if the focus is focus on the entire VP.

5.1 Diagnostics for focus

To test the above hypothesis, I use the question-answer congruence test following Hamblin (1973), as the main example of pragmatic focus emerges in question-answer congruence where a question indicates the communicative goal of the questioner. In context (32), the focus is on the entire VP. Thus, the direct object *ghodo* ('horse'), is licensed with dative case, as

seen in sentence (33). Here, the presence of dative case is evident, as the verb shows default neuter agreement and does not agree with the direct object; recall that I have argued that dative case does not control agreement in Surati Gujarati (See §1, example (3)).

Focus on the VP

(32) **Context**: Priyanka wanted to know the strength of her new horse. When her mother saw the horse gasping for breath she asked
Q1: What did Priyanka do?

(33) *Priyanka-e* *[ghoda-ne$_{Dat}$* *dawdaivu/*vo]$_F$*
Priyanka.F.SG-ERG horse.M.SG-DAT run.CAUS.PFV.N.SG/*M.SG
'Priyanka made a/the horse run.'

In context (34), the focus is on the direct object. As a result, it gets the accusative case, as seen in sentence (35). The presence of accusative case in this sentence is evident, since the verb in (35) agrees with the direct object, as accusative case is transparent to agreement in Surati Gujarati (see §2–§4).

Narrow focus (Focus on the direct object)

(34) **Context:** Priyanka had a goat and a horse. She was unsure whom she would send to the race. So Salman asked her mother
Q2: Whom did Priyanka pick to run in the race?

(35) *Priyanka-e* ***[ghoda-ne$_{Acc}$]$_F$*** ***dawdaivo/*vu***
Priyanka.F.SG-ERG horse.M.SG-ACC run.CAUS.PFV.M.SG/*N.SG
'Priyanka made the horse run.'

In following section, I discuss dative case assignment in causative structures of Surati Gujarati.

6 Dative case assignment in causative structures

The empirical facts presented in the paper point us to the following direction with respect to dative case assignment. In what follows, I propose that dative case in Surati Gujarati causative structures is assigned by a dependent case mechanism. I build on Baker (2015), who argues that dative case is a high dependent case in VP and a VP analog of ergative case. We

know from sentences (25)–(26) that dative objects occur to the right of the adverb *varamvaar* ('frequently'), which is exactly what we get in sentence (36), thus, it is safe to say that the object *vaagh-ne* ('tiger') in sentence (36) is dative. The dative in (36) seems to be structural case as the object *vaagh* ('tiger') has an agent theta role from the verb *kha-* ('eat'). Thus, *vaagh* ('tiger') does not have the theta role associated with dative case (e.g. recipient). The dependent case idea is of utility here as the causee of a transitive causative predicate like *kha* ('eat') is dative if and only if it c-commands a distinct NP in the same VP domain. Notice that the object *vaagh-ne*$_{Dat}$ ('tiger') and *bakri* ('goat') in (36) are in the same VP domain but the dative object c-commands the unmarked object *bakri* ('goat') and thus, dative case is a result of case competition between the two DPs in the same domain.

(36) *mhe (??vaagh-ne*$_{Dat}$*) (varamvaar)* $_{\text{VP}}$[*vaagh-ne*$_{Dat}$ *bakri*
I.ERG tiger.M.SG-DAT frequently tiger.M.SG-DAT goat.F.SG
khawdaavi]
eat.CAUS.PFV.F.SG
'I frequently fed a/the tiger a goat.'

Now let us extend this analysis to a causativised unergative predicate like *dod* ('run'). For this, I Follow Laka (1993) and Baker & Bobaljik (2017), via Bárány & Sheehan (2019), who propose that unergative predicates as opposed to unaccusatives are said to be underlyingly transitive with a null cognate object. I further propose that causativised unergative predicates like *dod* ('run') contain an implicit object in the VP domain, which acts as the case competitor for the causee. As a result the causee gets dative as a dependent case, as seen in (37).

(37) *priyanka-e (??ghoda-ne*$_{Dat}$*) varamvaar* $_{\text{VP}}$[*ghoda-ne*$_{Dat}$
Priyanka.F.SG-ERG horse.M.SG-DAT frequently horse.M.SG-DAT
∅ *dawdaivu*]
run.CAUS.PFV.N.SG
'Priyanka made a/the horse run.'

In line with the above-mentioned generalization, causativized unaccusative predicates do not have a null cognate object, and thus, do not

allow default agreement. I illustrate this in (38), where the causative predicate *paad* ('made to fall') has to obligatorily agree with the causee *ghoda* ('horse').

(38) *priyanka-e* *ghoda-ne*$_{Acc}$ *padaav-yo/*yu*
Priyanka.F.SG-ERG horse.M.SG-ACC fall.CAUS.PFV.M.SG/*N.SG
'Priyanka made the horse fall.'

Recall, I show in sentences (25) and (26) that the accusative case marked causee moves out of the VP domain and it is structurally higher than the dative argument. I discuss the mechanism of object movement in the following section in greater details and provide evidence for the same.

7 Object shift/scrambling in Surati Gujarati

The distinct agreement patterns and case alternation seem to be a result of object shift/scrambling. In this section, I will illustrate the difference in the syntactic positions of DOM marked object vs. ordinary Dative object with respect to adverbs, as seen in examples (40) and (42). I propose that in the causative examples that I focus on, an object with +FOCUS feature moves out of the VP to a position where it gets accusative case, whereas an object without the +FOCUS feature remains in-situ regardless of its specificity and is assigned dative case. The crucial data supporting the claim comes from the adverb placement test (Pollock 1989). Assuming that adverbs have fixed positions, the results of the adverb placement test have a direct implication for the syntactic analysis of the word order. I use the adverbs *lagbhag* ('probably'), *varamvaar* ('frequently'), and *vhel-* ('early') to test object movement. For the [+FOCUS] object to get accusative case it must move out of the VP, as in (40), where the direct object *Ram* occurs to the left of the adverb *varamvaar* ('frequently').

(39) **Context:** Priyanka had many servants of which only Ram was the one who looked very exhausted. Thus, Priyanka's mother asked her father the following question:
Q3: Whom did Priyanka frequently make run early?

(40) *Priyanka-e* ***lagbhag*** *ram-ne$_{i\,Acc}$* $_{VP}$[***varamvaar*** ***t$_i$***
Priyanka.F.SG-ERG probably Ram.M.SG-ACC frequently
vhel-lo *dawdaivo*]
early-M.SG run.CAUS.PFV.M.SG
'Priyanka probably frequently made Ram run early.'

The sentence is grammatically deviant if the accusative case is forced on the object in-situ for context (39), as in sentence (41). Since proper names that have a human referent must be case marked with the DOM marker, *-ne* cannot be omitted from *Ram*.

(41) ??*Priyanka-e* ***lagbhag*** ***varamvaar*** $_{VP}$[***(vhel-lo)***
Priyanka.F.SG-ERG probably frequently early-M.SG
ram-ne$_{Acc}$ ***(vhel-lo)*** *dawdaivo*]
ram.M.SG-ACC early-M.SG run.CAUS.PFV.M.SG
'Priyanka probably frequently made ram run early.'

Now let us consider a different question within the same context, as in (39), described in Q4. The response to the question in Q4 (with wide VP focus) is given in (42). If the object is dative case marked with default neuter agreement on the verb, the object without the [+FOCUS] feature remains in-situ, to the right of the adverb *varamvaar* ('frequently'), as seen in (42).

Q4: What did Priyanka do?

(42) *Priyanka-e* ***lagbhag*** ***varamvaar*** $_{VP}$[*ram-ne$_{Dat}$*
Priyanka.F.SG-ERG probably frequently Ram.M.SG-DAT
vhel-lu *dawdaivu*]
early-N.SG run.CAUS.PFV.N.SG
'Priyanka probably frequently made Ram run early.'

The sentence is judged to not be perfectly grammatical if the dative case marked object *Ram*, moves higher on the clausal spine out of the VP, as seen in sentence (43) which contrasts with (42).

(43) ?*Priyanka-e* ***lagbhag*** *ram-ne$_{i\,Dat}$* ***varamvaar*** $_{VP}$[***t$_i$***
Priyanka.F.SG-ERG probably Ram.M.SG-DAT frequently

vhel-lu *dawdaivu*]
early-N.SG run.CAUS.PFV.N.SG
'Priyanka probably frequently made Ram run early.'

To account for object shift (sentence (40)) I propose the following (illustrated in (44)): First, the ergative subject, moves to the specifier of TP in this case. Subject movement is evident from the position of the adverb *lagbhag* ('probably') in the structure as shown in (44). Second, the accusative case marked object moves out of the VP. Movement of the DP out of the VP has been previously argued by Diesing (1992) to target specific objects.

For now I assume that the accusative object originates in the complement of V and the dative object in [Spec, VP] (as shown in (46)). Third, the movement of the object out of the VP is connected to accusative case in causatives. The evidence from the adverb placement test clearly show us that the object has to be higher on the clausal spine for it to surface with accusative case. Fourth, only objects with +FOCUS feature move out of the VP and get accusative case, as seen in sentence (40). If objects with +FOCUS feature do not move out of the VP, it renders the sentence ungrammatical as seen in (41). I build on Diesing (1997) and claim that objects with +FOCUS feature escape the existential closure when they move out of the VP, thereby acquiring accusative case as a form of dependent case (Baker 2015). This is sketched in (44) for sentence (40), and in (45) for (41).[8]

To account for in-situ objects (see (42)), I propose the following, illustrated in (46): First, the ergative subject moves to the specifier of TP as previously argued. Second, the dative case marked object remains in-situ and does not move out like the accusative case marked object. Third, the in-situ position of the dative case marked object is confirmed by the adverb placement test in (42) and (43). The sentence is perfectly grammatical when the dative marked direct object remains in-situ as seen in (42). However, the sentence cannot be judged perfectly grammatical when the dative marked direct object moves out of the VP, as in (43). Fourth, only objects with +FOCUS feature can move out of the VP. Since, the object in (42)

[8]This movement of the object is plausibly to the specifier of the focus position as described by Jayaseelan (2008). Movement of the object to the specifier of the focus position is made explicit in the trees in (44)–(47) for better clarity.

(44) √

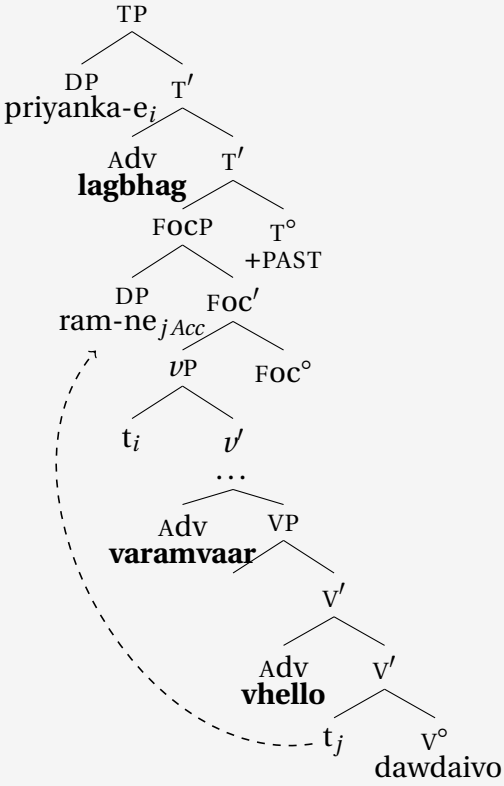
TP
DP
priyanka-e$_i$
T′
Adv
lagbhag
T′
FocP
T°
+PAST
DP
ram-ne$_j$ *Acc*
Foc′
*v*P
Foc°
t$_i$
v′
…
Adv
varamvaar
VP
V′
Adv
vhello
V′
t$_j$
V°
dawdaivo

(45) *

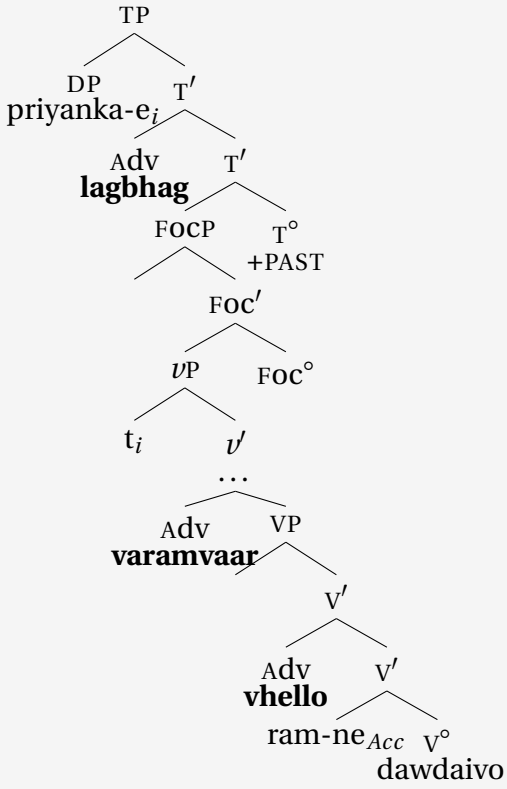
TP
DP
priyanka-e$_i$
T′
Adv
lagbhag
T′
FocP
T°
+PAST
Foc′
vP
Foc°
t$_i$
v′
…
Adv
varamvaar
VP
V′
Adv
vhello
V′
ram-ne$_{Acc}$
V°
dawdaivo

(46) √

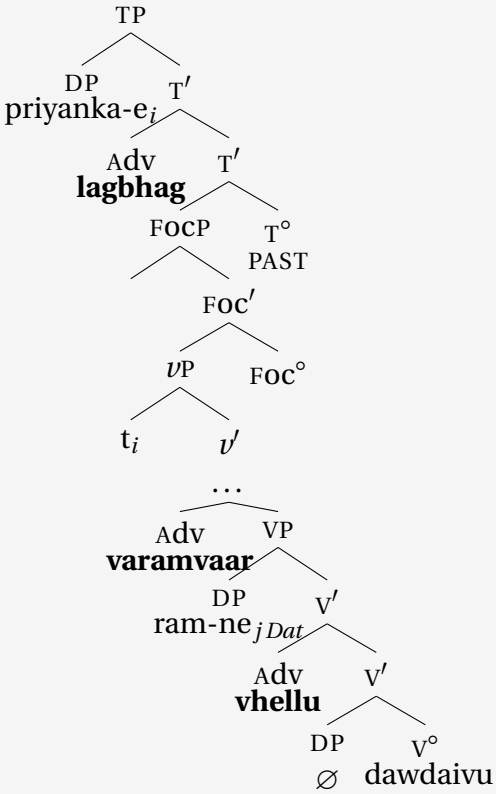
TP
DP
priyanka-e$_i$
T′
Adv
lagbhag
T′
FocP
T°
PAST
Foc′
*v*P
Foc°
t$_i$
v′
…
Adv
varamvaar
VP
DP
ram-ne$_j$ *Dat*
V′
Adv
vhellu
V′
DP
∅
V°
dawdaivu

(47) *

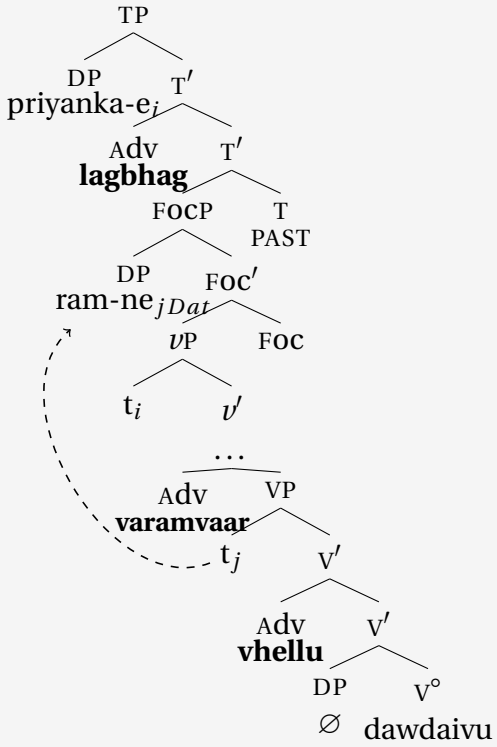
TP
DP
T′
priyanka-e$_i$
Adv
T′
lagbhag
FocP
T
PAST
DP
Foc′
ram-ne$_j$ *Dat*
*v*P
Foc
t$_i$
v′
…
Adv
VP
varamvaar
t$_j$
V′
Adv
V′
vhellu
DP
V°
∅
dawdaivu

does not have the +FOCUS feature due to focus on the entire VP, it remains in-situ and gets bound by existential closure inside the VP. The fact that the object in (42) is not the focused element of the sentence was illustrated in both (33) and (42). The structural analysis for sentence (42) is sketched in (46), and for sentence (43) in (47).[9] The existentially bound causee, which is in c-command relation with the null cognate object gets the dependent dative case. This is mainly because the causee is in a case competition scenario with the null cognate object in the same domain; dative case assignment within the VP is briefly mentioned in Baker (2015:133). Note however that Baker (2015) does not discuss dative case assignment with respect to information structure.[10]

8 Conclusion

In this paper, I have presented novel empirical evidence. I show that the difference in the agreement patterns in Surati Gujarati causative predicates, seems to be due to alternation of case on the objects. The next task was to define the conditions under which we observe case alternation. While pursuing this task, I determined the information structure and agreement relationship in Surati Gujarati. I showed that it is the presence or absence of a +FOCUS feature on the object, which seems to be the necessary condition for case alternation. Further, I propose that the dative case is assigned along the lines of Baker's (2015) dependent case proposal. To account for all of the above generalizations, I proposed an object shift analysis that derives both case alternation and the information structure-agreement relationship in Surati Gujarati. The requirement for object movement out of the VP is the +FOCUS feature, as objects without the +FOCUS feature remain in-situ. Such focus-driven object movement is the prerequisite for accusative case assignment. This idea is supported by the empirical data in the paper, in particular, (40) and (42). The empirical facts and the analysis in this article suggest that object movement for fo-

[9] Sentence (42) indicates that the Diesing effect, which holds for Hindi-Urdu as shown by Bhatt & Anagnostopoulou (1996), does not always hold for Surati Gujarati, as the dative case marked direct object in sentence (33) is ambiguous for specificity/definiteness object, but still does not move out of the VP.

[10] I remain agnostic about the status of the interaction of FocP with VP, when the focus is on the entire VP.

cus feeds differential object marking. This focus-driven object movement raises a bigger question on the status of A vs A-bar movement, as under the Minimalist Program (MP) framework, movement has been viewed as a concomitant of the operation 'Agree', (see Chomsky 2000) and focus movements are considered A-bar movements while movement related to case assignment is A movement. A-movement (for case) is usually assumed to precede A-bar movement (for focus). Here, the same movement serves both purposes. Thus, what drives A and A-bar movement operations in the derivation must be treated as a major empirical issue.

Acknowledgments I would like to thank the editors of EISS, Christopher Pinon and Laurent Roussarie, for their extensive support. The article has greatly benefitted from discussions with Patrick Grosz, Pritty-Patel Grosz, Ashwini Deo, Rajesh Bhatt, Omer Preminger, Jonathan Bobaljik, Sabine Iatridou, Shigeru Miyagawa, Elena Anagnostopoulou and Mark Baker. I am thankful to the audiences at the CSSP conference, Linglunch at MIT and Forum of Theoretical Linguistics reading group at University of Oslo and to the anonymous reviewers of CSSP 2019/EISS 13 for their questions, comments and suggestions. Thanks are also due to Bhargavi Munshi, Trusha Joshi and Ojas Bhatt for their extensive help during the fieldwork.

References

Aissen, Judith. 2003. Differential object marking: Iconicity vs. economy. *Natural Language & Linguistic Theory* 21(3). 435–483.

Baker, Mark. 2015. *Case: Its principles and its parameters*. Cambridge: Cambridge University Press.

Baker, Mark & Jonathan Bobaljik. 2017. On inherent and dependent theories of ergative case. In Jessica Coon, Diane Massam & Lisa deMena Travis (eds.), *The Oxford handbook of ergativity*, 111–134. Oxford: Oxford University Press.

Bárány, András & Michelle Sheehan. 2019. When dependent case is not enough. In *Paper presented GLOW42, Oslo 07–11 May 2019 Main colloquium*.

Bhatt, Rajesh & Elena Anagnostopoulou. 1996. Object shift and specificity: Evidence from ko-phrases in Hindi. In *Papers from the 32nd regional meeting of the Chicago Linguistics Society, Chicago Linguistics Society, Chicago*, vol. 1122.

Cardona, George. 1965. *A Gujarati reference grammar*. ERIC.

Cardona, George & Babu Suthar. 2003. Indo-Aryan Languages. In George Cardona & Dhanesh Jain (eds.), *The major languages of South Asia, the Middle*

East and Africa, 29–34. New York: Routledge.
Chomsky, Noam. 2000. Minimalist inquiries: The framework. In Roger Martin, David Michaels & Juan Uriagereka (eds.), *Step by step: Essays on minimalist syntax in honor of Howard Lasnik*, 89–155. Cambridge, MA: The MIT Press.
Demonte, Violeta. 1995. Dative alternation in Spanish. *Probus* 7(1). 5–30.
Diesing, Molly. 1992. *Indefinites* (Linguistic inquiry monographs 20). Cambridge, MA: The MIT Press.
Diesing, Molly. 1997. Yiddish VP order and the typology of object movement in Germanic. *Natural Language & Linguistic Theory* 15(2). 369–427.
Dryer, Matthew S. 1986. Primary objects, secondary objects, and antidative. *Language* 62(4). 808–845.
Hamblin, Charles L. 1973. Questions in Montague English. *Foundations of Language* 10(1). 41–53.
Haspelmath, Martin. 2005. Argument marking in ditransitive alignment types. *Linguistic discovery* 3(1). 1–21.
Jayaseelan, Karattuparambil A. 2008. Topic, focus and adverb positions in clause structure. *Nanzan Linguistics* 4. 43–68.
Joshi, Kinjal. 2019. Agreement puzzles in Surati Gujarati. Unpublished manuscript University of Oslo.
Laka, Itziar. 1993. Unergatives that assign ergative, unaccusatives that assign accusative. *MIT working papers in linguistics* 18. 149–172.
Manzini, M. Rita & Ludovico Franco. 2016. Goal and DOM datives. *Natural Language & Linguistic Theory* 34(1). 197–240.
Mistry, P. J. 1998. Objecthood and specificity in Gujarati. In Jane H. Hill, P.J. Mistry & Lyle Campbell (eds.), *The life of language : Papers in linguistics in honor of William Bright*, 425–442. Berlin: Mouton de Gruyter.
Mistry, PJ. 2004. Subjecthood of non-nominatives in Gujarati. *Typological studies in language* 61. 1–32.
Pineda, Anna. 2013. Double object constructions and dative/accusative alternations in Spanish and Catalan: A unified account. *Borealis–An International Journal of Hispanic Linguistics* 2(1). 57–115.
Pollock, Jean-Yves. 1989. Verb movement, universal grammar, and the structure of IP. *Linguistic inquiry* 20(3). 365–424.
Sigurðsson, Einar Freyr & Jim Wood. 2012. Case alternations in Icelandic 'get'-passives. *Nordic Journal of Linguistics* 35(3). 269–312.
Suthar, Babubhai Kohyabhai. 2005. *Agreement in Gujarati*. Philadelphia, PA: University of Pennsylvania dissertation.

Non-responsibility and narrow scope reading of positive polarity indefinites in negative imperatives and negated controlled infinitive complements

Makoto Kaneko

Abstract This study examines an apparently peculiar narrow scope reading of positive polarity indefinites in negative imperatives, as in *Do not question somebody's immigration status*, and in negated controlled infinitive complements, as in *I was trying not to be somebody's bitch*, and proposes, referring to Richardson (1985) and Szabolcsi (2010), that (I) non-responsibility induced by the predicate leads us (i) to project a subjunctive complement, (ii) to allow the negation to be interpreted in a clause-external position, and (iii) to coerce a monitoring predicate, *take care*, all of which amount to paraphrasing the above two examples by *Take care for it not to be the case that you question somebody's immigration status* and *I was trying to take care for it not to be the case that I be somebody's bitch*; (II) this analysis in terms of coercion is more compatible with a property-denoting approach to imperatives and controlled complements, than with a proposition-denoting approach.

Keywords non-responsibility · positive polarity item · negation · imperative · obligatory control · obviation

M. Kaneko, Aoyama Gakuin University, https://nrid.nii.ac.jp/en/nrid/1000000362947/

In Christopher Pinon & Laurent Roussarie (eds.), Empirical Issues in Syntax and Semantics 13, 25–54. Paris: CSSP. http://www.cssp.cnrs.fr/eiss13/

1 Introduction

This study aims first at examining an apparently peculiar narrow scope reading of indefinite pronouns, observed in some negative imperatives and negated controlled infinitive complements, and second at shedding new light on the discussion about imperatives and controlled infinitives, two constructions which "semantically show significant parallels" (Potsdam 1998: 216), and which have yielded parallel approaches in terms of denoting properties vs. propositions. The discussion is based mainly on data from English, but also from other languages, especially Japanese.

The indefinite pronoun *somebody* in (1a-b) is out-scoped by a clause-mate negation.[1,2] Indefinites with *some* are known as Positive Polarity Items (PPIs) which should take wide scope over a clause-mate negation. The narrow scope reading observed in (1a-b) is at first glance surprising.

(1) a. [How to avoid offending minority customers?] Do not touch them. [...] Do not use profanity or racial slurs, and do not question **somebody**'s immigration status. (COCA)
b. Basically, I was trying not to be **somebody**'s bitch for 10 years. (COCA)

In order to account for this reading, I invoke (i) Richardson (1985) who claims that some negative imperatives, like in (2a), convey not a prohibition but a command of monitoring activity, and proposes to paraphrase its meaning by the coerced monitoring predicate TAKE CARE, as in (2b), and (ii) Szabolcsi (2010) who argues that a controlled infinitive complement denoting a situation where the *responsibility* relation (defined below) doesn't hold, as in (3a), is semantically equivalent to a subjunctive complement and that the negation may be interpreted in a lowered but clause-external position, as in (3b).

(2) a. Don't fall off the ladder! (Richardson 1985: 247)
b. TAKE CARE not to fall off the ladder. (*ibid.*)

(3) a. I don't want to offend **someone**. (Szabolcsi 2004: 417, fn. 10)
b. I want for it not to be the case [that I offend **someone**]. (Szabolcsi 2010: 6)

Partly modifying their ideas, this study proposes that (I) in examples like

[1]The same is true for the Japanese indefinite pronoun *nanika* 'something' in (i).

(i) ***nanika**-o* *okiwasure-nai yooni-si-naitoikenai.*
something-ACC misplace-NEG so.as.to-do-should
'I should take care not to misplace **something**.'
(`https://tinyurl.com/y85csgor`)

[2]Glosses follow the Leipzig Glossing Rules. Abbreviations not included in their list are: CONJ = conjunctive verbal form; DISJ = disjunctive verbal form; RESP = responsibility; PPI = positive polarity item; TDL = to-do-list; SenP = Sentential Phrase.

(1a-b), the nature of the predicate leads (i) to coerce the predicate TAKE CARE, (ii) to project a subjunctive complement, and (iii) to allow negation to be interpreted in a clause-external position, all of which amount to paraphrasing (1a-b) by (4a-b); (II) this analysis in terms of a coerced predicate is more compatible with a property-denoting approach to imperatives and controlled complements, than with a proposition-denoting one.

(4) a. TAKE CARE for it not to be the case [that you question **somebody**'s immigration status].
b. I was trying to TAKE CARE for it not to be the case [that I be **somebody**'s bitch].

The outline of this study is as follows. §2 reviews Richardson's and Szabolcsi's analyses in view of the notion of responsibility. §3 presents my main proposals. Next, after having quickly reviewed two approaches to imperatives and to infinitive complements, §4 compares their compatibility with the analysis proposed here. §5 recapitulates my main claims.

2 Previous analyses

In this section, I first present the notion of responsibility introduced by Farkas (1988), which allows us to capture the semantic parallelism between imperatives and obligatory control (§2.1). I next review, from the perspective of responsibility, Richardson's (1985) analysis of negative imperatives (§2.2). Then, I discuss Szabolcsi (2010) and Zu (2018) who suggest that infinitive complements are structurally ambiguous depending on the nature of the predicate (§2.3 and §2.4).

2.1 The responsibility relation

Farkas (1988: 36) advocates, to account for the semantics of controlled infinitives, as in (5a), a responsibility relation, defined as a "two-place relation, RESP(*i*,*s*), holding between an individual *i* and a situation *s* just in case *i* brings *s* about, that is, just in case *s* is the result of some act performed by *i* with the intention of bringing *s* about". She then notes that the RESP relation is equally relevant to appropriate uses of the rational clause *in order to*, as in (5b), the adverb *intentionally*, as in (5c), and imperatives, as in (5d).

(5) a. Mary ordered John to leave. (Farkas 1988: 31)
b. John read Anna Karenina in order to impress Mary. (*idem*.36)
c. John fell off the ladder intentionally. (*idem*.39)
d. [To John] Be polite! (*ibid*.)

(5a-d) are acceptable since a RESP relation holds between John and his leaving, his reading 'Anna Karenina', his falling off the ladder or his being polite, while (6a-d) are strange since it is difficult in ordinary situations to assume a RESP relation between John and his resemblance to his father.

(6) a. #Pete ordered John to resemble his father. (*idem*.46)
b. #John resembles his father in order to annoy his grandmother. (*idem*.36)
c. #John resembles his father intentionally. (*idem*.39)
d. #[To John] Resemble your father. (*ibid*.)

2.2 Negative imperatives as a command of a monitoring activity

The RESP relation equally holds in negative imperatives expressing a prohibition, as in (7a). This example is in effect incompatible with the non-RESP-inducing adverb *inadvertently*, as in (7b).

(7) a. Don't trip!
b. #Don't inadvertently trip!

This is not the case with the negative imperatives in (8a-b): A RESP relation does not hold between the addressee and her not falling off the ladder or her not catching cold, as confirmed by the compatibility with *inadvertently*.

(8) a. [To a carpenter] Don't (inadvertently) fall off the ladder! (Adapted from (2a))
b. Don't (inadvertently) catch cold! (Adapted from Richardson: 1985)

According to Richardson (1985: 246), in (8a), "not falling off the ladder is for the carpenter not an act [...], but rather a state of affairs which he must try to maintain if he is going to accomplish whatever his positive intentions might be". The same is true for not catching cold in (8b).

In other words, negative imperatives like (8a-b) do not convey prohibitions but "commands to monitor against the events denoted by the VPs" (*idem*.247). Richardson proposes to capture this meaning by means of coercion of the monitoring predicate TAKE CARE, as in (9a-b).

(9) a. TAKE CARE not to fall off the ladder. (*ibid.*)
b. TAKE CARE not to catch cold. (*ibid.*)

Now, the negative imperative in (1a) equally conveys, in the relevant context where the speaker gives some advises to avoid offending minority customers, a command of a monitoring activity, which is confirmed by the acceptability of *inadvertently*, as in (10a), and by a possibility of the paraphrase in (10b).

(10) a. Do not inadvertently question somebody's immigration status. (Adapted from (1a))
b. TAKE CARE not to question somebody's immigration status.

But the motivation of coercion remains unclear, and the paraphrase in (10b) does not yet account for the narrow scope reading of *somebody* under negation. To tackle the second question, I review Szabolcsi's (2010) analysis in the next section.

2.3 The ambiguity of controlled infinitives, and obviation exemptions

Szabolcsi's (2010) analysis is based on Farkas's (1992) discussion on obviation. Obviation is a constraint requiring the subject of a subjunctive clause to be disjoint in reference from the matrix subject. Farkas (1992: 104) claims that there is a "canonical control case", where "both the complement [subject] and the matrix argument it is referentially dependent on bear the RESP relation to [the complement situation]" and that "obviation is strongest in case the semantic characteristics of canonical control are met". Thus, in Hungarian, when the complement of volition verbs denotes an intentional act, like *go to the movie* in (11a-b), the semantic characteristics of canonical control are met and obviation is imposed: when matrix and complement subjects are coreferential, the infinitive in (11a) is chosen, while the subject of a subjunctive clause cannot be coreferential with the matrix subject, as in (11b).

(11) a. *János moziba akar men-ni.* (Hungarian)
J. movie-into want.3SG go.INF
'János wants to go to the movie.' (Farkas 1992: 92)
b. *János$_1$ azt akarja, hogy pro$_{*1/2}$ moziba men-jen.*
J. it.ACC want.3SG that (he/she) movie-into go-SBJV
'János wants that he/she go to the movie.' (*ibid.*)

However, when the context forces the complement to receive a non-RESP reading, obviation may fail to obtain and a controlled subjunctive complement, as in (12a), is as acceptable as a controlled infinitive complement, as in (13a).[3]

(12) *Controlled subjunctive complement*
a. *Nem akarom, hogy lelőjek valakit.* (Hungarian)
not want.1SG that shoot.SBJV.1SG someone.ACC
'I don't want that I shoot someone.' (Szabolcsi 2010: 7)
b. [RESP reading] #You can trust me with a gun, because I want to shoot no one. (Paraphrase given by Zu (2018: 152))
c. [non-RESP reading] Take this gun from me because I want to shoot no one. (*ibid.*)

(13) *Controlled infinitive complement*
a. *Nem akarok lelőni valakit.* (Hungarian)
not want.1SG shoot.INF someone.ACC
'I don't want to shoot someone.' (Szabolcsi 2010: 7)
b. [RESP reading] #You can trust me with a gun, because I want

[3]Ruwet (1990) observes that, equally in French, where obviation is usually observed, as in (ia), when the complement verb is passivized and the matrix subject does not bear a RESP relation with the complement situation, as in (ib), this constraint is relaxed and the subjunctive becomes more acceptable, although it remains less acceptable than the infinitive in (ic).

(i) a. *Je veux que je parte. (French)
'I want for me to leave.' (Ruwet 1990: 18)
b. ?Je veux que je sois enterré dans mon village natal. (French)
'I want for me to be buried in the village of my birth.' (*idem*.20)
c. Je veux être enterré dans mon village natal. (French)
'I want to be buried in the village of my birth.' (*ibid.*)

to shoot no one. (Paraphrase given by Zu (2018: 152))

c. [non-RESP reading] Take this gun from me because I want to shoot no one. (*ibid.*)

Szabolcsi (2010) further observes that, both in subjunctive and infinitive complements, indefinite PPIs may be interpreted differently with respect to the matrix negation depending on the complement meaning: the narrow scope reading of *valakit* 'someone' is impossible when the complement situation receives a RESP reading, as in (12b) and (13b), but possible when it obtains a non-RESP reading, as in (12c) and (13c). Based on this parallel interpretation, Szabolcsi claims that (i) the infinitive complement is structurally different in canonical and non-canonical control; (ii) in the latter, the controlled infinitive, as in (13a), is structurally parallel to the controlled subjunctive, as in (12a).

She next shows that such a contrast between RESP and non-RESP readings is equally observed in English controlled infinitives. Thus, in (14a) with the RESP-inducing verb *call*, *someone* usually out-scopes the matrix negation, while in (14b) including the non-RESP-inducing verb *offend*, *someone* may take a narrow scope reading.

(14) a. *Controlled infinitive with RESP predicate*
I don't want to call **someone**. [*not > some]
(Szabolcsi 2004: 417, fn. 10)

b. *Controlled infinitive with non-RESP predicate*
I don't want to offend **someone**. [✓not > some] (= (3a))

c. I want for it not to be the case [that I offend **someone**]. (= (3b))

This author then claims that (i) just like in Hungarian, English infinitive complements are structurally different in canonical and non-canonical control; (ii) in the latter, as in (14b), the controlled infinitive is structurally parallel to the controlled subjunctive, as in (14c). Furthermore, in order to account for the narrow scope reading of *someone* in (14b), Szabolcsi (2010: 6-7) proposes the following hypotheses: (i) (14b) "doesn't mean *I have no desire to offend someone*, it means, or it definitely can mean *I want not to offend*", that is, the matrix negation is lowered; (ii) the non-RESP relation in (14b) leads to project an implicit extra layer that she calls "non-RESP marker", and that she represents by *for it to be the case that* in (14c); (iii)

the narrow scope reading of *someone* in (14b) is due to the fact that the non-RESP marker "shields from negation"; (iv) "the presence of non-RESP marker is decided locally, by looking at the complement, and not by looking at the relation between a participant of the matrix situation and the complement".

However, the details of her hypotheses raise some problems: (i) the first hypothesis presupposes, without discussion, that the volition verb *want* allows implicit lowering of the matrix negation. But the conditions of Neg-lowering should be examined in a more detailed way (see the discussion in §3.1); (ii) the idea of a non-RESP marker heading a syntactic projection, advocated in the second hypothesis, is not clear. Especially, while the RESP relation between an individual, *i*, and a situation, *s*, is expressed by means of a predicate, as in "*i* intentionally brings about *s*", this is not the case for the non-RESP relation; (iii) the third hypothesis, phrased in terms of "shielding", faces empirical problems;[4] (iv) the fourth hypothesis is based on Szabolcsi's observation that in Hungarian, the scope of indefinite PPIs in non-controlled subjunctive complements depends on the (non) RESP nature of the complement, as in (15b-c), just in the same way as in controlled infinitive complements, illustrated by (16b-c). Especially Szabolcsi argues that indefinite PPIs in non-controlled subjunctive complements conveying a RESP situation cannot be out-scoped by the matrix negation, as in (15b).

(15) *Non-controlled subjunctive complement*

a. *Nem akarom, hogy leugorjál valahonnan.* (Hungarian)
not want.1SG that jump.SBJV.2SG from.somewhere
'I don't want that you jump from somewhere (I don't want for you to jump from somewhere).' (Szabolcsi 2010: 7)

[4]Goncharov (2018) observes that on the one hand, in ordinary shielding cases, as in (i), where *always* serves as an intervener, PPI *some* is accepted, while NPI *any* is not; on the other hand, in the relevant construction in (ii), *any* is accepted. She then argues that the lack of intervention effect puts into question the analysis in terms of shielding.

(i) John doesn't *always* call {someone/ *anyone} (Szabolcsi 2004: 415-416)

(ii) I don't want to offend {someone / anyone}. (Goncharov 2018)

b. Jumping is a planned voluntary act. (RESP) [*not > some]
c. Jumping is due to an urge, or inability to resist a temptation. (non-RESP) [✓not > some]

(16) *Controlled infinitive complement*

a. *Nem akarok leugrani valahonnan.* (Hungarian)
not want.1SG jump.INF from.somewhere
'I don't want to jump from somewhere.' (*idem*.6)
b. Context: Look at those rocks. It would be fun to climb and jump. – I am not going. I don't want to jump from anywhere. (RESP) [*not > some]
c. Context: there's a great view from those rocks. Let's go climb them. – I have the fear of heights. I don't want to jump from anywhere. (non-RESP) [✓not > some]

Szabolcsi's analysis is thus recapitulated by the configurations in (17)–(19): in all three types of complements (controlled infinitive, controlled subjunctive and non-controlled subjunctive), RESP and non-RESP cases yield structural differences: the non-RESP marker projects an extra layer serving to create an anti-licensing domain for PPIs.

(17) *Controlled infinitive complement*

a. [DP1 … not INF$^{\text{RESP}}$ … PPI] [*not > some]
b. [DP1 … not [INF$^{\text{non-RESP}}$ PPI]] [✓not > some]

(18) *Controlled subjunctive complement*

a. *[DP1 … not pro SBJV$^{\text{RESP}}$ PPI] [obviation]
b. [DP1 … not [pro SBJV$^{\text{non-RESP}}$ PPI]] [✓not > some]

(19) *Non-controlled subjunctive complement*

a. [DP1 … not DP2 SBJV$^{\text{RESP}}$ PPI] [*not > some]
b. [DP1 … not [DP2 SBJV$^{\text{non-RESP}}$ PPI]] [✓not > some]

However, French PPIs like *quelqu'un* 'someone' in non-controlled subjunctive complements, as in (20a), and English PPIs like *someone* in non-controlled infinitive complements, as in (20b),[5] take narrow scope under

[5]Given that it is possible to add the prepositional complementizer *for* before the infinitive subject, I assume that the latter is inside the complement and that this construction

the matrix negation, when the complements include the RESP-inducing verb *kill*.

(20) a. *Non-controlled subjunctive complement*
Eliott, je ne veux pas que tu tues quelqu'un … Une nouvelle vie s'ouvre à toi … ce serait bien qu'elle ne se déroule pas en prison … (French)
'Eliott, I don't want you to kill someone … A new life is open to you … it would be nice if your new life did not go on in prison …' (M. Carrieu 1990. *Le Regarder en Face*.)

b. *Non-controlled infinitive complement*
Then why do you want to spare her? – I have many reasons. First, I don't want you to kill someone. Second, I'm opposed to murder in general. And third, I'm curious about that phantom girl. (`https://tinyurl.com/y9wmgmqp`)

Taking into account these data, it seems to be more appropriate to draw a line between canonical-control cases on the one hand and non-canonical control and non-control cases on the other hand, as shown in (21a-c).

(21) a. *Canonical controlled infinitive*
[DP … not INF$^{\text{RESP}}$ PPI] [*not > some]

b. *Non-canonical controlled infinitive or subjunctive*
[DP … not [INF$^{\text{non-RESP}}$ PPI / pro SBJV$^{\text{non-RESP}}$ PPI]] [✓not > some]

c. *Non-controlled infinitive or subjunctive*
[DP1 … not [(for) DP2 INF PPI / DP2 SBJV PPI]] [✓not > some]

In the next subsection, I review Zu (2018) who proposes a detailed analysis in this direction, based on the conjunct vs. disjunct verbal markings in Newari.

2.4 Canonical and non-canonical control: Zu (2018)

The conjunct vs. disjunct distinction in Newari shows similarities with and differences from the infinitive vs. subjunctive distinction in Hungarian. On the one hand, when matrix and complement subjects are coreferen-

corresponds to a non-controlled subjunctive complement in French.

tial, and the complement is of canonical control type, including the RESP-inducing predicate *go there*, the conjunct suffix (corresponding to the infinitive) should be used, as in (22a), and the disjunct one (corresponding to the subjunctive) is excluded, as in (22b).

(22) *Canonical controlled complement*

a. *wõ: wa ana wan-ã dhakã: dhãla.* (Newari)
(s)he.ERG (s)he there go-PST.CONJ that said
'(S)he$_1$ said that (s)he$_1$ went.CONJ there.' (Zu 2018: 143)

b. **wõ: wa ana wan-a dhakã: dhãla.* (Newari)
(s)he.ERG (s)he there go-PST.DISJ that said
*'(S)he$_1$ said that (s)he$_1$ went.DISJ there.' (*ibid.*)

On the other hand, in non-canonical control cases, as in (23) including the non-RESP inducing predicate *accidentally run into someone*, acceptability is reversed: the conjunct suffix is excluded, as in (23a), differently from a Hungarian infinitive form in (13b), and the disjunct suffix should be used, as in (23b). In other words, the conjunct vs. disjunct distinction in Newari more transparently reflects the distinction between canonical and non-canonical control than the infinitive vs. subjunctive distinction in Hungarian.

(23) *Non-canonical controlled complement*

a. **Shyam-a dhãla ki wõ: masika shun*
Shyam-ERG said that he.ERG accidentally someone
nãpalãt-ã
meet-PST.CONJ
*'Shyam$_1$ said that he$_1$ accidentally ran.CONJ into someone.'
(*idem.*146)

b. *Shyam-a dhãla ki wõ: masika shun*
Shyam-ERG said that he.ERG accidentally someone
nãpalãt-a
meet-PST.DISJ
'Shyam$_1$ said that he$_1$ accidentally ran.DISJ into someone.'
(*ibid.*)

Zu (2018) proposes the following analysis of this distribution. In the

left periphery of the complement clause is projected a Sentential Phrase (SenP) which conveys a perspective from which the complement clause is presented. The SenP forms a local domain and normally functions as a phase for syntactic operations. However, in canonical control, as in (24a), where the matrix subject (coreferential with the complement subject) bears a RESP relation with the complement situation, "the Sen head contains unchecked phi-features" (Zu 2018: 157), which "must be checked against a contextually salient individual" (*idem.*158), provided by the matrix subject. The domain is thus not closed off (that is, suspended), and the SenP does not form a local domain, as in (24a). But, "in non-canonical control [...] the Sen head enters the derivation with no phi-features" (*ibid.*). In the same vein, "the Sen head in non-control does not have uninterpretable features" (*idem.*159). The SenP therefore serves as a phase, as in (24b-c).

(24) a. *Canonical control* (Simplified from Zu (2018: 154))
[$_{\text{CP1 (phase1)}}$ DP ... [$_{\text{SenP [uninterpretable phi-features]}}$ pro verb.CONJ]]
b. *Non-canonical control*
[$_{\text{CP1(phase1)}}$ DP ... [$_{\text{SenP(phase2)}}$ pro verb.DISJ]]
c. *Non-control*
[$_{\text{CP1(phase1)}}$ DP1 ... [$_{\text{SenP(phase2)}}$ DP2 verb.DISJ]]

According to Zu (2018), the different scope readings of PPIs in canonically and non-canonically controlled English infinitives in (25a-b) are also due to the presence or absence of uninterpretable phi features. The narrow scope reading of PPIs in non-controlled infinitives, as in (25c), may also be accounted for in the same way

(25) a. *Canonical control*
[$_{\text{phase1}}$ I don't want PRO to call someone]. [*not > some] (*idem.*158)
b. *Non-canonical control*
[$_{\text{phase1}}$ I don't want [$_{\text{phase2}}$ PRO to offend someone]]. [✓not > some] (*ibid.*)
c. *Non-control*
[$_{\text{phase1}}$ I don't want [$_{\text{phase2}}$ (for) you to kill someone. [✓not > some]

The dichotomy *canonical control* vs. *non-canonical control* and *non-control* is further supported by the fact that an explicit occurrence of the infinitive subject is forced in non-control, as in (26a), and is not so bad in non-canonical control, as in (26b), but is totally unacceptable in canonical control, as in (26c).

(26) a. *Non-control*
I want for you to come back in my life. (Zu 2018: 153)
b. *Non-canonical control*
(?)I want for me to be in shape for tomorrow's game. (*ibid.*)
c. *Canonical control*
*I want for me to call my mom. (*ibid.*)

Thus, Zu's analysis is essentially based on the different feature setting of the Sen head: it bears uninterpretable phi-features in canonical control, but not in non-canonical control and non-control. This hypothesis may be valid for Newari where the dichotomy is explicitly observed through the different verbal suffixes, but sounds ad hoc for English where morphological support is lacking. Furthermore, unlike Szabolcsi (2010), Zu (2018) does not invoke Neg-lowering in order to account for the narrow scope reading of PPIs in (25b). Therefore, the narrow scope of PPIs with respect to the matrix negation in (25b) and with respect to the clause-mate negation in (1a-b) are essentially independent phenomena, and the analysis of the former case cannot be applied to the latter in spite of the common non-RESP nature of the complement.

If we however assume, with Szabolcsi, that a non-RESP predicate serves to make infinitive complements equivalent to subjunctive complements, rather than to prevent domain suspension, the similar narrow scope reading of PPIs in the two different contexts may be analyzed in a parallel way. This possibility will be pursued in §3.

2.5 Recapitulation

In this section, I present (i) Richardson's analysis according to which negative imperatives including a non-RESP predicate, like in (27a), are paraphrased by means of the coerced monitoring predicate TAKE CARE, as in (27b), and (ii) Szabolcsi's analysis according to which non-RESP controlled infinitive complements, as in (28a), are equivalent to subjunctive ones. Sz-

abolcsi further suggests that the matrix negation is lowered into a position where it clause-externally scopes over the complement, as in (28b).

(27) a. Do not question somebody's immigration status. (= (1a))
b. TAKE CARE not to question somebody's immigration status.

(28) a. I do not want that I offend someone.
b. I want for it not to be the case [that I offend **someone**].

There remain some questions. As regards Richardson's analysis, the motivation of coercion is not clear, and it is not clear how PPIs are licensed in (27b). Concerning Szabolcsi's analysis, it is not clear when the matrix negation is interpreted in a lowered position, and why it is not interpreted inside the complement, as in ordinary Neg-lowering cases (e.g. "It is not likely that she will go" paraphrasable by "It is likely that she will not go."), but in an intermediate position, as in (28b).

3 Proposals

To answer these questions, I propose the following hypotheses.
(I) The matrix negation in (29a) may be implicitly lowered down into the complement when the volition predicate conveys an *effective preference*.
(II) The non-RESP nature of the complement predicate allows not only the controlled infinitive complement to be reanalyzed as a subjunctive one, but also the negation once lowered to be interpreted in a clause-external position, as in (29b). Furthermore, a monitoring predicate is coerced.
(III) In negative imperatives, as in (30a), the non-RESP nature of the predicate brings about the projection of the subjunctive complement, the clause-external interpretation of the negation, and the coercion of the monitoring predicate, all of which allow a PPI in the complement to be licensed by a clause-external negation, as in (30b). Essentially the same operations occur in negated controlled complements, as shown in (31a-b).

(29) a. I don't want to offend someone.
b. I want$^{\text{effective preference}}$ to TAKE CARE for it not to be the case [that I offend **someone**].

(30) a. Do not question somebody's immigration status. (= (1a))
b. TAKE CARE for it not to be the case [that you question **some-**

body's immigration status].

(31) a. I was trying not to be somebody's bitch for 10 years. (= (1b))
b. I was trying to TAKE CARE for it not to be the case [that I be **somebody**'s bitch].

In what follows, after having introduced the background of Hypothesis (I), I present evidence in favor of each of the three hypotheses, based on Japanese data.

3.1 Background of Hypothesis (I)

Levinson (2003) observes that one and the same reasonable person can utter (32a) and (32b) to reply to the same question in (32), within a short time and without having changed her mind.

(32) Do you want to play tennis?
a. I really want to play, but I have to teach.
b. No [= I don't want to], I have to teach.

Levinson (2003: 222-223) claims that this is because *want* is ambiguous: it denotes a mere desire "as a matter of psychological fact" in (32a), but a "desire accompanying intentional action" in (32b). Condoravdi & Lauer (2012) call the second meaning "effective preferences". Grano (2018) proposes to paraphrase the two readings by *would like* and *intend*. I hereafter call the two readings *want*$^{\text{would-like}}$ and *want*$^{\text{intend}}$. One test to disambiguate them is provided by *anankastic conditionals*, roughly defined by (33a) and illustrated by (33b).

(33) a. For an agent *a* and predicates *P* and *Q*,
if *a* wants to *P*, *a* must *Q* = *a* must *Q* in order to *P*
(*Q* is a necessary condition for *P*)
b. If you want to go to Harlem, you must take the A train.
= You must take the A train in order to go to the Harlem.
(Condoravdi & Lauer 2016: 2)

Now, the compatibility of *want* with the anankastic conditional in (34a) indicates that it is interpreted here as *want*$^{\text{intend}}$. In the paraphrase in (34b), negation scopes over the infinitive clause, which suggests that it is

also interpreted in (34a) in a lowered position where it scopes over the infinitive complement.

(34) a. If you don't want to offend someone, you must watch your behavior well.
b. = You must watch your behavior well in order not to offend someone.

To further support the correlation between the *want*$^{\text{intend}}$ reading and the availability of the lowered-Neg interpretation, I refer to Horn's (2001: 320) analysis in terms of *(in)tolerance*, informally represented by (35). For example, on the one hand, *likely* is intolerant, as shown by the incongruity of (36a). The disjunction in (36b) is exclusive. When the disjunction is shared between discourse participants, if the speaker asserts that the first disjunct is false, the disjunction being exclusive, the hearer can infer the truth of the second disjunct, where the negation is lowered.

(35) For a predicate P and a proposition p, (i) when $P(p) \vee P(\neg p)$ is shared by discourse participants, and (ii) when P is intolerant [i.e. $^{*}P(p) \wedge P(\neg p)$], if the speaker utters $\neg P(p)$, the hearer can infer $P(\neg p)$.

(36) a. #It's likely she'll go and likely she won't go. (Horn 2001: 320)
b. (It's likely she will go)∨(It's likely she won't go) [exclusive]

On the other hand, *possible* is tolerant, as shown by the acceptability of (37a). Therefore, in the inclusive disjunction in (37b), even if the speaker denies the truth of the first disjunct, the hearer cannot conclude the truth of the second disjunct. Now, intolerant *likely*, but not tolerant *possible*, yields a lowered-Neg reading, as in (38a-b).

(37) a. It's possible she'll go and possible she won't go.
b. (It's possible she will go) ∨ (It's possible she won't go) [inclusive]

(38) a. It's not likely she will go. ≒ It's likely she will not go.
b. It's not possible she will go. ≠ It's possible she will not go.

Now, are the two readings of *want* intolerant or tolerant? Condoravdi &

Lauer (2016: 28) point out that *want* allows contradictory wishes only for *want*$^{\text{would-like}}$: "the consistency of [(39)] is dependent on a contextual resolution for *want* where the targeted preference is 'mere desire' [...]. While [(39) with *want*$^{\text{would-like}}$] is coherent (and simply attributes indecision to John), [(39) with *want*$^{\text{intend}}$] sounds contradictory (or attributes a certain amount of irrationality to John)".

(39) [John said] I want to move in with my girlfriend, but I also want to keep living alone. (Condoravdi & Lauer 2016: 28)

In this example, the complement of the second conjunct, *to keep living alone*, amounts to the negation of the complement of the first conjunct, *to move in with my girlfriend*. (39) thus suggests that *want*$^{\text{would-like}}$ is tolerant, while *want*$^{\text{intend}}$ is intolerant. We have now confirmed, from a theoretical viewpoint, Hypothesis (I), according to which the matrix negation may be lowered when the volition predicate conveys an effective preference. In effect, with a non-RESP predicate denoting a situation to avoid for social or moral reasons, like *offend* in (3a), *not want to* tends to convey not a simple absence of desire (which sounds half-hearted in these contexts), but the expression of an intention to monitor against an occurrence of the relevant situation.

3.2 Evidence for Hypotheses (I) and (II)

Some Japanese data provide empirical support for Hypotheses (I) and (II). The Japanese volitional form *tai*, as in (40a), requires the coreference of the matrix and complement subjects. When the two subjects are disjoint, another form *tehosii* with a dative-marked subject, is used, as in (40b).

(40) a. *watasi-wa gaisyutusi-tai.* (Japanese)
me-TOP go.out-want
'I want to go out.'
b. *watasi-wa anata-ni gaisyutusi-tehosii.*
me-TOP you-DAT go.out-want
'I want (for) you to go out.'

Japanese also possesses a strong NPI *daremo* which is licensed by a clausemate negation, as in (41a), but not by a clause-external negation, as in

(41b).

(41) a. *John-wa **daremo** kidutuke-nai.*
John-TOP anyone hurt-NEG
'John does not hurt **anyone**.'
b. **watasi-wa [John-ga **daremo** kidutukeru-to]-wa*
me-TOP John-NOM anyone hurt-COMP-*top*
omowa-nai.
think-NEG
'I don't think [that John hurts **anyone**].'

Now, when volitional *tai* is negated, the strong NPI *daremo* as well as the indefinite PPI *dareka*, as in (42a-b), can appear in the complement (in these examples and in the examples below, I only focus on the narrow scope reading of *dareka*). In both cases, an anankastic conditional reading is available, as shown in (43), which indicates that volitional *tai* conveys here an effective preference, and that the matrix negation may therefore a priori be interpreted in a lowered position.

(42) a. *daremo kizutuke-taku-nai.*
anyone hurt-want-NEG
'I do not want to hurt anyone.'
(`https://tinyurl.com/ybdawhod`)
b. *dareka-o kizutuke-taku-nai.*
someone-ACC hurt-want-NEG
'I do not want to hurt someone.'
(`https://tinyurl.com/ybukoyjf`)

(43) *{daremo/dareka-o} kizutuke-taku-nai-nara kotoba-ni*
{anyone/someone-ACC} hurt-want-NEG-COND mouth-DAT
tyuuisiro.
watch.IMP
'If you don't want to hurt anyone/someone, watch your mouth.' =
You must watch your mouth in order not to hurt anyone/someone.[6]

[6]In fact, the wide scope reading of *dareka* over the negation is preferred here. The narrow scope reading is forced by adding another noun phrase after the indefinite, as in *dareka midikana hito* 'someone nearby'.

But the position where the negation is interpreted turns out to be different in *daremo* cases and in *dareka* cases. As in (44a) and (44b), in the complement of negated *tai*, *daremo* cannot follow the adverb *ukkari* 'inadvertently' which highlights the non-RESP nature of the complement, but can precede it. The same is true in simple clauses, as in (45a) and (45b).

(44) a. ??*ukkari* *daremo kizutuke-taku-nai.*
inadvertently anyone hurt-want-NEG
'I don't want to inadvertently hurt anyone.'
b. *daremo ukkari* *kizutuke-taku-nai.*
anyone inadvertently hurt-want-NEG
'I don't want to hurt anyone inadvertently.'

(45) a. ??*ukkari* *daremo kizutuke-nai.*
inadvertently anyone hurt-NEG
??'Inadvertently, I hurt no one.'
b. *daremo ukkari* *kizutuke-nai.*
anyone inadvertently hurt-NEG
'No one, I inadvertently hurt.'

I propose to attribute the unacceptability of (45a) to the semantic incongruity of the wide scope reading of the non-RESP-inducing adverb over the negation, as shown by that of the English translation. (45b) is acceptable since the scrambling of *daremo* allows the negation to out-scope the adverb, just as in the English translation including topicalized *no one*. Now, combining Hypothesis (I) advocating a lowered-Neg interpretation with Szabolcsi's idea according to which the non-RESP nature of the predicate allows a subjunctive complement, we can account for the acceptability contrast between (44a) and (44b) by paraphrasing them as in (46a) and (46b), where the matrix negation is interpreted inside the complement, and takes narrow or wide scope with respect to the RESP-inducing adverb (as regards the coercion of the causative predicate BRING IT ABOUT, see the discussion in §4.1).

(46) a. ??I want to BRING IT ABOUT that I inadvertently not hurt anyone.
b. I want to BRING IT ABOUT that no one I inadvertently hurt.

Next, (47a) shows that, contrary to *daremo*, *dareka* can follow the adverb

ukkari.[7] Hypothesis (II), which advocates a clause-external interpretation of the negation, allows us to paraphrase (47a) by (47b) and to account for the different acceptability of (44a), with the strong NPI *daremo*, and (47a), with the PPI *dareka*.

(47) a. *ukkari dareka-o kizutuke-taku-nai.*
inadvertently someone-ACC hurt-want-NEG
'I don't want to inadvertently hurt someone.'
b. I want to TAKE CARE for it not to be the case [that I inadvertently hurt **someone**].

The different positioning of the lowered negation in *daremo* cases and in *dareka* cases is further confirmed by their different compatibility with another strong NPI, *kessite* 'at all', which requires to be licensed by a clause-mate negation, as in (48a).

(48) a. *kessite {daremo/??dareka-o} kizutuke-taku-nai.*
at all {anyone/someone-ACC} hurt-want-NEG
'I don't want to hurt {anyone/someone} at all.'
(Adapted from `https://tinyurl.com/yb229p6y`)
b. *kessite {daremo/??dareka-o} kizutuke-nai.*
at all {anyone/someone-ACC} hurt-NEG
'I never hurt {anyone / someone}] at all.'

Under the proposed hypothesis, *daremo* in (48a) is acceptable since the negation lowered into the complement can also license the strong NPI *kessisite*, just as in (48b) where the two strong NPIs are in the same clause as the negation; *dareka* is unacceptable here since it requires a clause-external negation while *kesssite* requires a clause-mate negation.

[7] Its compatibility with the non-RESP-inducing auxiliary, *tesimau* 'happen to', also distinguishes *dareka* from *daremo*, as in (i), which indicates that the negation is interpreted in a higher position in *dareka* cases than in *daremo* cases.

(i) *{??daremo/dareka-o} kizutuke-tesimai-taku-nai.*
{anyone/someone-ACC} hurt-happen.to-want-NEG
??'I want to happen to hurt no one / I want not to happen to hurt someone.'

3.3 Evidence for Hypothesis (III)

The 'inadvertently' (*ukkari*) test also supports Hypothesis (III). In the negated complement in (49a) and the negative imperative in (50a), this adverb can precede the PPI *nanika* 'something', but not the strong NPI *nanimo* 'anything'.

(49) a. *ukkari* *{??nanimo/nanika-o}* *kowasa-nai*
Inadvertently {anything/something-ACC} break-NEG
yoonisi-ta.
try-PST
??'I tried to inadvertently break nothing / I tried not to inadvertently break something.'

b. I tried to TAKE CARE for it not to be the case [that I inadvertently break **something**]

(50) a. *ukkari* *{??nanimo/nanika-o}* *kowasu-na.*
inadvertently {anything/something-ACC} break-IMP.NEG.
??'Inadvertently break nothing! / Don't inadvertently break something!'

b. TAKE CARE for it not to be the case [that you inadvertently break **something**].

Hypothesis (III) accounts for these contrasts as follows. On the one hand, in (49a) and (50a), *nanimo* is unacceptable since the negation is lowered into the subjunctive complement (projected due to the non-RESP nature of the predicate), and remains there to license the strong NPI. But when this NPI follows the adverb *ukkari*, the negation is out-scoped by the adverb, which yields a semantic incongruity. On the other hand, *dareka* is acceptable since the negation may be interpreted in a clause-external position in order to avoid an incongruous scope relation, just as in the paraphrases in (49b) and (50b).

Thus, we can now answer two of the three remaining questions: (i) when is the matrix negation lowered into the complement of a volitional predicate?; (ii) why is the negation interpreted clause-externally when the predicate is of non-RESP nature? But we still don't have an answer to the third question, that is, the motivation for coercion of the monitoring predicate. I will try to tackle this problem by examining previous approaches to im-

peratives and obligatory control.

4 Previous analyses of imperatives and obligatory control

This section reviews two main approaches to the semantics of imperatives and obligatory control, that is, the property-denoting approach (§4.1) and the proposition-denoting one (§4.2). Meanwhile, I also discuss how each approach deals with cases where the addressee or the attitude holder doesn't bear a RESP relation with the denoted situation, for example, third person imperatives, as in (51a), where "the person or persons who are supposed to carry out the action the imperative specifies are not [...] those to whom it is addressed" (Hamblin 1987: 53), and uncontrolled complements, as in (51b), where the matrix and complement subjects are distinct.

(51) a. Your guards be the diversion while we sneak in! (Potsdam 1998: 208)
b. John intended for Bill to leave. (Grano 2015: 242)

4.1 The property-denoting approach

According to Portner (2005; 2007; 2012), three major sentence types (declaratives, interrogatives and imperatives) correspond to different semantic types (proposition, set of propositions and property) and are related to different forces (Assertion, Asking and Requiring) and to different discourse components (Common Ground, Question Set and To-Do List). Declaratives denote a proposition, and serve to make an Assertion which boils down to adding a proposition into the Common Ground; interrogatives denote a set of propositions, and serve to make an Asking which amounts to add a set of propositions into the Question Set; imperatives denote a property, and serve to make a Requiring which boils down to adding a property to the To-Do List (set of properties), abbreviated by TDL below. Thus, an imperative sentence expresses a property, "which can only be true of the addressee" (Portner 2007: 358), as in (52).

(52) $[\![\textit{Sit down!}]\!] = [\lambda x.\ x \text{ sits down}] \in \text{TDL(you)}$
(Adapted from Portner 2007: 358)

Assuming, with Chierchia (1988), "that controlled complements denote properties", Grano (2015: 244) directly applies Portner's analysis of imper-

atives to the semantics of *intend*, which consists, as in (53a), in putting a property denoted by the complement into the matrix subject's Private TDL. Thus, (53b) means that the matrix subject, John, adds a property of leaving to his Private TDL.

(53) a. ⟦*intend*⟧ = $\lambda P_{\langle e,t\rangle}$ λx.P ∈ Private TDL(x) (Grano 2015: 244)
b. ⟦*John intends to leave*⟧ = [λx. x leave] ∈ Private TDL(j) (*ibid.*)

As regards non-controlled complements, Jackendoff & Culicover (2003) propose, for (54a) including another obligatory control verb *plan*, an analysis in terms of coercion of the silent causative predicate BRING IT ABOUT THAT, as in (54b).

(54) a. Hilary plans for there to be more light in here.
(Jackendoff & Culicover 2003: 542)
b. Hilary plans to BRING IT ABOUT THAT there is more light here.
(*ibid.*)

(55) a. ⟦CAUSE⟧ = $\lambda p\, \lambda x.$ *x* brings it about that *p*. (Grano 2015: 246)
b. ⟦*John intends for Bill to leave*⟧ (*ibid.*)
= [λx. x BRINGS IT ABOUT THAT Bill leave] ∈ Private TDL(j)

Formalizing the semantics of the causative predicate, as in (55a), Grano (2015) also proposes the analysis in terms of coercion, as in (55b) and further clarifies its motivation: a controlled complement denotes a property, while a non-controlled complement containing an explicit subject denotes a proposition. Coercion is required to repair type-mismatch.

Zanuttini et al. (2012: 1266, fn. 46) analyze essentially in the same way third person imperatives in Bhojpuri, as in (56).

(56) *kha :y* (Bhojpuri)
eat-IMP.3SG
'Let him eat.' (Zanuttini et al. 2012: 1252)

(57) a. [λx: x=addressee(c).[λw. he eats in w]] (*idem.*1266, fn. 46)
b. ⟦*kha :y*⟧ = [λx. x BRINGS ABOUT IT THAT he eats] ∈ TDL(you)

These authors then represent the semantics of (56) by (57a), suggesting that "intuitively, this places a requirement on the addressee that is only

satisfied if the referent of *he* eats." But (57a) is at odds with TDL canonical meaning since the role of the argument of the property is not specified. Here, following Hamblin's (1987) remarks according to which "the imperative: ⟨X do so-and-so⟩, addressed to *Y*, is really an elliptical plain-predicate imperative, ⟨Bring it about that X does so-and-so⟩", and inspired by Grano's (2015) analysis on uncontrolled infinitives, I propose that a causative predicate, BRING IT ABOUT is coerced to specify the relation between the hearer and the denoted proposition, as in (57b).[8]

Next, as regards the negative imperative in (58a) and the negated controlled complement in (58a), because of a non-RESP nature of the predicates, they are semantically interpreted as equivalent respectively to a third person subjunctive imperative and an uncontrolled subjunctive complement, both of which denote a proposition. To allow them to be appropriately put into the hearer's or the matrix subject's TDL and further repair type-mismatch, I propose, inspired by Richardson's (1985) analysis of negative imperatives, to introduce a monitoring predicate TAKE CARE, rather than a causative predicate BRING IT ABOUT (which is associated with a RESP situation, as argued by Farkas 1988), as in (58b-c) and (59b-c).

(58) a. Do not question somebody's immigration status. (= (1a))
b. TAKE CARE for it not to be the case [that you question **somebody**'s immigration status].
c. [λx. x TAKE CARE $\neg\exists$p [p = λw ($\exists$y. you question y's immigration status in w)]] $\in$ TDL(you)

(59) a. I was trying not to be somebody's bitch for 10 years. (= (1b))
b. I was trying to take care for it not to be the case [that I am **somebody's bitch**].
c. [λx. x TAKE CARE $\neg\exists$p [p = λw ($\exists$y. I am y's bitch in w)]] $\in$ Private TDL(I)

Thus, the property-denoting approach allows us to clarify the motivation for coercion of the monitoring predicate.

[8] I owe the account of this part to Laurent Roussarie's insightful remarks (p.c.).

4.2 The proposition-denoting approach

Next, I review the proposition-denoting approach to imperatives and controlled infinitives. As regards imperatives, Kaufmann (2012) analyzes them as denoting a modalized proposition expressing a necessity. For example, the imperative in (60a) boils down to a modalized proposition, very roughly represented by (60b).

(60) a. Get up! (Kaufmann 2012: 131)
b. ∀w ∈ Modal Background. you get up in w
c. [λw. you get up in w] ∈ the speaker's effective preferences

From a slightly different perspective, Condoravdi & Lauer (2012: 48) claim that imperatives convey that the relevant proposition is among the speaker's effective preferences, as roughly represented in (60c). The negative imperative in (61a) paraphrased by (61b) may be analyzed as involving coercion of the monitoring predicate, as in (62a) (Kaufmann's approach) or as in (62b) (Condoravdi & Lauer's approach). But the motivation of coercion is not clear under their framework.

(61) a. Do not question somebody's immigration status. (= (1a))
b. TAKE CARE for it not to be the case that you question somebody's immigration status.

(62) a. ∀w. you TAKE CARE in w ¬∃p[p = λw′(∃y. you question y's immigration status in w′)]
b. [λw. you TAKE CARE in w ¬∃p [p = λw′(∃y. you question y's immigration status in w′)]] ∈ the speaker's effective preferences

As regards controlled infinitives, while Grano (2015) invokes Portner's property-denoting approach of imperatives, Grano (2018) refers to Condoravdi & Lauer's proposition-denoting theory. According to Grano (2018), (63a) means that John's leaving is among his effective preferences, as in (63b).

(63) a. John intends to leave.
b. [λw. John leaves in w] ∈ John's effective preferences.

Grano (2018) further proposes a detailed analysis of non-controlled infini-

tives. This author first advocates a parallelism between *intend* and its adverbial form *intentionally*, which necessarily induces a RESP relation. He then claims that a RESP relation is not occasionally coerced, as proposed by Grano (2015), but inherently included in the semantics of *intend*. Thus, according to Grano (2018: 616), "intention reports that instantiate syntactic control, like [64a], have essentially the same status as intention reports that do not instantiate syntactic control, like [65a]. [...] In both cases, John names [...] the individual who bears the RESP-relation with *p* [= the complement proposition]. The only difference is that in [64a], in virtue of being a control sentence, John also names the participant associated with the subject of *break the window*, whereas in [65a], Bill takes this place."

(64) a. John intended to break the window. (Grano 2018: 616)
b. John intended to BRING IT ABOUT THAT he break the window.

(65) a. John intended for Bill to break the window. (*ibid.*)
b. John intended to BRING IT ABOUT THAT Bill break the window.

From this perspective, Grano (2018) paraphrases (64a) and (65a) in a parallel way by (64b) and (65b) and calls this approach "COERCION FREE semantics for intention reports". As a counterexample to Grano's (2015) approach, Grano (2018: 624) cites (66), which includes the non-RESP predicate *resemble his father*.

(66) a. John intended to resemble his father.
b. John intended to BRING IT ABOUT THAT he resemble his father.

In order to naturally interpret (66a), the causative predicate is needed to establish a RESP relation between John and his resemblance to his father. But if the controlled complement in (66a) denotes a property, no type-mismatch should take place, and we find no motivation to coerce the causative predicate. On the contrary, the semantics of (66b) is naturally obtained by the coercion-free approach according to which the verb *intend* always introduces the causative predicate by its lexical characteristics. However, if we assume, with Szabolcsi (2010), that the non-RESP nature of the complement hinders obviation, and reanalyzes a controlled infinitive as a subjunctive complement denoting a proposition, the coercion-based approach naturally makes sense of the semantics in (66b), and (66a)

no longer contradicts this approach.

Furthermore, if we apply the coercion-free approach to imperatives, the semantics of the order in (67a) boils down to (67b) including the coerced predicate, just as the Italian third person imperative (68a) involving a subjunctive verbal form, represented by (68b).

(67) a. Get up! (= (65a))
b. [λw. you BRING IT ABOUT THAT you get up in w] $\in$ the speaker's effective preferences

(68) a. *Che venda anche lui!* (Italian)
that come.SBJV also he
'See to it that he comes as well.' (Zanuttini et al. 2012: 1251)
b. [λw. you BRING IT ABOUT THAT he come in w] $\in$ the speaker's effective preferences

But such a parallel treatment obscures the marked status of third person imperatives.[9] In sum, the proposition-denoting approach to imperatives and to controlled infinitives is forced either (i) to invoke, without clarifying its motivation, coercion of a causative or monitoring predicates or (ii) to overgeneralize the use of these predicates.[10]

[9]Kaufmann (2012) effectively puts into question Potsdam's (1998) view according to which the referent of the imperative subject "is not restricted to [...] the addressee", and suggests to analyze examples like (68a) as subjunctive clauses, concluding that "imperative subjects are restricted to refer to [...] the addressee(s)" (Kaufmann 2012: 122).

[10]von Fintel & Iatridou (2017) give two arguments against the proposition-denoting approach to imperatives; (i) indifferent reading, as in (ia); (ii) "non-endorsing *Imperative and Declarative* construction", as in (ib).

(i) a. Go left! Go right! I don't care. (von Fintel & Iatridou 2017: 291)
b. Ignore your homework and you will fail this class. (*idem.* 297)

(ii) a. #You should go left! You should go right! I don't care. (Adapted from *idem.* 294)
b. #You should ignore your homework and you will fail this class.

According to Kaufmann, imperatives are inherently disposed with performatively interpreted *should*. First, (ia) would then be paraphrased by (iia), which is not the case: the first and the second sentences of (iia) are contradictory. Next, a non-endorsing *Imperative and Declarative* construction, like (ib), aims at warning the addressee NOT to perform the action expressed by the imperative (e.g. NOT to ignore the homework in (ia)).

5 Concluding remarks

This study tries to account for a seemingly peculiar narrow scope reading of PPI indefinites under a clause-mate negation in negative imperatives, as in (69a), and in negated controlled infinitives, as in (69b).

(69) a. Do not question **somebody**'s immigration status. (= (1a))
b. I was trying not to be **somebody**'s bitch for 10 years. (= (1b))

Observing that the predicate in these examples denotes a non-RESP situation, I referred to (i) Richardson (1985), who claims that the interpretation of negative imperatives involves coercion of the monitoring predicate TAKE CARE, and (ii) Szabolcsi (2010), who argues that the non-RESP nature of the predicate amounts to projecting a subjunctive clause. Applying these two analyses to the examples in (69a-b), I further claimed that negation may be interpreted in a clause-external position to avoid an incongruous scope relation with respect to the non-RESP complement.[11] The projection of a subjunctive clause, the clause-external interpretation

The paraphrase by (iib) is therefore inappropriate.

According to Condoravdi & Lauer (2012), the speaker's preference meaning is built in the semantics of imperatives. Therefore a paraphrase by *I want to …* should always be possible. This analysis, however, does not hold for (ia) and (ib), as shown in (iiia-b).

(iii) a. #I want you to go left. I want you to go right. I don't care. (*idem.*)
b. #I want you to ignore your homework and you will fail this class.

[11]The distribution of Dutch PPI predicates, like *in een lasting parket* 'in a difficult / awkward situation' seems at odds with the proposed analysis. These predicates are anti-licensed by an external negation, as in (i). But they are licensed in the complement of negated *willen* 'want', as in (ii).

(i) **Ik denke niet dat we in een lasting parket zitten.*
I think not that we in a tough spot sit.
'I don't think that we are in a difficult situation.' (Hoeksema 2018)

(ii) *Ik wil u niet in een lasting parket brengen.*
I want you not in a tough spot bring.
'I don't want to put you in an awkward position.' (*ibid.*)

A possibility is that the narrow scope reading is due to a contrast between *you* and someone else. A detailed examination of these data is a subject of future research.

of negation and the coercion of the monitoring predicate amount to paraphrasing (69a-b) by (70a-b), which allows us to account for the licensing of PPI indefinites.

(70) a. TAKE CARE for it not to be the case [that you question **somebody**'s immigration status].
b. I was trying to TAKE CARE for it not to be the case [that I be **somebody**'s bitch].

I further showed that the property-denoting approach to imperatives and controlled infinitive complements accounts for the motivation of coercion in terms of type-mismatch, while the proposition-denoting approach either cannot account for the motivation of coercion or leads to overgeneralization of causative or monitoring predicates.

Acknowledgments I am grateful to the two editors, especially to Laurent Roussarie, who gave me instructions about how to prepare texts in LaTeX, and even modified a lot of errors included in earlier versions. I also thank the anonymous reviewer who raised many insightful questions and Anne Zribi-Hertz who kindly proofread the final draft. Of course, all errors are mine.

References

Chierchia, Gennaro. 1988. *Topics in the syntax and semantics of infinitives and gerunds* (Outstanding Dissertations in Linguistics). Garland Publishing, Inc.

Condoravdi, Cleo & Sven Lauer. 2012. Imperatives: Meaning and illocutionary force. In Christopher Pinon (ed.), *Empirical Issues in Syntax and Semantics*, vol. 9. 37–58. `http://www.cssp.cnrs.fr/eiss9/eiss9_condoravdi-and-lauer.pdf`.

Condoravdi, Cleo & Sven Lauer. 2016. Anankastic conditionals are just conditionals. *Semantics and Pragmatics* 9. 1–69.

Farkas, Donka. 1988. On obligatory control. *Linguistics and Philosophy* 11. 27–58.

Farkas, Donka. 1992. On obviation. In Ivan A. Sag & Anna Szabolcsi (eds.), *Lexical matters*, 85–109. Stanford, CA: CSLI Publications.

Goncharov, Julie. 2018. Intentionality and polarity sensitivity. Paper presented at the 51st annual meeting of the Societas Linguistica Europaea, Tallinn, August 29– September 01.

Grano, Thomas. 2015. Getting your to-do list under control: Imperative semantics and the grammar of intending. In Thuy Bui & Deniz Özyildiz (eds.), *Pro-*

ceedings of NELS 45, vol. 1. 241–251.
Grano, Thomas. 2018. The logic of intention reports. *Journal of Semantics* 34(1). 587–632.
Hamblin, Charles Leonard. 1987. *Imperatives*. Oxford: Basil Blackwell.
Hoeksema, Jack. 2018. Positive polarity predicates. *Linguistics* 56(2). 361–400.
Horn, Laurence. 2001. *A natural history of negation*. Stanford, CA: CSLI Publications 2nd edn.
Jackendoff, Ray & Peter W. Culicover. 2003. The semantic basis of control in English. *Language* 79. 517–556.
Kaufmann, Magdalena. 2012. *Interpreting imperatives*. Dordrecht: Springer.
Levinson, Dmitry. 2003. Probabilistic model-theoretical semantics for *want*. In Thuy Bui & Deniz Özyildiz (eds.), *Proceedings of SALT*, vol. 13. 222–239.
Portner, Paul. 2007. Imperatives and modals. *Natural Language Semantics* 15. 351–383.
Potsdam, Eric. 1998. *Syntactic issues in the English imperative* (Outstanding Dissertations in Linguistics). Garland Publishing, Inc.
Richardson, John F. 1985. Agenthood and ease. In *Proceedings from the 21st meeting of CLS*. 241–251.
Ruwet, Nicolas. 1990. *Je veux partir* / **Je veux que je parte*: On the distribution of finite complements and infinitival complements in French. In John A. Goldsmith (ed.), *Syntax and human experience*, The University of Chicago Press.
Szabolcsi, Anna. 2004. Positive polarity negative polarity. *Natural Language and Linguistic Theory* 22(2). 409–452.
Szabolcsi, Anna. 2010. Infinitives vs. subjunctives: What do we learn from obviation and from exemptions from obviation? Manuscript. `https://annaszabolcsi.files.wordpress.com/2020/07/szabolcsi_obviation_march_2010.pdf`.
von Fintel, Kai & Sabine Iatridou. 2017. A modest proposal for the meaning of imperatives. In Ana Arregui, María Luisa Rivero & Andrés Salanova (eds.), *Modality across syntactic categories*, 288–319. Oxford: Oxford University Press.
Zanuttini, Raffaella, Miok Pak & Paul Portner. 2012. A syntactic analysis of interpretive restrictions on imperative, promissive, and exhortative subjects. *Natural Language and Linguistic Theory* 30. 1231–1274.
Zu, Vera. 2018. *Discourse participants and the structural representation of the context*. New York: New York University dissertation.

Perception verbs and finite complement clauses

Kalle Müller

Abstract In this paper, I propose analyses for the difference between inferential and reportative readings of perception verbs with finite complement clauses in contrast to perception verbs with bare infinitives. The inferential analysis draws on the new observation that German perception predicates with *dass*-clauses can be accompanied by prepositional *an*-phrases which are highly restricted within this use and are only compatible with inferential readings.

Keywords perception verb · clausal complement · epistemicity · evidentiality

K. Müller, University of Tübingen, kalle.mueller@uni-tuebingen.de

In Christopher Pinon & Laurent Roussarie (eds.), Empirical Issues in Syntax and Semantics 13, 55–79. Paris: CSSP. http://www.cssp.cnrs.fr/eiss13/

1 Introduction

In many languages, perception verbs (henceforth PVs) can take various kinds of complements resulting in often nuanced semantic differences. In (1), this is illustrated for German with a bare infinitive (henceforth BI) and a finite complement clause (henceforth FC).

(1) a. *Margarete sah / hörte Marie den König töten.*
'Margaret saw / heard Mary kill the king.'
b. *Margarete sah, dass Marie den König tötete.*
'Margaret saw that Mary killed the king.'
c. *Margarete hörte, dass Marie den König tötete.*
'Margaret heard that Mary killed the king.'

If the sentences in (1) are interpreted differently, how do the different kinds of perception relate to each other? Do FCs trigger a different kind of perception, one where the PV is interpreted in a metaphorical manner, as in the English phrase *I see*, which is used to convey *I understand*, the phrase *I hear you*, or Goethe's famous dictum *Die Botschaft hör ich wohl* ' 'Tis true, I hear the message'?

While sentences like (1a) have sparked much interested in event and situation semantics where they have played an important role, this paper

focusses on the intricacies of sentences like (1b) and (1c): even though both are syntactically similar, (1c) with *hören* in German or *hear* in English can receive a slightly different reading, one where Margaret has heard a rumor.

The paper draws upon the observation illustrated in (2): in German, PVs with a FC can be accompanied by *an*-PPs.

(2) *Margarete sieht an dem blutigen Messer, dass Maria den König umgebracht hat.*
Margaret sees at the-DAT bloody-DAT knife-DAT that Mary the-ACC king-ACC killed has
'Margaret sees from the bloody knife that Mary killed the king.'

The main questions of this paper are: How can the different readings in (1a) and (1b) be captured in formal semantic analyses? And how do these relate to PVs with a nominal complement as in *I see a cat* or with a bare infinitive as in (1a) above, both of which receive a purely visual reading, but would not be compatible with an *an*-PP in German?

There are basically two options. In the first option, perception is the same for (1a) and (1b), but the latter expresses additional meaning beyond perception. In the second, the act of seeing in (1a) is different from the act of seeing in (1b) because perception differs for objects of different ontological types and seeing a cat is different from seeing a proposition or fact.

The aims of this paper are to argue for the first option for (1b) as well as for the necessity of assuming a different one which is closer to the second option for the most prominent interpretation of (1c), which is one based on hearsay. Furthermore, the present paper aims at presenting further data that analyses of these phenomena need to deal with and to present an analysis for each option on the basis of the data discussed.

This paper is organized in two main sections, together with this introduction and a conclusion. In the first main section, §2, I will lay out the differences between sentences like (1a) and (1b) and the challenges that these different semantic and syntactic restrictions pose to any analysis. They encompass selectional restrictions for predicate type and tense, epistemicity, evidentiality, the abovementioned *an*-PPs, and presupposi-

tions. In the second main section, §3, I will use this data to discuss the problems for previous semantic analyses and propose two separate analyses for the inferential and hearsay interpretations of (1b) and (1c), respectively. They are then compared to each other and discussed with respect to their differences. I will use English examples to illustrate more general points that apply to English and German alike and German examples if the point made might only apply to German.

2 Differences

2.1 Selectional restrictions

Perception verbs with bare infinitives allow only for events and Davidsonian states (Maienborn 2005), but not for statives like in (3).

(3) a. I see her come.
b. *I see her know Margaret.
c. *I see her have red hair.

Perception verbs with *that*-clauses underly no such restrictions, as can be seen in (4).

(4) a. I see that she is coming.
b. I see that she knows Margaret.
c. I see that she has red hair.

Furthermore, bare infinitives have to be co-temporal with the matrix event, which is illustrated in (5), while *that*-clauses can have any tense, as in (6).

(5) a. I see her go on vacation.
b. *I see her have gone on vacation.
c. *I see her will go on vacation.

(6) a. I see that she goes on vacation.
b. I see that she has gone on vacation.
c. I see that she will go on vacation.

2.2 Epistemic load

Perception verbs with finite *that*-clauses carry a certain epistemic load, while perception verbs with bare infinitives are epistemically neutral (Bayer

1986; Hintikka 1969b; Kratzer 2017; Maienborn 2011). In the following example adapted from Maienborn (2011), it is conceivable for (7a) and (7b) that Anna perceived the event of rose cutting and the speaker knows the rose cutting agent to be Heidi, but Anna did not recognize her or maybe does not even know her. However, the sentence in (7c) is only felicitous if Anna recognized Heidi, which is why the part in parentheses cannot be added.

(7) a. Anna saw Heidi cut the roses.
b. Anna saw Heidi cut the roses, but she didn't recognize it was Heidi who cut the roses.
c. Anna saw that Heidi was cutting the roses (#but she didn't recognize that it was Heidi who was cutting the roses).
(Adapted from Maienborn 2011: 808, (11))

This epistemic effect, which requires the subject of the PV to fully recognize the content of the FC, also renders the inference in (9) invalid, whereas the inference in (8) is valid (Kratzer 2017).

(8) *First premise*
Beryl saw Meryl sprinkle the white powder on Cheryl's dinner.
Second premise
The white powder was the most deadly poison.
Conclusion (valid)
Beryl saw Meryl sprinkle the most deadly poison on Cheryl's dinner.

(9) *First premise*
Beryl saw that Meryl sprinkled the white powder on Cheryl's dinner.
Second premise
The white powder was the most deadly poison.
Conclusion (invalid)
Beryl saw that Meryl sprinkled the most deadly poison on Cheryl's dinner. (Taken from Kratzer 2017: (3) and (4))

Only in the case of FCs does the perceiver necessarily know what he is perceiving and the descriptions attributed to the event and to the individual participants in the *that*-clause match the perceiver's epistemic state. In the case of BIs, however, they do not necessarily match the perceiver's

epistemic state.

Hence, the conclusion in (8) is valid because it does not make a statement about whether Beryl knew the white powder to be the most deadly poison. In contrast, the invalid conclusion in (9) entails that Beryl knew about the poison, which is not given by the premises.

2.3 Types of evidence

Crosslinguistically, many languages around the world employ grammatical markers for the category *evidentiality*, marking the evidential grounds the speaker has for asserting a main proposition. Regarding the types of evidential grounds, a central distinction can be made between direct evidence, in which case the speaker has directly perceived an event, and indirect evidence, which divides further into inference and report (Willett 1988).

PVs with a BI refer to direct perception or evidence, as in (10), while PVs with a FC can usually refer to indirect knowledge or evidence, as in (11) and (12) (Aikhenvald 2007; Kratzer 2017), but they can also be based on a direct perception context. In the contexts given in (11) and (12), Margaret has not witnessed the event of Mary killing the king directly. In (11) she draws an inference and in (12) she has been told so.[1] In both contexts the BI is not acceptable.

(10) *Direct perception context*
Margaret, Mary, and a king were present. Mary killed the king and Margaret saw / heard the event.

a. Margaret saw / heard Mary kill the king.
b. Margaret saw / heard that Mary killed the king.

(11) *Inference context*

a. Margaret knew that Mary wanted to kill the king and saw a bloody knife.

[1] As indicated above, the reportative interpretation is much more prominent for *hear* as well as for German *hören*. As a reviewer emphasizes, the inferential is very rare and would have to be forced by context. It is important at this point, however, to acknowledge that these readings exist and that *see* and *hear* as well as there counterparts in German do not completely split into inferential and reportative readings, respectively.

Margaret saw that Mary killed the king.
#Margaret saw Mary kill the king.

b. Margaret knew that Mary went out to kill the king and heard a loud scream from the king as he fell down from the castle wall.
Margaret heard that Mary killed the king.
#Margaret heard Mary kill the king.

(12) *Report context*
Someone told Margaret that Mary killed the king.

a. Margaret heard that Mary killed the king.
b. #Margaret heard Mary kill the king.

2.4 Prepositional *an*-phrases

In German, PVs with a FC can be accompanied by a prepositional *an*-phrase that indicates the source of the inference, as in (13).[2] While *hören* with a FC often receives a reportative interpretation, the presence of such an *an*-phrase forces an inferential reading, as in (14).[3]

(13) a. *An dem blutigen Messer sah Margarete, dass*
at the-DAT bloody-DAT knife[DAT] saw Margaret that
Maria den König umgebracht hat.
Mary the-ACC king[ACC] killed has
'Margaret sees from the bloody knife that Mary killed the king.'

b. *An dem Geräusch hörte Margarete, dass Maria den*
at the-DAT sound[DAT] heard Margaret that Mary the-ACC
König umgebracht hat.
king[ACC] killed has
'Margaret heard from the sound that Mary killed the king.'

[2]They are already attested as indicators of inference in Old High German (Axel-Tober & Müller 2017; Müller 2019).

[3]The word *daran* (see (14)) is a contracted PP headed by *an* with the literal meaning 'at it' or 'at that'.

(14) a. *Es gab einen Schrei. Daran hörte Margarete, dass*
it give.PST a scream there.at heard Margaret that
Marie den König getötet hat.
Mary the-ACC king[ACC] killed has
'There was a scream. From that Margaret heard that Mary killed the king.'

b. *Margarete sprach mit Marion. #Daran hörte Margarete,*
Margaret talk.PST with Marion there.at heard Margaret
dass Marie den König getötet hat.
that Mary the-ACC king[ACC] killed has
'Margaret talked to Marion. From that Margaret heard that Mary killed the king.'

In this use, *an*-phrases are possible with all kinds of perception predicates with *dass*-clauses in German, e.g. *sichtbar / hörbar / spürbar, dass* 'visible / audible / sensible that', but neither with nouns or BIs in a direct perception reading, as seen in (15), nor with belief predicates, as seen in (16).

(15) a. **Daran sehe ich eine Katze.*
there.at see I a-ACC cat[ACC]
'From that I see a cat.'

b. **Daran sehe ich dich kommen.*
there.at see I you.ACC come-INF
'From that I see you come.'

(16) **An dem Messer glaube ich, dass …*
at the-DAT knife[DAT] believe I that
'From the bloody knife I believe that …'

However, there is a certain type of noun that can be accompanied by an *an*-phrase and that is the trope denoting type. In these cases, we observe the same kind of inferential reading – as opposed to an extensional reading as in *Ich sehe eine Katze* 'I see a cat' – that we observe with FCs. Accordingly, the content of the noun can be rephrased as a FC, as illustrated

in (17b).[4,5]

(17) a. *An seinem Gesichtsausdruck sah sie seine Schuld.*
at his expression see.PST she his guilt
'From his expression, she saw his guilt'

b. *An seinem Gesichtsausdruck sah sie, dass er schuldig war.*
at his expression see.PST she that he guilty was
'From his expression, she saw that he was guilty.'

In contrast to pure perception predicates like *sehen* 'see' and *sichtbar* 'visible', *erkennen* 'recognize' and its corresponding modal adjective *erkennbar* generally allow for an *an*-phrase together with a nominal complement, as in (18).

(18) *Daran erkennt / *sieht man einen Betrüger.*
there.at recognize see one a fraud
'Thereby you recognize / *see a con man.'

Yet, *erkennen* cannot take a BI as its complement:

(19) **Ich erkenne dich kommen.*
I recognize you.ACC come-INF
'I recognize you come.'

In the translations above, the most literal counterpart *at* is used for glossing the *an*-phrase, which would not be used in English. Instead, the preposition *from* is accepted by at least some speakers with perception predicates.[6] However, there are differences. In contrast to German *an*, English

[4]The English part of the example in (17) was pointed out to me by a reviewer. It works slightly worse in German, but I think well enough to illustrate the point.

[5]Again, this construction and reading is already attested in Old High German (Axel-Tober & Müller 2017; Müller 2019).

[6]A native speaker suggested this; however, another native speaker did not agree. Another possible candidate is the English preposition *by*, which was also used with *recognize* above. A first enquiry has not found it to be substantially better than *from*, though. This is corroborated by a query a reviewer made in the Corpus of Contemporary Ameri-

from is also compatible with pure predicates of inference like *conclude* or *infer*, as in (20), and even more so than with perception predicates. In German, *(dar)aus* has to be used with such predicates. So, while English *from* seems to generally indicate a source of knowledge with any knowledge predicate, German *an* is restricted specifically to perception predicates.

(20) a. From that I conclude / infer that …
b. *Daran(/Daraus) schließe / schlussfolgere / inferiere ich …

For French, a possible candidate is the preposition *à*, as in (21).

(21) a. *J'ai vu à son air qu'il était fâché.*
I-have seen at his appearance that-he was angry
'I have seen from his appearance that he was angry'
b. *J'ai vu au couteau sanglant que Marie avait tué le roi.*
I-have seen at.the knife bloody that Mary had killed the king
'I saw from the bloody knife that Mary killed the king.'

While more data from other languages is still to be collected, this shows at least for German that the *an*-phrase is tied specifically to PVs in an inferential reading with a FC or a trope denoting noun or certain predicates like *erkennen* 'recognize'. Such predicates are sometimes treated as perception predicates (e.g. Rau 2011). However, they have a similar inferential meaning already incorporated in a way that makes it available for nominal complements without restrictions regarding their semantic type, but prohibits BI complements.

2.5 Entailments and presuppositions

The complements of PVs are usually entailed such that *I see a cat* entails the existence of a cat and *I see someone come* entails an event of someone coming. Regarding FCs, the truth of the complement of *see* is entailed, but the truth of the complement of *hear* often is not. It might be tempting to attribute this difference to a different reliability of the senses such that visual perception is more reliable than auditory perception. However, this

can English (= COCA), who found five examples for *from*, but only one for *by*.

difference is not due to sensory type but actually due to evidence type, as is shown in the following examples. The example in (22) shows that one cannot say that someone saw that it rained if one knows that it did not rain, even if the subject of the main clause believes that it rained and has reasonable grounds for assuming so. The same holds for *hear* in (23a) with an inferential reading. Only in the reportative reading in (23b) the proposition in the subclause may be known to be false.

(22) *Heidi saw a wet street. It didn't rain.*
#Heidi saw that it rained.

(23) a. *Heidi heard the door open downstairs and thought it was her father. But it was her mother.*
#Heidi heard that her father came home.
b. *Someone told Heidi that her friend Peter was a spy. But he wasn't.*
Heidi heard that Peter was a spy.

Turning to the question of presuppositions, the set of diagnostics called *family of sentences* by Chierchia & McConnell-Ginet (1990) is most established. One of these diagnostics is embedding under questions. In the following German examples, the question in (24a) does not imply that anyone actually came; however, the question in (24b) is only felicitous in a context where the speaker knows that someone did. Hence, the BI is not presupposed, whereas the FC is presupposed.[7]

(24) a. *Hast du jemanden kommen sehen?*
have-2.SG.PRS you someone come-INF seen
'Did you see anyone come?'

[7]However, one might imagine a court room situation, where a witness is questioned whether he actually *saw* that something happened, implying that if he did not see it, it might not have happened at all. Moreover, in the wake of Simons et al. (2010), the traditional notion of presupposition has come into question especially for complements of factive verbs. Tonhauser et al. (2018) report that verbs do not behave in strict accordance with the distinct categories of factive and non-factive, but factive presuppositions are subject to gradiency. Nonetheless, even in the experiments reported in Tonhauser et al. (2018), *see* is usually on the upper end of the scale. The courtroom example seems to be a special case and presuppositions to be more prone to cancelling in interrogations. In German, a non-presuppositional complement of a PV with *dass* has to be marked with the subjunctive mood.

b. *Hast du gesehen, dass jemand gekommen ist?*
have-2.SG.PRS you seen that someone come-PRF is
'Did you see that someone came?'

3 Analyses

Most analyses have focused on the combination of perceptions verbs with bare infinitives, which have been widely discussed in event and situation semantics (see e.g. Barwise 1981; Barwise & Perry 1983; Higginbotham 1983; 1999; Vlach 1983). Rau (2011) provides an account for PVs with FCs which relies on a relation between two events, the SEE-event and the event described in the FC. Hintikka (1969b), on the other hand, extends his famous analysis for propositional attitudes from predicates of knowledge and belief to perception predicates.

Assuming a Davidsonian event-based account for BIs can explain the restrictions from §2.1 if we posit that not all verbs describe a Davidsonian event (Maienborn 2005; 2011) and that for a perceptual relation to hold between two individuals, there must be a temporal overlap between both. Assuming a similar account for FCs as well not only faces the problem of epistemic load described in §2.2, which might be remedied by introducing an additional epistemic function, but also needs to explain why these restrictions hold in one case but not the other.

In (25), a pure propositional attitude analysis along the lines of Hintikka (1969a;b) is shown.

(25) a. Margaret saw / heard that Mary killed the king.
b. For every w' such that w' is compatible with what Margaret saw / heard in w, Mary killed the king in w'.

This analysis captures the lack of restrictions for the predicate within the FC as well as the accompanying epistemicity. Since this analysis has been deemed inadequate for the bare infinitive constructions by most authors, the question arises whether there is any link between a PV with a nominal argument or a BI and the same PV with a FC.[8] In the case of *hear* / *hören*, a general analysis describing a set of worlds which is compatible

[8]However, Saarinen (1983) defends a Hintikkan treatment also for BI complements.

with what someone has heard is compatible with both inferential and reportative readings. However, it is not clear how these readings would be distinguished. Moreover, it is not explained how these readings relate systematically to the availability of an *an*-phrase together with the presence or absence of factive presuppositions.

In the following I will first elaborate on the role of the perceiving subject and the perceived object for an analysis. Based on the data discussed I will then propose an analysis for all inferential readings of PVs with a FC and explain its relationship to standard extensional uses of PVs. Finally, I will show a different analysis for the reportative readings of *hear* / *hören* by extending the analysis given by Kratzer (2016) for verbs of belief and verbs of speech.

3.1 Subjects and objects of perceptions

The perceiving subject The perceiving subject of the main clause is responsible for the conclusion described by the FC. Imagine the context in (26).

(26) *A woman consults Holmes and Watson. Both see the same things. For instance, both see elm leaves on her boots. Both recognize they are elm leaves. Holmes knows that elms grow only on East End next to the river Thames. Watson does not know this.*

a. Holmes saw that she was coming from East End.
b. #Watson saw that she was coming from East End.

Though both men see exactly the same things, only the sentence in (26a) would be acceptable in this context. In a strictly Hintikkan analysis like (25), it is not clear why the FC should not be compatible with what Watson saw. In fact, it is compatible with what both saw, but not with what both know.

This example shows that these PVs do not rely solely on the object of perception nor on general rules about the world, but they also interact with the previous knowledge of the attitude holder about the world. Only Holmes knows that the elm leaves grow on East End, so only he knows them to be evidence for the embedded statement. Furthermore, what can be evidence for someone for a certain conclusion varies across worlds, since there would be possible worlds in which elms do not grow on East

End, as well as possible worlds in which they do but Holmes does not know about it.

The perceived object The perceived object of the inferential readings is not an intensional object itself. In reportative readings, *hören* can take a noun like *Gerücht* 'rumor' as a complement. For inferential readings of *sehen* or *hören*, nouns meaning 'fact' or 'conclusion' cannot act as a complement of the PV. However, English *see* behaves differently. In English, *fact* is a perfectly adequate nominal complement.[9]

(27) a. *Heidi hörte das Gerücht, dass …*
Heidi hear-PST the rumor that
'Heidi heard the rumor that …'

b. **Heidi sah die Tatsache / Schlussfolgerung, dass …*
Heidi see.PST the fact conclusion that
'Heidi saw the fact / *conclusion that …'

While the previous paragraph might suggest that the FC describes the content of a conclusion made by the subject of the main clause, this conclusion does not seem to be present as an individual argument in German or, at least, it does not act directly as the internal argument of the perception predicate.

Furthermore, the perceived object is not necessarily described by any part of the FC. It is, however, necessarily described by the *an*-phrase. This is illustrated in (28) and (29). Out of (29a-e), only (29e) would always be a false statement in the context.

[9] I am indebted to a reviewer who made me aware of this fact. In fact, he found that *fact* is the third most frequent noun returned by the query *[see] the [n*] that* in the COCA. In the German Reference Corpus (= DeReKo), there was only one example of this type, which is part of a speech by Angela Merkel. In German, the combination of *sehen* and *Tatsache* only occurs within constructions that are usually considered small clauses but are different from BIs in German, e.g. *Ich sehe die Sache nicht so dramatisch.* 'I don't see the matter as dramatic.'

(28) *You know that every time Hannah is working, Anna hangs out a red ribbon for her lover.*
An dem roten Band sehe ich, dass Hannah arbeitet.
at the-DAT red-DAT ribbon.[DAT] see I that Hannah works
'I see (/ know) from the red ribbon that Hannah is working.'

(29) a. I don't see Hannah work.
b. I don't see anyone work.
c. I don't see Hannah.
d. I don't see anyone.
e. #I don't see a red ribbon.

This shows that at least in German, inferential readings of PVs with a FC still involve regular perception, only what is perceived is not expressed in the FC, but can be expressed optionally with the German *an*-phrase.

3.2 Inferential readings

In the spirit of event semantics, I presuppose that standard readings of PVs with noun phrases or BIs involve an extensional relation between three individuals, the perception event, the perceiving subject and the perceived object, which may be another event.[10] What is its relationship to the inferential readings of a PV with a FC?

As shown above, the perceiving subject is necessarily the one who draws the conclusion described by the FC. It was also argued for German that this construction still involves an act of actual sensory perception, only that the object which is perceived in this act is independent of the FC, but can be overtly realized. We can thus conclude that an analysis of the inferential readings in German needs to include a standard perception predicate as described above and an additional part which introduces a conclusion and relates it to the perception event. This part is provided in (30) as a function from a piece of evidence, an attitude holder, and a world of evaluation to a set of worlds.

(30) $I_{inference}(y)(x)(w) = \{w' \in W : w'$ is compatible with the conclusions x draws (wrt. to x's knowledge in w) from y in $w\}$

[10] Since this is not the focus of this paper, I will leave the discussion aside. The interested reader is referred to the articles mentioned above.

This inferential function is not supposed to replace a standard perception predicate, but to augment it. I assume a predicate SEE_1 for standard readings of *sehen* with noun phrases and BIs which encodes an extensional eventive relation between a perceiving subject and a perceived object. It shares two arguments with the inferential function such that the perceived object functions as a piece of evidence from which a conclusion can be drawn in a world and the perceiving subject is the attitude holder to draw the conclusion from it. The presupposition is rendered by a partial function which requires the embedded proposition to be true not only in all worlds w' of the inferential function, but also in w. The complete analysis for *sehen* with a FC can be seen in (31):

(31) $\lambda p \lambda x \lambda e \lambda w : p(w).\exists y[\text{SEE}_1(y)(x)(e)(w) \wedge$
$\forall w'[w' \in I_{nference}(y)(x)(w) \rightarrow p(w')]]$

Returning to our initial problems posed by the differences described in §2, the presupposition and the evidence type is wired directly into the analysis. The valid and invalid conclusions of (8) and (9), respectively, can be explained as well. Given the extensional analysis for PVs with a BI we expect premises and conclusions to be evaluated with respect to the same parameters. So, if x sees an event involving an individual z in w in the first premise and the same individual z is poison in w in the second premise, we can conclude that x sees an event involving z and z is poison in w, given that x and z are assigned the same values.[11] For PVs with a FC, however, the inferential function in the conclusion would require an individual z to be poison not only in w as given by the second premise but in w' as well, the epistemic world of the perceiver, which is not given by the premises. The selectional restrictions exemplified in §2.1 can be explained by the assumptions that stative predicates like *know* or *have red hair* do not provide a Davidsonian event argument for the perception predicate (Maienborn 2005; 2011). Though we maintain the same perception predicate as a part of our analysis for the inferential interpretation, the same restrictions

[11] For the sake of simplicity, I neglect the assignment function in my examples. Dynamic semantics would be another possibility to achieve this. In any case, it is obvious that the conclusion can only hold if the white powder in the first premise is the same as in the second premise.

do not hold, since it does not require the perceived object to be the event or any individual of the embedded clause.

In the following examples (32)–(35), I spell out the analyses for *sehen* with a noun phrase, a BI, a FC, and a FC with an additional *an*-phrase (where tense is neglected).[12]

(32) a. *Margarete sah einen Marder.*
'Margaret saw a marten.'
b. $\lambda w \exists e \exists y [\text{SEE}_1(y)(\text{Mrgt})(e)(w) \wedge \text{marten}(y)(w)]$

(33) a. *Margarete sah Marie den König töten.*
'Margaret saw Mary kill the king.'
b. $\lambda w \exists e \exists e' [\text{SEE}_1(e')(\text{Mrgt})(e)(w) \wedge \text{kill-theking}(\text{Mary})(e')(w)]$

(34) a. *Margarete sah, dass Marie den König getötet hatte.*
'Margaret saw that Mary had killed the king.'
b. $\lambda w : \exists e' [\text{killed-theking}(\text{Mary})(e')(w)].$
$\exists e \exists y [\text{SEE}_1(y)(\text{Mrgt})(e)(w) \wedge$
$\forall w' [w' \in I_{nference}(y)(\text{Mrgt})(w) \rightarrow$
$\exists e'' [\text{killed-theking}(\text{Mary})(e'')(w')]]]$

(35) a. *An einem blutigen Messer sah Margarete, dass Marie den König getötet hatte.*
'From a bloody knife Margaret saw that Mary had killed the king.'
b. $\lambda w : \exists e' [\text{killed-theking}(\text{Mary})(e')(w)].$
$\exists e \exists y [\text{SEE}_1(y)(\text{Mrgt})(e)(w) \wedge \text{bloody}(y)(w) \wedge \text{knife}(y)(w) \wedge$
$\forall w' [w' \in I_{nference}(y)(\text{Mrgt})(w) \rightarrow$
$\exists e'' [\text{killed-theking}(\text{Mary})(e'')(w')]]]$

One of the main benefits from such an analysis for the inferential construction is that its relationship to the other reading becomes perspicuous. It is not necessary to involve metaphoric mechanisms to change what seeing means in this context or to resort to fact perception as a different kind of perception. In contrast, the perception part stands unaltered but is related to an epistemic part by sharing arguments.

[12] In (34) and (35), e' in world w has to be counterpart-identical with e'' in world w'. I ignore this issue in these analyses.

3.3 Reportative readings

Verbs of hearing with finite clauses, however, are generally ambiguous between inferential and reportative interpretations. I argue that reportative interpretations arise if the internal argument is satisfied by an informational object with which the *that*-clause is associated.

Relativization, modal anchors and content functions Generally, *that*-clauses occur as complements for a variety of verbs, e.g. verbs of speech, thought, and belief. However, *that*-clauses can also complement or modify nouns like *idea*, *possibility*, or *thought*, denoting their content. As a consequence, we can associate propositional content with informational objects, i.e. objects like ideas, stories or rumors. Such informational objects may be modeled by taking an individual argument as a modal anchor (Hacquard 2006). Modal functions can project sets of possible worlds from such a modal anchor, i.e. from an individual argument. Kratzer (2016) argues that the contribution of mood is the introduction of a free variable ranging over domain projection functions, as in (36).

(36) $[\![Mood]\!] = \lambda p \lambda x \forall w'[w' \in f(x) \rightarrow p(w')]$ (Kratzer 2016)

One possible assignment for $f(x)$ in (36) is a content-related domain projection function with defeasible normalcy conditions which can be introduced by the *that*-clause via relativization relating x to an individual in the main clause, as in (37) (Kratzer 2016).

(37) $\lambda x.C_{ontent}(x) = \{w' \in W : w'$ is a world that is compatible with the content of $x\}$
Undefined if x doesn't have intensional content (Kratzer 2016)

In (38), this is exemplified with a noun phrase.

(38) a. the [rumor that Ortcutt is a spy]
b. $\lambda x \lambda w[\text{rumor}(x)(w) \wedge \forall w'[w' \in C_{ontent}(x) \rightarrow \text{spy}(\text{Ortcutt})(w')]]$
(Adapted from Kratzer 2016)

Syntactically, it has been argued that what were previously considered to be sentential complements actually involve relative structures (Arsenijevic 2009; Kayne 2008; Moulton 2009; diachronically for German: Axel 2009; Axel-Tober 2017). Kratzer (2016) draws upon these ideas and proposes an

analysis like (39) for speech act verbs like *say*. In (39), the speech event produces an argument which acts as a modal anchor for the content function in (37) relating it to the speech content.

(39) a. Margaret [says that Mary killed the king].
b. $\lambda x \lambda e \lambda w \exists y [\text{say}(y)(x)(e)(w) \wedge$
$\forall w' [w' \in C_{content}(y) \rightarrow \exists e' [\text{killed-theking}(\text{Mary})(e')(w')]]]$

In other words, if I say some*thing*, this *thing* is an informational object with a propositional content. The *that*-clause acts as a relative clause restricting its content to a specific set of worlds.

Application to reportative If verbs of saying relate to an object associated with the content of the speech act, it can be assumed that this object is received by the addressee. Hence, the analysis from (39) naturally carries over to reportative readings of *hear*. In (40) Margaret hears something which is related to the worlds in which Mary killed the king via the content function.

(40) a. Margaret [heard that Mary killed the king].
b. $\lambda x \lambda e \lambda w \exists y [\text{hear}(y)(x)(e)(w) \wedge$
$\forall w' [w' \in C_{content}(y) \rightarrow \exists e' [\text{killed-theking}(\text{Mary})(e')(w')]]]$

While this analysis is intensional, it is not epistemic, since it does not relate to what Margaret knows, believes, or thinks about the world. Accordingly, there is no argument or parameter for any attitude holder present in the content function.

In §2.5, it was noted that the truth of the embedded proposition is not entailed or presupposed and no presupposition or entailment follows from this analysis. One can expect some variation in the degree to which the received information is probable or reliable. This can be explained by the fact that the informational object x in (40) is not further qualified. Pragmatic enrichment may lead to different specifications of x. If x is taken to be a rumor as in *She heard the rumor that …*, the associated proposition will receive more doubt than it would if x is taken to be news.

Both *news* and *rumor* are compatible among other nouns as direct objects of *hear*. In contrast, the object of perception is never realized as the direct object of an inferentially interpreted PV, as argued above, but indi-

rectly with the preposition *an* in German. Hence, when the object of perception is realized, existentially bound and described by an *an*-phrase, the listener / reader knows that it can no longer be an object associated with the *that*-clause and a reportative interpretation is excluded.

3.4 Comparison between both analyses

Two very different analyses have been proposed for a seemingly single phenomenon, i.e. PVs with finite complement clauses, one analysis for inferential and another for reportative readings. Aside from the difference in evidence type, this has been motivated by the accompanying differences in presupposition, the prepositional phrase, and the availability of abstract nouns as direct objects to which the finite complement clause would be related.

Furthermore, the two analyses extend to very different phenomena. The reportative analysis carried over from Kratzer (2016) extends to a very wide range of predicates. It also extends to specific cases of *see*, as in *Did you see the memo that* . . . , but in this case, the interpretation is not much different from *Did you get the memo that* . . . and analogue to *I sent out a memo that* . . . . It can be viewed as a special case of the reportative interpretation and can occur with various visual objects that carry information like memos, notes, and billboards. It is not specific to PVs.

The inferential interpretation, however, is very specific. Here, the prepositional *an*-phrases from German constitute an important piece of evidence. As shown, they do not generally occur with PVs in all syntactic constructions. Neither do they generally occur with verbs with clausal complements. And they do not generally occur with PVs with clausal complements, which is evident from the reportative constructions. They appear with PVs with clausal complements with an inferential reading. This indicates that something must have changed for the PV which is particular to this exact combination. What has changed is that additional meaning components have been introduced.

So far, an extensional core of visual perception could be upheld, while the intensionality carried by the FC was integrated by introducing additional meaning components. However, since – as argued – the *an*-phrase relates to the actual object of perception, in contrast to the FC, it has to be noted that this prepositional phrase can relate to a propositional ob-

ject as well. This can be achieved by embedding a second FC under the preposition with an integrated pronominal, as in (41).

(41) *Daran, dass die Aktienkurse gestiegen sind, sehe ich, dass es*
there.at that the stock.prices risen AUX see I that it
dem Konzern gut geht.
the corporation well goes
'From the fact that stock prices went up I can tell that the corporation is doing well.'

Examples like this one can be analyzed by combining both analyses. The first *dass*-clause in (41), however, does not relate content to an information bearing object, rather it refers to a fact. Remember that according to Kratzer (2016) mood introduces a free variable for domain projection functions and $C_{content}(x)$ is only one possible assignment. Kratzer (2016) also argues for a factual domain projection $f_{act}(x)$ mapping x onto a set of possible worlds which have a counterpart of x. One argument for not employing the same analysis for both *dass*-clauses in (41) is given in (42): the noun *Tatsache* cannot act as the direct object of *sehen* in (42a), but it can be the complement of *an* in (42b).

(42) a. **Ich sehe die Tatsache, dass …*
I see the fact that
'I see the fact that …'

b. *An der Tatsache, dass …, sieht man (*die Tatsache), dass …*
at the fact that see one the fact that
'From the fact that … you can tell that …'

Using these tools, we can analyze (41) as in (43).[13]

(43) $\lambda w : \exists e_1[\text{do-well}(\text{the-corp})(e_1)(w)].\exists e \exists y[\text{PERCEIVE}(y)(\text{ego})(e)(w) \wedge$
$\forall w'[w' \in f_{act}(y) \rightarrow \exists e'[\text{go-up}(\text{stock-prices})(e')(w')]] \wedge$
$\forall w''[w'' \in I_{nference}(y)(\text{ego})(w) \rightarrow \exists e''[\text{do-well}(\text{the-corp})(e'')(w'')]]]$

[13] Though this is not entirely satisfying, only the presupposition of the second *dass*-clause is modeled as a partial function because the function f_{act} is supposed to yield the factivity of the first one.

Note that in this case a more general predicate PERCEIVE is employed instead of SEE, because *sehen* may no longer refer to visual perception in this example. It is important to acknowledge, however, that this change from visual perception in a literal sense to a more general, metaphorical kind of perception is not triggered by *sehen* taking a FC, but by *an* taking a FC. In German, the non-metaphoric restriction for *sehen* is not absolute, but quite strong. Consequently, *sehen* with a FC usually involves actual visual perception.

In English, however, this is different. Beside the fact that English *see* can take *the fact that …* as a complement, there are cases like *I see*, *I can see that*, or *I see your point*, where no visual perception need be involved at all. Does this mean that the proposed analysis for inferential readings applies only to German?

If we compare a non-visual example like (44) with an example that suggests a visual context, as in (45), there seems to be a crucial difference. In (44), the problem might be that the person addressed has never even thought about it or that they simply refuse to acknowledge the truth. The sentence does not express an inference or a conclusion but rather an insight and an awareness. One might offer certain situations as evidence for the change, but the problem is not that person addressed lacks the evidence, but rather that they lack the insight. The example in (45), on the other hand, would be adequate in a situation where the person addressed is already looking around and they are asked to keep looking. In this situation, it seems odd to insert *the fact* before the FC.

(44) You still can't see the fact that you have changed.

(45) You still can't see that I have cleaned the kitchen.

This suggests that, though English behaves differently in some respects, the proposed analysis for inferential interpretations in German applies to English as well. But since English has a broader variety of metaphorical, non-sensory readings for PVs without a FC, it does so for PVs with a FC as well and these might often be the preferred readings. Depending on the particular example, a more adequate analysis could be achieved by replacing the visual predicate SEE with a cognitive one in the inferential analysis or by taking the structure of the reportative analysis but employing a fact

projection function rather instead of the content projection function and adjusting the visual predicate for metaphorical processes.

4 Conclusion

In this paper, two different semantic analyses of perception verbs with finite complement clauses have been proposed, one for inferential readings and one for reportative readings. While the former are mostly present with *sehen* / *see* and the latter mostly with *hören* / *hear*, both readings are generally available for both verbs. In case of *see*, the reportative reading arises in contexts where the subject has read the information somewhere. The reportative analysis extends the analysis by Kratzer (2016) for verbs of speech, where a function projected from a modal anchor maps intensional content onto an information bearing individual.

It has been argued that this analysis cannot be carried over to inferential readings. Instead, I have proposed an analysis where a basic perception predicate may stay unaltered and is augmented by an inferential function from evidence, an attitude holder, and a world of evaluation to a conclusion. Both parts, the perception predicate and the inferential function relate to each other by sharing their arguments. The perceived object acts as evidence from which a conclusion can be drawn and the perceiving subject is the one drawing the conclusion. That is, you do not see a fact, but you see something and draw a factual conclusion from this object. By augmenting the perception predicate, a common core can be upheld and the relationship between readings with nominal complements and readings with finite complements becomes transparent.

The paper has also argued that only perception predicates and only inferential readings of these can be accompanied by a PP introduced by the preposition *an* in German. It is clear that something must have changed for perception verbs with finite complements. The question arises how these additional meaning components are introduced and how the change in valency is achieved. One option is to assume a silent derivational morpheme augmenting its meaning with the inferential function, changing its syntactic and semantic argument structure. Other verbs like *erkennen* 'recognize' have already incorporated this part and can exhibit intensionality with nouns as well, but cannot take BI complements. At this point, however, I must leave this issue to further research.

There are more syntactic configurations and more interpretations than the ones considered in this paper. One of these are intensional interpretations of *see* with nominal complements. An illustration of this would be the utterance of *I see a cat* in a situation where you are describing a picture by Piet Mondrian or where you are performing a Rorschach test. In these cases, the *an*-phrase is not available, which separates this phenomenon from the ones discussed.

Acknowledgments I would like to thank Angelika Kratzer, Claudia Maienborn, Sebastian Bücking, Sam Featherston, Christopher Gauker, my colleagues from C6, the audiences of CSSP 2019 as well as SPE11, and the reviewer and editors for valuable discussions and helpful comments. This research was partly funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 75650358 – SFB 833.

References

Aikhenvald, Alexandra Y. 2007. Information source and evidentiality: What can we conclude? *Rivista di Linguistica* 19(1). 209–227.

Arsenijevic, Boban. 2009. Clausal complementation as relativization. *Lingua* 119(1). 39–50.

Axel, Katrin. 2009. Die Entstehung des *dass*-Satzes – ein neues Szenario. In Veronika Ehrich, Chirstian Fortmann, Ingo Reich & Marga Reis (eds.), *Koordination und Subordination im Deutschen* (Linguistische Berichte Sonderheft 16), 21–41. Hamburg: Buske.

Axel-Tober, Katrin. 2017. The development of the declarative complementizer in German. *Language* 93(2). e29–e65.

Axel-Tober, Katrin & Kalle Müller. 2017. Evidential adverbs in German: Diachronic development and present-day meaning. *Journal of Historical Linguistics* 7(1). 9–47.

Barwise, Jon. 1981. Scenes and other situations. *Journal of Philosophy* 78(7). 369–397.

Barwise, Jon & John Perry. 1983. *Situations and attitudes*. Cambridge, MA: The MIT Press.

Bayer, Joseph. 1986. The role of event expressions in grammar. *Studies in Language* 10. 1–52.

Chierchia, Gennaro & Sally McConnell-Ginet. 1990. *Meaning and grammar*. Cambridge, MA: The MIT Press.

Hacquard, Valentine. 2006. *Aspects of modality*. Cambridge, MA: MIT dissertation.

Higginbotham, James. 1983. The logic of perceptual reports: An extensional alternative to situation semantics. *The Journal of Philosophy* 80. 100–127.

Higginbotham, James. 1999. Perceptual reports revisited. In Kumiko Murasugi & Robert Stainton (eds.), *Philosophy and linguistics*, 11–33. Boulder, CO: Westview.

Hintikka, Jaakko. 1969a. Semantics for propositional attitudes. In *Models for modalities: Selected essays*, 87–111. Dordrecht: D. Reidel Publishing Company.

Hintikka, Jaakko. 1969b. On the logic of perception. In *Models for modalities: Selected essays*, 151–183. Dordrecht: D. Reidel Publishing Company.

Kayne, Richard. 2008. Antisymmetry and the lexicon. *Linguistic Variation Yearbook* 8. 1–31.

Kratzer, Angelika. 2016. Evidential moods in attitude and speech reports. Slides of a talk given at University of Siena, University of Pennsylvania, and the University of Connecticut. Available at `https://works.bepress.com/angelika_kratzer/10/`.

Kratzer, Angelika. 2017. Situations in natural language semantics. In Edward N. Zalta (ed.), *The Stanford encyclopedia of philosophy*, Metaphysics Research Lab, Stanford University winter 2017 edn.

Maienborn, Claudia. 2005. On the limits of the Davidsonian approach: The case of copula sentences. *Theoretical Linguistics* 31(3). 275–316.

Maienborn, Claudia. 2011. Event semantics. In Claudia Maienborn, Klaus von Heusinger & Paul Portner (eds.), *Semantics: An international handbook of natural language meaning*, vol. 1, 802–829. Berlin: Mouton de Gruyter.

Moulton, Keir. 2009. *Natural selection and the syntax of clausal complementation*. Amherst, MA: University of Massachusetts dissertation.

Müller, Kalle. 2019. *Satzadverbien, Evidentialität und Non-at-issueness*. Tübingen: University of Tübingen dissertation.

Rau, Jennifer. 2011. *Semantic and syntactic differences between finite and infinitival complements in German*. Tübingen: University of Tübingen dissertation.

Saarinen, Esa. 1983. On the logic of perception sentences. *Synthese* 54(1). 115–128.

Simons, Mandy, Judith Tonhauser, David Beaver & Craige Roberts. 2010. What projects and why. In *Proceedings of SALT 20*. 309–327.

Tonhauser, Judith, David Beaver & Judith Degen. 2018. How projective is projective content? Gradience in projectivity and at-issueness. *Journal of Semantics*

35(3). 495–542.
Vlach, Frank. 1983. On situation semantics for perception. *Synthese* 54. 129–152.
Willett, Thomas. 1988. A cross-linguistic survey of the grammaticalization of evidentiality. *Studies in Language* 12(1). 51–97.

Distributional semantics and the conceptual foundations of verb meaning: How neural word embeddings memorize the unaccusative hypothesis

Tillmann Pross

Abstract I investigate whether and how neural word embeddings can be understood to encode not only idiosyncratic aspects of word meaning but also the kind of general and abstract concepts that are central to theoretical approaches of lexical semantics. To this end, I compute the difference between general-purpose embeddings of intransitive verbs, and task-specific embeddings of the same verbs that capture their similarity according to the unaccusative hypothesis. I show that the difference that retraining makes is neither trivial nor random but captures surpisingly well the cues for unergativity and unaccusativity that have been proposed in the theoretical literature. The study presented thus suggests that word embeddings may provide a novel and empirically grounded perspective on the conceptual underpinnings of verb meaning.

Keywords unaccusativity · word embedding · thematic role · lexical semantics

T. Pross, University of Stuttgart, tillmann.pross@ims.uni-stuttgart.de

In Christopher Pinon & Laurent Roussarie (eds.), Empirical Issues in Syntax and Semantics 13, 81–107. Paris: CSSP. http://www.cssp.cnrs.fr/eiss13/

1 Introduction

A theory of lexical representation is key to compositional theories of the meaning of phrases and sentences. One of the main challenges for the development of such a theory is the identification and specification of the primitive meaning components out of which verb meanings are constructed. While such a theory of the primitive components of verb meaning is of central importance to the development of a general theory of the lexical semantics of verbs – "[t]he important theoretical construct is the notion of meaning component, not the notion of verb class" (Levin 1993: 18) – the determination of the components of verb meaning is a central, yet unsolved research problem in both theoretical and computational approaches to verb meaning. By means of a case study on the meaning com-

ponents of intransitive verbs, this paper ties in with recent work (such as Asher et al. 2016; McNally & Boleda 2017; Pross et al. 2017, see e.g. Boleda 2020 for a review of the current state of the art) that aims to show that this situation is to the detriment of both theoretical and computational approaches to lexical semantics. Accordingly, the goal of the present paper is to show that striving for a combination of theoretical and computational approaches to lexical semantics is a natural way to deal with problems which can not be solved by introspection or corpus statistics alone.

1.1 Distributional vs. decompositional similarity

According to the distributional hypothesis, words that occur in similar contexts tend to have similar meanings (Firth 1957; Harris 1954; Miller & Charles 1991; Turney & Pantel 2010; Clark 2015). Since this hypothesis can be operationalized through approximation of word meaning "by the patterns of co-occurrence of words in corpora from statistical semantics" (Baroni et al. 2014: 241), the distributional hypothesis has become the main starting point for current research in computational semantics. In traditional count-based approaches to the distributional similarity of words, a word's meaning is typically a point in a high-dimensional vector space, where the dimensions of the vector correspond to context items, e.g. co-occurring words, and the coordinates of the vector are defined by the strength of these context items, e.g. co-occurrence counts. Contextual – or more precisely, distributional – similarity then becomes proximity of word meanings in the vector space. In this paper, I use distributional semantic models extracted from corpus data with neural network architectures that are referred to as "word embeddings" (Mikolov et al. 2013; Pennington et al. 2014). The relevant differences between word count models and word embeddings are that (i) at the quantitative level, word count models are high-dimensional while embedding models are low-dimensional and (ii) qualitatively, the dimensions of count models correspond to actual words, while the dimensions produced by embedding models can be thought of as soft clusters of context items that do not correspond to actual words (Levy & Goldberg 2014).

In a prototypical distributional model of word meaning, words that are most similar in meaning to the verb *laugh* are *funny*, *cry* and *tear*. The co-occurrence contexts of *laugh* thus reflect that a laughing event usu-

ally takes place in a funny situation, and often goes along with tears and crying. But the distributional similarity of words is not the only way in which words, and in particular verbs, can be judged to be similar. For example, according to the so-called unaccusative hypothesis (Perlmutter 1978), verbs like *laugh*,*work*, *sleep* and *run* are semantically similar in that they share the lexical entailment of an agentive meaning of their single argument – as opposed to verbs like *stumble*, *die*, *arrive* and *crystallize*, which are semantically similar because they entail a patient-like meaning of their argument.[1] But since verbs like *laugh*, *work*, *sleep* and *run* are highly unlikely to co-occur in similar contexts, unergative verbs cannot be semantically similar according to the distributional hypothesis.

In contrast to the distributional similarity of words, which can be read off the surface distribution of words in a corpus, the semantic similarity of intransitive verbs is semantically determined at a covert level of lexical representation and reflected by syntactic properties (Levin & Rappaport Hovav 1995) like e.g. auxiliary selection in the present perfect in German (Wunderlich 1985; Grewendorf 1989). Unergative verbs (1a) select HAVE in the present perfect, whereas unaccusative verbs (1b) select BE.

(1) a. *Maria hat gelacht.*
Maria HAVE laugh
'Maria has laughed.'
b. *Maria ist gestolpert.*
Maria BE stumble
'Maria has stumbled.'

[1]While there seem to be robust intuitions that there is a semantic difference between intransitive verbs that entail an agent-like meaning of their single argument and intransitive verbs that entail a patient-like meaning of their single argument, the grammatical realization of this intuitive dichotomy has been subject to debate, in particular with respect to the question for whether or not the unaccusative hypothesis pertains to a binary distinction or not, see e.g. Sorace (2000) for a gradient analysis of the unaccusative hypothesis based on data from auxiliary selection and Levin & Rappaport Hovav (1995) for a more general overview of the problem set. In this paper I make the more or less standard but simplifying assumption that the unaccusativity hypothesis makes a claim about a split in the syntax and semantics of intranstive verbs. Whether or not the methodology presented in this paper supports a specific view on the nature of the unaccusative hypothesis is a question I leave to further research.

Semantically, unaccusativity is determined by an intuition Dowty (1991: 605) characterizes as follows: "intransitive predicates argued to be unaccusative on syntactic grounds usually turned out to entail relatively patient-like meanings for their arguments [...], while those argued to be syntactically unergative were usually agentive in meaning." Levin & Rappaport Hovav (1995: 91) propose that unergative verbs describe internally caused events in which "inherent properties of the single argument like will, volition, emotion or physical characteristics are "responsible" for bringing about the eventuality" that the verb describes. Unaccusative verbs describe externally caused events for which an agent, an instrument, a natural force or a circumstance has "immediate control over bringing about the eventuality described by the verb" (Levin & Rappaport Hovav 1995: 92).

The semantic distinction between unergative verbs like *laugh* and unaccusative verbs like *arrive* is often represented through lexical decompositions as in (2) (see e.g. Levin & Rappaport Hovav 1995).

(2) a. laugh $\rightsquigarrow$ DO laugh
b. arrive $\rightsquigarrow$ BECOME arrived

In (2), the lexical meaning of the verbs *laugh* and *arrive* is decomposed into recurrent (word-overarching; blue) and idiosyncratic (word-specific; red) meaning elements.

Conceptual constants specify idiosyncratic properties of the event described by a verb. For example, *laugh* sets apart the manner of a laughing action from the manner of other actions, and *arrived* sets apart the state of having been arrived from other states. The unique vector representations assigned to words by vector space models of word meaning under the assumption of the distributional hypothesis capture exactly this kind of word-specific semantic information, given that what sets apart the manner of a laughing action from other actions are the specific contexts in which a specific a verb like *laugh* is used to describe the manner of a laughing event.

The predicates DO and BECOME are conceptual primitives that are not specific to a single verb, but recur in the decomposition of whole classes of verbs. DO is a meaning component that is present in the lexical decompo-

sition of all verbs that describe an action and entail an agentive meaning. BECOME is present in all verbs that describe a change of state and entail a patient-like meaning.

1.2 The open question for the components of word meaning

Against the background of the unaccusative hypothesis about intransitive verbs, the general problem with which the present paper is concerned is that work on the interpretation of distributional semantic representations like Levy & Goldberg (2014) fosters relatively clear-cut intuitions about how distributional semantic representations account for the meaning of specific words (and in particular nouns). But there are no similarly clear-cut intuitions about how distributional semantic representations of verb meaning may be understood to characterize concepts such as agency, volition or control, on the one hand and patiency or change of state, on the other. It thus remains an open question whether and how vector space models are also able to represent general concepts that are relevant to the lexical semantics of more than one word. A similar conclusion holds for theoretical approaches to verb meaning, since the type, number and determination of the recurring conceptual features relevant to verb meaning is a central, yet unsolved research problem. While theoretical work on lexical semantics assigns meaning components like agency or change of state a central role in the definition of thematic roles, it appears that "[t]here is perhaps no concept in modern syntactic and semantic theory which is so often involved in so wide a range of contexts, but on which there is so little agreement as to its nature and definition, as THEMATIC ROLE" (Dowty 1991: 547).

1.3 Methodological remarks

A combination of theoretical and computational approaches to lexical semantics poses at first a methodological challenge. Distributional semantics considers the main problem of lexical semantics to be a problem of empirical breadth: "[t]he problem of lexical semantics is primarily a problem of size: even considering the many subregularities found in the content lexicon, a hand-by-hand analysis is simply not feasible for the thousands of elements that populate the content word lexicon" (Baroni et al. 2014: 246). Accordingly, distributional semantic models are assessed ac-

cording to their extrinsic value, which can be measured by the performance of a certain model in a downstream task that evaluates a model's predictions against a large set of data annotated according to a gold standard. But the extrinsic assessment of computational approaches to semantics leaves open the question for whether and how computational models not only accomplish a certain task but also have an intrinsic explanatory value. As Lenci (2014) argues with a case study on the distributional classification of Italian verbs, the reproduction of a theoretically defined gold standard of verb classification does not indicate what the concepts or semantic features are like that underlie the classification. In contrast, theoretical approaches consider the main problem of lexical semantics to be a problem of analytical depth (Levin & Pinker 1991: 1):

> If *to paint* means 'cause to be covered with paint', why isn't it *painting* when a paint factory explodes or when Michelangelo dips his brush into the can [...]? These particular definitions can be patched up, but sceptics foresee a never-ending need for such patching, with no real increase in watertightness.

As a way out of this problem, theoretical lexical semantics proposes to focus on recurrent components in the decomposition of word meaning that can be used to define the meaning of verbs in terms of verb classes. But since there is no principled way to pin down the meaning components of a verb, the components of lexical decomposition have to be stipulated rather than being derived them from empirical observations, as Van der Leek (1996) argues with a case study on the conative alternation.

Given that theoretical and computational approaches to lexical semantics have clearly defined and widely agreed methodological standards, a middle ground between theoretical and computational semantics is likely to fall short of the established expectations of both theoretical and computational approaches to lexical semantics. The different focus of theoretical and distributional approaches to word meaning is easily mistaken as excluding an interoperability of the two methods, since bold and simple, from the viewpoint of theoretical lexical semantics, distributional approaches to word meaning fail to account for systematic, non-idiosyncratic aspects of word meaning, whereas from the viewpoint of distributional approaches to word meaning, theoretical lexical semantics fails to account

for idiosyncratic aspects of word meaning. The goal of the present paper is to widen the view on both theoretical and computational lexical semantics by bringing together the benefits of both approaches in the examination of the question whether, and if yes, how, general concepts like agency and patiency (represented by primitive predicates like DO and BECOME), respectively, are reflected in the embeddings of intransitive verbs.

1.4 Outline of the paper

I approach the problem of identifying recurrent meaning components in distributional semantic representations with the following steps.

First, the learning objective for word embeddings is distributional similarity, not the kind of decompositional similarity underlying the unaccusative hypothesis. I deal with this problem in §2 by fine-tuning general-purpose word embeddings of German and English intransitive verbs with the objective of learning to distinguish between unaccusative and unergative verbs. To identify the specific "surplus" that fine-tuning makes to the embeddings, I compute the difference vectors that represent the contribution of fine-tuning by subtracting task-specific intransitive verb embeddings from the corresponding general-purpose embeddings.

Second, since the difference vectors computed in §2, like general word embeddings, are dense and continuous, they are uninterpretable by humans. I address this problem by rendering the word embeddings interpretable through approximation of their meaning with their nearest neighbors in the embedding space (where the nearest neighbors of a given vector are the most proximate word embeddings in the embedding space).

Third, we do not know a priori how general concepts like those represented by DO and BECOME might be represented in the difference vectors. To account for this problem, in §4 I search the nearest neighbors of the difference vectors for meaningful and diagnostic linguistic cues that have been proposed in the theoretical literature to indicate an agentive or patient-like meaning. This qualitative inspection reveals that the nearest neighbors correspond to prototypical linguistic realizations of unergativity, like *-er* nominals and intentional actions for the unergative difference vector and descriptions of unintentional change of state events for the unaccusative difference vector. I conclude that the effect of retraining is neither trivial nor random, but instead captures surprisingly well the proper-

ties that linguists have found to be characteristic of agentive and patient-like verb meanings. In turn, the more general insight that I discuss in §5 is that the concepts that make up the meaning components of verbs are not directly encoded in embeddings by increasingly abstract terms as in ontologies like WordNet (Fellbaum 1998), but through patterns of word formation.

Finally, to further narrow down the concepts acquired through fine-tuning, in §6 I abstract away from specific difference vectors by estimating a linear regression model of the correlation between retrained and baseline unergative and unaccusative embeddings. The regression model allows to transfer the effect of retraining to other linguistic domains where the same distinction between DO and BECOME has been argued to be relevant: agentive *-er* and patient-like *-ion* nominals, respectively. The regression model also captures more complex conceptual generalizations, e.g. when the regression model for atelic DO and telic BECOME is used to transform the embeddings of mass and count nouns (following Bach (1986)). §7 concludes.

2 Data and methods

2.1 German and English word embeddings

The basis for the results reported in the present paper are German and English word embeddings learned with Word2Vec (Mikolov et al. 2013; SGNS: skip-gram with negative sampling, CBOW: continuous bag of words with hierarchical softmax, 300 dimensions, using 10w(ord) and 5w windows). The German embeddings were learned from SdeWac, a 0.88 billion word corpus of parsable German web data (Faaß & Eckart 2013), where the embedding space has a vocabulary of 237615 words. The English embeddings were learned from ukWac, a 1.3 billion word corpus of English web data (Ferraresi et al. 2008), where the embedding space has a vocabulary size of 70950 words. In the following, I refer to the German and English Word2Vec verb embeddings described in this section as "baseline embeddings". BERT embeddings (Peters et al. 2018) for German and English verbs were extracted from the multi-cased L-12/H-768/A-12 model,[2] using mean reduce as a pooling strategy on the last two output layers. In

[2]Downloaded from `https://github.com/google-research/bert`

the present paper I use BERT embeddings only for comparison for the following reasons. First, out-of-the-box BERT embeddings do not perform better in capturing unaccusativity than Word2Vec embeddings (see §2.4). Second, contextualized embeddings like BERT represent word meaning only relative to a context, which makes it difficult to pin down meaning components like DO and BECOME that words have independent of their context. Third, to rule out an interference of overt syntactic properties of unaccusativity like auxiliary selection during training of the embedding model, the Word2Vec embeddings used in this paper were won from corpus data filtered to consist only of content words. In comparison to BERT, Word2Vec is a computationally cheap way to learn embeddings and thus Word2Vec models can be trained from scratch without investing huge amounts of computation time.

2.2 Intransitive Verb Dataset

SdeWac was parsed with the syntactic and semantic dependency parser described in Björkelund et al. (2010). I extracted verbs that the parser saw more than 90 percent in an intransitive construction together with the semantic role label of the single argument (grammatical subject or grammatical object). I manually corrected the semantic role labels, using auxiliary selection in the present perfect (see (1)) as a diagnostic. In a further step of cleaning, I removed two classes of intransitive verbs that have been argued to involve an unaccusativity mismatch (Zaenen 1988) and thus are not unambiguously unergative or unaccusative, so-called verbs of emission and particle verbs of directed movement. In total, I ended up with a vocabulary of 972 unergative and 840 unaccusative German verb embeddings. Since English doesn't have reliable markers of unaccusativity (such as auxiliary selection), determining whether an English intransitive verb is unergative or unaccusative is more involved. I thus relied on existing lexical resources and used a subset of the unambiguously internally and externally caused verbs listed in the appendix of Levin & Rappaport Hovav (1995). As examples of externally caused verbs, I chose the classes of "alternating change of state verbs" and "cooking verbs" (251 verbs, class labels according to Levin 1993). For internally caused verbs, I chose "run verbs", verbs that partake in the unspecified object alternation and verbs that alternate with a cognate object construction (275 verbs). I used the

Embedding	linear SVM
baseline-SGNS-10w-de	fail to converge
retrained-SGNS-10w-de	1.0
baseline-CBOW-10w-en	0.89
retrained-CBOW-10w-en	0.99
BERT multi-cased-de	fail to converge
BERT multi-cased-en	fail to converge

Table 1 F1-score (10-fold cross-validation) of the binary classification of unergative/unaccusative verb embeddings with linear SVM.

classification of Levin & Rappaport Hovav (1995) to label intransitive verbs as subcategorizing either a grammatical subject or a grammatical object. The annotations for English nouns as "regular count nouns" and "regular mass nouns" were taken from the Bochum Countability Lexicon (Kiss et al. 2016).

2.3 Fine-tuning of verb embeddings through retraining

I fine-tuned the baseline embeddings for the German and English intransitive verbs through retraining the embeddings on a binary classification task on the unergative/unaccusative distinction, using RMSprop as the optimization algorithm and binary crossentropy as a loss function. For the retraining, I used a simple neural network architecture consisting of an embedding layer fully connected to a single output neuron. The resulting embedding layer is of size 1812 × 300 for German and 526 × 300 for English. The output neuron has a sigmoid activation function to yield a continuous probability distribution over the binary labels. To make sure the embeddings memorize the distinction between unergative and unaccusative verbs, the embeddings were overfit to 1.0 accuracy on the intransitive verb dataset.

2.4 Effect of fine-tuning: Linear separation

The effect of retraining is a linear separation of unergative and unaccusative verb embeddings, which becomes clearly visible in Figures 1 and 2, where the embeddings are projected down to two dimensions with PCA. In numbers, the linear separability achieved through retraining manifests itself in the improvement of the F1 score of a linear SVM classification task to

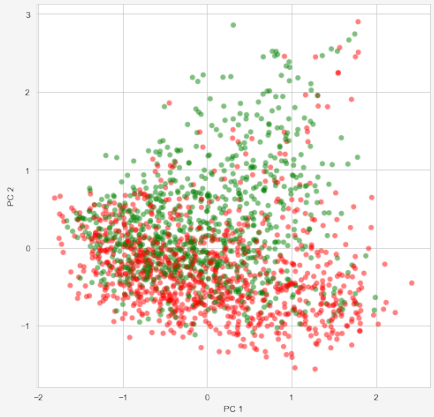

Figure 1 Baseline intransitive verb embedding space projected onto two dimensions with PCA (green = unaccusative verbs, red = unergative verbs)

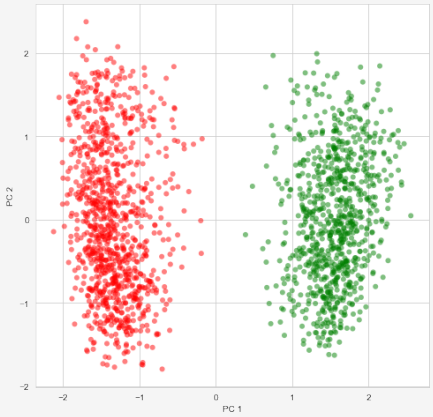

Figure 2 Retrained intransitive verb embedding space projected onto two dimensions with PCA.

near perfect, as reported in Table 1. The already good performance of linear SVM classification for the baseline English verb embeddings may be due to the selection of intransitive verbs, where I took great care to select only verbs that Levin & Rappaport Hovav (1995) classify unambiguously as unergative or unaccusative. It should be noted that state-of-the-art contextual embeddings like out-of-the-box BERT do not perform better than the classical Word2Vec embeddings, and thus that there is a genuine advantage in using contextualized embeddings over Word2Vec embeddings.

3 Quantitative evaluation

With respect to the quantitative effect of retraining, the first thing to rule out is that the weight updates through retraining are trivial. The weight updates would be trivial if the fine-tuned embeddings would have been changed in a way such that unergative and unaccusative embeddings are shifted to clearly distinct regions of the embedding space, but where the retrained embeddings would no longer be semantically similar to the baseline embeddings from which they are derived. To rule out that the weight updates are trivial in this sense, I consider the overlap in nearest neighbors between baseline and retrained embeddings. Here, and in what fol-

lows, I follow Levy & Goldberg (2014) and calculate the nearest neighbors of a word embedding with the dot-product, and correlate the magnitude of the dot-product with semantic similarity.

The overlap in the 10 nearest neighbors of the untrained and retrained verb embeddings is 76% for the German 10w SGNS model, and 97% for the English 10w CBOW model. These differences between German and English are likely to result from the difference in vocabulary size. Given that the German vocabulary is much bigger than the English vocabulary, we expect that for a given word, there is a high number of nearest neighbors with a small distance, and thus that even small changes in the embedding weights leads to a difference in the calculation of nearest neighbors. This explanation is supported by the fact that 96% of the top neighbors of the baseline German verb embeddings are also neighbors of the corresponding retrained verb embedding, whereas 87% of the top nearest neighbors of English baseline embeddings are neighbors of the corresponding retrained verb embeddings. The retrained verb embeddings thus basically retain the position of the baseline embedding in the embedding space, which suggests that retraining optimizes only a small subspace of the full embedding space without affecting the distributional similarities encoded in the embedding space as a whole. Injection of linguistic knowledge into word embeddings through retraining thus constitutes a cheap and effective way to improve the performance of word embeddings in classification problems that are difficult to approach on the basis of the distributional hypothesis alone, a finding the further exploration of which I leave to future research.

4 Qualitative interpretation

Given that the quantitative effect of retraining is non-trivial, the consequent question is whether the linear separation of unergative and unaccusative embeddings achieved through retraining is also systematic and non-random, i.e. whether the weights updates through retraining can be given a qualitative explanation that is related to the lexical decomposition of intransitive verbs. For this qualitative investigation of the retrained embeddings, I isolate the “surplus” that the retraining makes to the embeddings by subtracting for each of the embeddings of our German and English vocabulary the weights of the baseline embedding from the weights

of the retrained embedding as in (3) and then approximate the difference vectors with their nearest neighbors in the embedding space of the baseline embeddings.

(3) difference vector = retrained embedding − baseline embedding

One might wonder, as a reviewer does, about the mathematical soundness of applying operations of linear algebra on embedding spaces, given that embedding spaces like Word2Vec (and the retrained embeddings used in this paper) are learned through the application of non-linear operations like the sigmoid function. In fact, it is one of the surprising properties of word embedding models like Word2Vec that the model learned by the algorithm encodes certain semantic properties that can be revealed through the application of linear algebra to word embeddings. Most famously, word embedding models can be employed to solve word analogy tasks like "man is to woman like king is to ?" with methods of linear algebra. Linear algebra has also been used to model semantic aspects of morphological derivation (Padó et al. 2016) or verb formation (Pross et al. 2017). In a manner of speaking, the method employed in the present paper can thus be understood as an attempt to reveal with methods of linear algebra the properties of embedding spaces that correspond to the unaccusative hypothesis. The difference vectors of unergative and unaccusative embeddings represent the relevant semantic information that is responsible for the linear separability reported in Table 1. The first interesting observation to be made is that the approximation of the difference vectors with their nearest neighbors crystallizes in a small and restricted set of shared and recurrent nearest neighbors. For example, if we calculate the 10 nearest neighbors for each intransitive verb, there are 18120 possible nearest neighbors for the German verbs. But the approximation of the difference vectors in the 10w-SGNS German model results in a total of 10 shared and recurring nearest neighbors for the unergative difference vectors and 16 shared and recurring neighbors for the unaccusative difference vectors. The same observations holds for the English difference vectors. Out of a total of 5260 possible 10 nearest neighbors of the English verbs, in the 10w-CBOW English model the difference vectors for the unergative embeddings consist of 15 recurring neighbors and and 24 neighbors for the unaccusative difference

vectors. In the next section, I consider the nearest neighbor approximation of difference vectors in more detail.

4.1 Unergative difference vectors

Consider first the approximated German and English difference vectors for German 10w-SGNS embeddings and English 10w-CBOW embeddings in (4) and (5), where I list the six most frequent nearest neighbors of the unergative and unaccusative difference vectors. I cleared the neighbor lists of proper names and derivationally related words. The chars following the neighbor represent the POS-Tag (n=noun, v=verb, a=adjective) and the numbers following the POS-Tag indicate how often that neighbor occurred in the approximations of the difference vectors of the set of verbs in question.

(4) *Prüferinnen.n.972 Informatikkaufmann.n.972*
'examiners' 'IT merchant'
Gruppenarbeitsraum.n.972 Diätassistent.n.972
'group office' 'diet assistant'
Kulturmanager.n.952
'culture manager'

(5) beginner.n.251 beginners.n.251 sewing.n.251 salary.n.251
crafty.a.251 ceilidh.n.246.0

An obvious cue of an agentive meaning in the neighbor approximations of the unergative difference vectors are *-er* nominals like *Prüfer* ('examiner'), *beginner* and job descriptions like *Kaufmann* ('merchant') or *Assistentin* ('assistant') , but also descriptions of agentive actions like *sewing* and related adverbials like *crafty*. The general tendency of an agentive meaning that appears in the approximated unergative difference vectors is stable across the statistical variation to retraining introduced by optimization with RMSprop. (6) and (7) are the six overall most frequent neighbors of the difference vectors of 50 runs of retraining the German 10w-SGNS model and the English 10w-CBOW model, respectively. Again, we encounter clear indicators of agentive meaning, job descriptions like *Redakteur* ('editor') or *illustrator*, agentive verbs like *rülpsen* ('burp') or agentive nouns like *sleeps*.

(6) *Transgender.n.1976 Redakteur.n.1944 Chef.n.1944*
'transgender' 'editor' 'chief'
Familiensonntag.n.1944 rülpsen.v.1942 Akustiker.n.1650
'family sunday' 'burp' acoustician'

(7) prance.v.3491 toy.n.1408 groovy.n.1384 illustrator.n.1296 sleeps.n.1170 replies.n.1006

Indication of an agentive meaning is also stable across window size. (8) and (9) are the approximations of the unergative difference vectors using five word window SGNS embeddings for German and 5w CBOW embeddings for English. Indicators of agentive meaning are again *-er* nominals like *Comiczeichner* ('cartoonist') or *tinker* and job descriptions like *illustrator*, but also intentional verbs like *freuen* ('be glad') and mental adjectives like *stinksauer* ('spitting mad').

(8) *Comiczeichner.n.972 freuen.v.972 klingeltöne.n.972*
'cartoonist' 'be glad' 'ring tones'
Mitgliedermagazin.n.972 rülpsen.v.971 stinksauer.a.677
'member magazine' 'burp' 'spitting mad'

(9) warranties.n.251 tinker.n.251 cordless.n.251 illustrator.n.251 difficulties.n.250 enquiries.n.250

4.2 Unaccusative difference vectors

In contrast to the indicators of agentive meaning that are present in the approximations of the unergative difference vectors, the unaccusative difference vectors for the German 10w-CBOW and English 10w-SGNS model in (10) and (11) are characterized by verbs that describe non-intentional change-of-state events like *ionisieren* ('ionize'), *purify* or *refract*, *-ion* nominals like *separation*, *eruption*, inanimate yet causally active substances like *toxin* and adjectival nominalizations such as *redness*.

(10) *Kontinentalplatte.n.840 ionisieren.v.840.0 Bremsvorgang.n.840*
'continental plate' 'ionize' 'braking process'
Vorderreifen.n.840 aushärten.v.839 Ladungstrennung.n.825
'front tire' 'harden' 'charge separation'

(11) purify.v.274 refract.v.265 redness.n.265 adverse.a.135 toxin.n.37

eruption.n.20

As for the unergative difference vectors, indicators of patient-like meanings in the unaccusative difference vectors are stable across window sizes. (12) and (13) are the nearest neighbors of the unaccusative difference vectors for five word window SGNS (German) and CBOW (English) embeddings. Again, we encounter *-ion* nominals like *reaction*, *circulation* and nominals that describe non-intentional change-of-state events like *heat* as well as inaminate yet causally active substances (*oxygen*, *metastasis*) and their properties (*flammable*).

(12) *Kernreaktion.n.840 Tochtergeschwülste.n.840*
'nuclear reaction' 'metastasis'
Unfallereignis.n.840 Verwitterung.n.840 Katarakt.n.840
'event of accident' 'weathering' 'cataract'
Ausschwemmung.n.836
'washout'

(13) heat.n.275 circulation.n.275 oxygen.n.275 flammable.a.275
starch.n.274 inefficient.a 268

The general tendency of a patient-like meaning that appears in the approximated unaccusative difference vectors is also stable across the statistical variation to retraining introduced by RMSprop. (14) and (15) are the six overall most frequent neighbors of the difference vectors of 50 runs of retraining the German 10w-SGNS embeddings and the English 10w-CBOW embeddings, respectively. Again, we see clear indicators of a non-intentional change-of-state meaning, like *Dickenwachstum* ('growth in girth') or *warming*.

(14) *einspritzen.v.2523 Dickenwachstum.n.2431 Zugspannung.n.1682*
'inject' 'growth in girth' 'tension'
Gerinnsel.n.1682 Verbuschung.n.1682 Schwerefeld.n.1682
'clot' 'scrub encroachment' 'gravitation field'

(15) moisture.n.144 temperature.n.114
warming.n.55 unenforceable.a.52 cremation.n.50 detoxify.v.40

In sum, an informed linguist is able to interpret the nearest neighbors of

the unergative and unaccusative vectors as characterizations of the distinction between agent- and patient-like meanings by detecting word formation patterns connected to unergativity, such as *-er* nominals, and shared lexical entailments of proto-agent and proto-patient properties in the sense of Dowty (1991). Since these observations can be obtained independently for both English and German intransitive verbs, this suggests that the representations of the semantic correlates of the unaccusative hypothesis by the approximated difference vectors of intransitive verbs are not random outliers but rather point towards a systematic effect of retraining of intransitive word embeddings. One explanation for this systematic effect may be that retraining of the embeddings is a method for strengthening those latent dimensions of the embedding space that involve the same recurrent meanings relative to the retraining objective. In the next §, I consider these latent dimensions of the embedding space in more detail.

5 Word formation and concept representation

In §4, I demonstrated that the quantitative effect of retraining qualitatively amounts to the addition of non-trivial and non-random information to the baseline embeddings. Inspection of the approximated difference vectors of unergative and unaccusative embeddings furthermore suggested that this information does not correspond to the meaning of single words but instead is related to general concepts like agency and intentionality for the unergative difference vectors and non-intentional change of state for the unaccusative difference vectors. An eye-catching observation in this respect is that these concepts are represented in the nearest neighbors of the difference vectors through systematic patterns of word formation. One particularly telling example are nominals derived with *-er*/*-or* for the unergative difference vectors and *-ion*/*-ung* nominals for the unaccusative difference vectors. The observation that word formation, and in particular the semantic distinction expressed by English nominals derived through suffixation with *-er*/*-ion* is highly diagnostic is reinforced by the overall proportion of these nominals in the nearest neighbor approximation of the difference vectors. For example, in the difference vectors of 50 runs of retraining the English 10w CBOW intransitive verb embeddings, there are 443 *-er* nominals in the unergative difference vectors, but only two *-ion* nominals; and conversely, there are 458 *-ion* nominals in

the unaccusative difference vectors, but only 17 *-er* nominals. One way to explain this observation is that the appearance of systematic patterns of word formation is no coincidence, since derivational morphology is one way in which languages overtly mark word-overarching meaning components. In fact, the word formation patterns of *-er*/*-ion* nominals have been argued in the theoretical literature to correlate with the distinction between an agentive and a patient-like meaning. As seen in (16), *-er* nominalizations have an agentive semantics and can only be derived from DO-verbs (Levin & Rappaport 1988).

(16) a. dancer, worker, dreamer
b. *faller, *asleeper, *arriver

In contrast, *-ion* nominalizations as in (17) have a passive semantics and can only be derived from (transitive) BECOME-verbs (Alexiadou 2001):

(17) a. destruction, explosion

These considerations suggest that the concepts that make up recurrent meaning components are not directly encoded in embeddings by increasingly abstract terms as in ontologies like WordNet (Fellbaum 1998), but through patterns of word formation. If this hypothesis is on the right track, we expect to see the same effect in other domains where patterns of word formation are semantically indicative. This expectation is borne out for the case of mass/count nouns. (18) and (19) present the 10 most frequent recurring neighbors of the difference vectors for English word embeddings of 1270 regular count and 2233 mass nouns. The neighbors in (18) and (19) were calculated using the same pipeline as for the unergative and unaccusative verbs, consisting of retraining noun embeddings with a binary classification task (on the mass/count distinction), computation of difference vectors and approximation with their nearest neighbors. As for the unergative/unaccusative verbs, the nearest neighbors of mass/count difference vectors consist of a small and restricted set of recurrent elements. The approximation of the 10w-CBOW difference vectors of English mass nouns consists of 19 recurrent nearest neighbors out of 22330 total possible nearest neighbors, and the approximation of the 10w-CBOW difference vectors of English count nouns consists of 35 recurrent neighbors out

of a total of 12700 possible neighbors.

(18) *10 most frequent neighbors of the difference vectors of English mass nouns*
tannin.n.1270 absorption.n.1269 sensuality.n.1269 warmth.n.1260 tightness.n.1256 metabolism.n.1235 cheques.n.1025 acidity.n.900 solubility.n.836 irritability.n.726

(19) *10 most frequent neighbors of the difference vectors of English count nouns*
keypad.n.2025 aisle.n.2025 doorway.n.2025 door.n.2024 verifiers.n.2019 policeman.n.2013 impale.v.2004 lock.v.1931 baccalaureate.n.746 row.n.520

Notably, in (18) we again encounter word-formation patterns as being indicative of the concepts of a mass noun, here in the form of *-ity* and *-ness* nominalizations. These nominalization formation patterns have been argued in the theoretical literature to derive abstract properties, "qualities" which are ontologically similar to mass noun (Francez & Koontz-Garboden 2017). The concept of a count noun, in contrast, seems to be encoded by the simple fact that the nearest neighbors in (19) – in contrast to the neighbors in (18) – can be counted. I take these observations to provide additional evidence for the claim that retraining adds non-trivial and non-random conceptual information to general-purpose word embeddings.

6 Lexical decomposition with distributional semantics

In §5, I showed through qualitative inspection of difference vectors that the surplus that retraining makes to word embeddings of verbs and nouns is neither trivial nor random, but instead captures the relevant high-level conceptual distinction on which the embeddings were fine-tuned. The observations I made, however, pertained to difference vectors of specific verbs and nouns. To represent the effect of retraining in a more general way independent of specific words, in this section I approximate the relation between the baseline and retrained intransitive English verb embeddings with a linear regression model that captures the effect of retraining as a linear transformation on embeddings. Casting the effect of retraining in such a general and abstract way allows me to investigate in more detail

the kind of concept acquired through retraining by applying the regression model to word embeddings that have been argued in the theoretical literature to involve the same kind of meaning components as intransitive verbs.

6.1 The regression model

To capture the effect of retraining, I estimated two regression models for the relation between (i) the unergative baseline and retrained embeddings and (ii) the unaccusative baseline and retrained embeddings. In (20), B is a coefficient matrix (more precisely, a diagonal matrix with the regression weights as diagonal elements) and Err an error term. Verb embeddings $x \in X_{baseline}$ are the independent variables of the model, and verb embeddings $y \in Y_{retrained}$ the dependent variables. I estimated the coefficient matrices and error terms of the linear regression equation with partial least squares regression on matrices X and Y, using the number of embedding dimensions as latent variables. Since I want to capture the effect of retraining, I overfit the estimated linear models for unergative and unaccusative verbs to an R^2 score of 1.0.

(20) $$Y_{retrained} = X_{baseline} * B + Err$$

I propose to use the estimated regression model for the relation between unergative and unaccusative baselines and retrained verb embeddings as an interpretation of the lexical decomposition of unergative and unaccusative intransitive verbs. (21) and (22) illustrate lexical decomposition with distributional semantics with the examples of unergative *laugh* and unaccusative *arrive*.

(21) $$\text{to laugh} \rightsquigarrow \text{DO } laugh \rightsquigarrow B_{do} * laugh + Err_{do}$$

(22) $$\text{to arrive} \rightsquigarrow \text{BECOME } arrived \rightsquigarrow B_{become} * arrive + Err_{become}$$

In (21) and (22), the conceptual primitives DO and BECOME are interpreted as the coefficients of the estimated linear regression model. As discussed in §1, word embeddings excel in capturing idiosyncratic aspects of word meaning. Thus, I propose that the interpretation of the idiosyncratic constants *laugh* and *arrived* is provided by the corresponding verb embeddings of *laugh* and *arrive*.

6.2 Evaluation

I evaluate the linear regression model by applying the linear transformation described by the model to linguistic domains that have been argued in the theoretical literature to involve the same kind of conceptual distinction. In the first evaluation, I apply the regression model for DO and BECOME to *-er*/*-ion* nominals, and in the second evaluation to mass and count nouns.

To evaluate the regression model against *-er*/*-ion* nominals, I collected all the 3812 "agentive" *-er* and 2357 "patient-like" *-ion* nominalization types contained in ukWac. Linear SVM classification of the English 10w-CBOW baseline embeddings for these nominals fails to converge. I then applied the linear regression model to the *-er*/*-ion* nominal baseline embeddings, as illustrated for the nominals *beginner* and *eruption* in (23a) and (23b), respectively.

(23) a. beginner $\rightsquigarrow B_{do} * \textit{beginner} + Err_{do}$
b. eruption $\rightsquigarrow B_{become} * \textit{eruption} + Err_{become}$

When the regression model for DO is used to transform the embeddings of *-er* nominals and the regression model of BECOME to transform the embeddings of *-ion* nominals as in (23), this improves the F1 score of the transformed English noun embeddings in a linear SVM classification task on the *-er*/*-ion* distinction to 0.97.

In the second evaluation, I test my regression model with a more far-fetched conceptual correlation than the distinction between agentive and patient-like nominals and agentive and patient-like verbs. Bach (1986) proposed a correlation between the concepts underlying the distinction between mass and count nouns, on the one hand, and DO and BECOME verbs on the other: mass is to atelic verbs (i.e. unergative verbs involving DO, see Zaenen (1988)) what count is to telic verbs (i.e. unaccusative and transitive verbs involving BECOME) as in (24).

(24) *atelic* :: *mass* = *telic* :: *count*

Linear SVM classification on the 10w-CBOW English baseline embeddings of the 3000 count and mass nouns from §5 fails to converge. But application of the DO regression model to the mass noun embeddings and the

BECOME regression model to the count nouns yields an F1 score of 0.97 for linear SVM classification of the transformed embeddings. To ensure this is not a random effect, I did a countercheck with interchanged regression models. When the DO-model is applied to count nouns and the BECOME-model to mass nouns, linear SVM classification of the transformed embeddings still fails to converge.

7 Conclusion and outlook

The main goal of the present paper was to bring together linguistic intuitions about word meaning with the empirical groundedness of computational approaches to lexical semantics derived from large corpora.

Before I summarize the findings of this paper and put into them into the context of larger research goals, I briefly return to the methodological issues of approaching questions of lexical semantics with methods of machine learning I already addressed in §1.3. A reviewer notes that the interpretation of word embeddings through approximation with their nearest neighbors, and in particular the interpretation of nearest neighbors as providing support of a certain hypothesis, makes the analysis "merely impressionistic", where one just uses one's own intuition to highlight neighbors that confirm the hypothesis, whereas neighbors that have nothing to do with unaccusativity are dismissed. I believe this is an important point because it pertains to the very question of linguistic theorizing. Broadly speaking, the word embeddings discussed in this paper can be understood as compressed representations of word usage derived from an empirical basis the size of which is beyond the limits of human capacity. Traditional linguistic theory is standardly based on a small set of carefully crafted examples. But exceptions to the rule of course also occur when theorizing about carefully selected examples, and the unaccusative hypothesis is one such topic in modern linguistics where there are countless exceptions to the rule, to the extent that a unified theory of intransitive verbs seems to become impossible (see Alexiadou et al. 2004 for an overview of the current state of the debate). The "impressionistic" method of interpreting approximated word embeddings through the eyes of an informed linguist may thus be considered as reproducing exactly the difficulties of theorizing about linguistic observations, in that they provide an unbiased, yet empirically grounded view of the complexities of linguistic

reality. Since it is difficult to determine which method of picturing natural language is the right one, instead of contrasting the methodological standards of the two approaches, I propose that striving for a combination of insights gained with methods of traditional linguistic theory and methods of machine learning may provide a more accurate picture of the inner workings of natural language,even if the picture may seem fuzzy and ad-hoc. This being said, the purpose of this paper is not to prove a certain point but rather to make an attempt at investigating how word embeddings of a certain data set relate to linguistic theorizing about the same data set.

From a theoretical point of view, the results obtained in this paper suggest that vector models of word meaning may in fact capture the high-level conceptual distinction related to the conceptual primitives DO and BECOME. The possibility to transfer the information gained through retraining of intransitive-verb embeddings to other domains shows that the same meaning component may in fact be present in intransitive verbs, derived *-er*/*-ion* nominals and mass and count nouns. Since word embeddings are learned from data independent of theoretical bias and intuitions, the case study presented in this paper suggests that word embeddings can be understood as providing a fresh perspective on the conceptual foundations of lexical semantics. The consequent challenge for the theoretical linguist is to make use of this new perspective on lexical meaning, since the observations made in the present paper to spell out in more detail the ontology of the parallels between the verbal and nominal domain that in the tradition of Bach have mainly been assessed in terms of a structural algebra. I leave the question for a unified ontology of verbs and nouns suggested by the case study in this paper as a challenge to future research.

From a computational point of view, it is important to note that I did not feed the relation between DO and BECOME and derived *-er*/*-ion* nominals or mass and count nouns into the model. The regression model thus captures general aspects of the linguistic intuition underlying the conceptual distinction between DO and BECOME. A possible explanation for this observation is that retraining strengthens latent semantic features of the baseline word embeddings. In the approach to semantic similarity based on the dot-product, feature strengthening finds expression in the magni-

tude of a vector. In fact, the regression weight matrices of the regression models for DO and BECOME are not just diagonal matrices. They are near-perfect scalar matrices, where the mean of the DO diagonal vector is 0.84 (with a variance of 0.00091) and the mean of the BECOME diagonal vector is 0.936 (with a variance of 0.00037). One main effect of the application of the regression model is thus a stretching of embeddings by a certain scalar factor. From this point of view, the results presented in §6 suggest that the scalar factors we estimated with the regression model are indeed semantically meaningful. In turn, one may speculate whether general concepts are represented in vector space models of meaning like those that can be derived with the Word2Vec algorithm as scalar operations on embeddings. I leave a further exploration of this idea to future research.

The more general hypothesis that arises from the discussion in the present paper concerns the nature of embedding spaces by themselves. Embeddings of specific words only occupy a small part of the full (continuous and dense) embedding space. This paper showed that some of the vectors and operations in the embedding space that do not correspond to words are actually meaningful: they represent abstract concepts like those associated with DO and BECOME. The more general goal of future research that this paper envisages is to use linguistic insights as a roadmap of the terra incognita of embedding spaces in which meaningful vectors live that do not correspond to specific words. Ever since Fillmore (1968), who characterized the covert concepts encoded by verbs as "a set of universal, presumably innate, concepts which identify certain types of judgments human beings are capable of making about the events that are going about around them, judgments about such matters as who did it, who it happened to, and what got changed" (Fillmore 1968: 45), linguists have been working on developing frameworks in which concepts that refer to meanings that are not associated with words. Linguistic theory thus plays a central role in the identification of meaningful non-word vectors in embedding spaces by providing frameworks in which non-word concepts can be systematically investigated. In turn, bringing together the compact and efficient representations that have been developed in theoretical lexical semantics with the empirical grounding of computational approaches to word meaning in the spirit of the distributional hypothesis may provide a way out of the ever-growing resource demands of modern natural lan-

guage processing.

Acknowledgments I would like to thank the reviewers and participants of CSSP 2019 for valuable feedback on earlier versions of this paper. The research reported in this paper has been supported by DFG grant PR1860/1.

References

Alexiadou, Artemis. 2001. *Functional structure in nominals – nominalization and ergativity*. Amsterdam: John Benjamins.

Alexiadou, Artemis, Elena Anagnostopoulou & Martin Everaert (eds.). 2004. *The unaccusativity puzzle: Explorations of the syntax-lexicon interface*. Cambridge University Press.

Asher, Nicholas, Tim Van de Cruys, Antoine Bride & Márta Abrusán. 2016. Integrating type theory and distributional semantics: A case study on adjective–noun compositions. *Computational Linguistics* 42(4). 703–725.

Bach, Emmon. 1986. The algebra of events. *Linguistics and Philosophy* 9. 5–16.

Baroni, Marco, Raffaela Bernardi & Roberto Zamparelli. 2014. Frege in space: A program of compositional distributional semantics. *LiLT* 9. 241–346.

Björkelund, Anders, Bernd Bohnet, Love Hafdell & Pierre Nugues. 2010. A high-performance syntactic and semantic dependency parser. In *Coling 2010: Demonstrations*. 33–36. Coling 2010 Organizing Committee. `http://aclweb.org/anthology/C10-3009`.

Boleda, Gemma. 2020. Distributional semantics and linguistic theory. *Annual Review of Linguistics* 6. 213–234.

Clark, Stephen. 2015. Vector space models of lexical meaning. In Shalom Lappin & Chris Fox (eds.), *The handbook of contemporary semantic theory*, 493–522. Wiley Blackwell 2nd edn.

Dowty, David R. 1991. Thematic proto-roles and argument selection. *Language* 67(3). 547–619.

Faaß, Gertrud & Kerstin Eckart. 2013. SdeWaC – a corpus of parsable sentences from the web. In Iryna Gurevych, Chris Biemann & Torsten Zesch (eds.), *Proceedings of the International Conference of the German Society for Computational Linguistics and Language Technology*.

Fellbaum, Christiane. 1998. *Wordnet: An electronic lexical database*. Cambridge, MA: Bradford Books.

Ferraresi, Adriano, Eros Zanchetta, Marco Baroni & Silvia Bernardini. 2008. Introducing and evaluating ukWaC, a very large web-derived corpus of English. In *Proceedings of the 4th Web as Corpus Workshop (WAC-4)*. 47–54.

Fillmore, Charles J. 1968. The case for case. In Emmon Bach & Robert T. Harms (eds.), *Universals in linguistic theory*, 1–90. Holt, Rinehart, and Winston.

Firth, John R. 1957. *Papers in linguistics, 1934–1951*. Oxford: Oxford University Press.

Francez, Itamar & Andrew Koontz-Garboden. 2017. *Semantics and morphosyntactic variation*. Oxford: Oxford University Press.

Grewendorf, Günther. 1989. *Ergativity in German*. Dordrecht: Foris.

Harris, Zellig. 1954. Distributional structure. In *Papers in structural and transformational linguistics*, vol. 10, 1456–1462. Dordrecht: Springer. doi:10.1007/978-94-017-6059-1_36.

Kiss, Tibor, Francis Jeffry Pelletier, Halima Husic, Roman Nino Simunic & Johanna Marie Poppek. 2016. A sense-based lexicon of count and mass expressions: The Bochum English countability lexicon. In *Proceedings of LREC*. 2810–2814. Portorož, Slovenia: European Language Resources Association (ELRA).

Lenci, Alessandro. 2014. Carving verb classes from corpora. In Raffaele Simone & Francesca Masini (eds.), *Word classes: Nature, typology and representations*, 17–36. John Benjamins.

Levin, Beth. 1993. *English verb classes and alternations: A preliminary investigation*. Chicago: University of Chicago Press.

Levin, Beth & Steven Pinker. 1991. Introduction. In *Lexical & conceptual semantics*, 1–8. Blackwell.

Levin, Beth & Malka Rappaport. 1988. Non-event *-er* nominals: A probe into argument structure. *Linguistics* 26. 1067–1083.

Levin, Beth & Malka Rappaport Hovav. 1995. *Unaccusativity at the syntax-semantics interface*. Cambridge, MA: The MIT Press.

Levy, Omer & Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence & K. Q. Weinberger (eds.), *Advances in neural information processing systems 27*, 2177–2185. Red Hook, NY: Curran Associates, Inc.

McNally, Louise & Gemma Boleda. 2017. Conceptual vs. referential affordance in concept composition. In James A. Hampton & Yoad Winter (eds.), *Compositionality and concepts in linguistics and psychology*, 245–267. Cham: Springer International Publishing.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S. Corrado & Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of NIPS*. 3111–3119.

Miller, George A. & Walter G. Charles. 1991. Contextual correlates of semantic

similarity. *Language and Cognitive Processes* 6. 1–28.

Padó, Sebastian, Aurélie Herbelot, Max Kisselew & Jan Šnajder. 2016. Predictability of distributional semantics in derivational word formation. In *Proceedings of coling*. 1285–1296. Osaka, Japan. `http://www.aclweb.org/anthology/C/C16/C16-1122.pdf`.

Pennington, Jeffrey, Richard Socher & Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of EMNLP*. 1532–1543.

Perlmutter, David M. 1978. Impersonal passives and the unaccusative hypothesis. In *Proceedings of the 4th Annual Meeting of the Berkeley Linguistics Society*. 157–190. Berkeley Linguistics Society.

Peters, Matthew E., Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee & Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of NAACL*. 2227–2237.

Pross, Tillmann, Antje Roßdeutscher, Sebastian Padó, Gabriella Lapesa & Max Kisselew. 2017. Integrating lexical-conceptual and distributional semantics: A case report. In Alexandre Cremers, Thom van Gessel & Floris Roelofsen (eds.), *Proceedings of the 21st Amsterdam Colloquium*. 75–85.

Sorace, Antonella. 2000. Gradients in auxiliary selection with intransitive verbs. *Language* 76. 859–890.

Turney, Peter D. & Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research* 37. 141–188.

Van der Leek, Frederike. 1996. The English conative construction: A compositional account. In *Proceedings of the Chicago Linguistic Society 32*, vol. 32. 363–378. Chicago Linguistic Society.

Wunderlich, Dieter. 1985. Über die Argumente des Verbs. *Linguistische Berichte* 97. 183–227.

Zaenen, Annie. 1988. Unaccusatives in Dutch and the syntax-semantics interface. CSLI Report 123 Center for the Study of Language and Information Stanford, CA.

Attributive *wrong* in underspecified semantics

Manfred Sailer

Abstract Attributive *wrong* as in *Alex opened the wrong bottle* shows a non-local reading, that is, its meaning is not local to the noun phrase but interacts with the meaning of the rest of the sentence. I argue that previous accounts did not assume the correct semantics for attributive *wrong* and do not account adequately for its restriction to the definite article. I show that there is a second non-local reading and use data from Papiamentu to show that *wrong* noun phrases are semantic uniques. I present an analysis within a framework of underspecified semantics that (i) treats *wrong* as an ordinary adjective in how it combines with the head noun, (ii) captures its non-local readings, (iii) accounts for the definiteness restriction, and (iv) can address the parallels and differences between non-local and local readings.

Keywords definiteness restriction · Lexical Resource Semantics · non-local adjective · Papiamentu · underspecified semantics

M. Sailer, Goethe-University Frankfurt a.M., sailer@em.uni-frankfurt.de

In Christopher Pinon & Laurent Roussarie (eds.), Empirical Issues in Syntax and Semantics 13, 109–138. Paris: CSSP. http://www.cssp.cnrs.fr/eiss13/

1 Introduction

Haïk (1985), Larson (2000), and Schwarz (2006) argue that attributive *wrong* systematically shows non-local readings. Schwarz' running example is given in (1), together with his paraphrase for the intended reading.

(1) I opened the wrong bottle of wine.
'I opened a bottle that it was wrong for me to open.'

Schwarz (2020) revises this to the following paraphrase.

(2) Liz underlined the wrong number.
'The number that Liz underlined is not the number she was supposed to underline.'

I will provide and justify a slightly different meaning of non-local *wrong*

in §2. In particular, I will argue for the existence of a so far unnoticed reading. I will also include aspects of the meaning of non-local *wrong* that have not been explicitly discussed in the literature: the discourse anaphoric potential noun phrases with *wrong*, the type of definiteness found with such noun phrases – for which point I will use data from a language with a different system of definiteness marking – and, finally, the type of modality attested associated with non-local *wrong*. I will, then, decompose attributive *wrong* into various parts: a set-inclusion statement, negation, a modal operator, and an iota-operator. The contributed operators show variable relative scope, which captures the two readings that I will argue for.

I will show that the readings I postulate are non-local in the sense of Schwarz (2020) in §3. I will discuss some challenges of previous approaches in §4. In §5, I will formulate my analysis in a framework of underspecified semantic combinatorics, *Lexical Resource Semantics* (LRS). This will allow me to treat *wrong* just like an intersective adjective from the point of view of the syntax-semantics interface. Following the literature, I will assume that non-local *wrong* is only found in definite noun phrases. I will, however, look at attribute *wrong* in indefinite noun phrases in §6, where I will try to connect non-local and local attributive *wrong*.

2 The meaning of non-local attributive *wrong*

In this section, I will propose a new semantic analysis of non-local attributive *wrong*. In particular, I claim that there are two readings, which I will call the *police reading* (P-reading) and the *Bluebeard reading* (B-reading), based on the subjects of the prototypical examples in (3) and (4).

(3) The police arrested the wrong person. (P)
'The person that the police arrested is not (among) the person(s) that the police should have arrested.'

(4) Bluebeard's wife opened the wrong door. (B)
'The door that Bluebeard's wife should not open is among the doors that Bluebeard's wife opened.'

In the P-reading in (3), there is a particular person that got arrested. There is also a (often singleton) set of persons that should have been arrested. The arrested person is not in this set.

The B-reading refers back to the French folktale *Bluebeard* (*Barbe bleue* in the original): Bluebeard allows his wife to go into all rooms of their palace except for one. She opens the door to exactly this room and finds the corpses of Bluebeard's former wives in it. In other words, in the B-reading in (4), there is a particular door that should not have been opened. However, this door is in the set of doors that got opened.

Previous discussions of non-local *wrong* only considered the P-reading. However, the example in (4) shows that the B-reading exists as well. The two readings can also be found with example (2), even though Schwarz (2020) only discusses the P-reading.

(5) Liz underlined the wrong number.
   a. 'The number that Liz underlined is not among the numbers Liz was supposed to underline.' (P)
   b. 'The number that Liz was not supposed to underline was among the numbers Liz underlined.' (B)

In this section, I will justify the existence of the P- and the B-readings by going through various aspects of their meaning: uniqueness, discourse-anaphoric potential, definiteness, and type of modal operator. Finally, I will present my formal rendering of the above paraphrases.

2.1 Uniqueness

For a sentence like (2), Schwarz (2020) looks at three sets: (i) the *actual set*, referring to the numbers that Liz underlines, (ii) the *required set*, which are the numbers that need to be underlined, and (iii) the *excluded set*, i.e., the numbers that must not be underlined. According to Schwarz, sentence (2) comes with a uniqueness and existence presupposition for the actual set and the required set, but not for the excluded set.

I will briefly summarize the arguments for the existence and uniqueness of the actual set for the P-reading and add the corresponding examples for the B-reading. Schwarz uses the convenient notation in (6), where the required set is underlined on the left side of the pipe, "|", and the actual set is underlined on the right side. I use "✓" when a sentence is true in the given scenario and "✗" when it is false. The symbol "#" marks an uninterpretable sentence, pointing to a violation of a presupposition.

(6) Liz underlined the wrong number. (P)

a. ✓4 7 <u>9</u> | 4 <u>7</u> 9
b. ✗4 7 <u>9</u> | 4 7 <u>9</u>
c. # 4 7 <u>9</u> | 4 7 9
d. # 4 7 <u>9</u> | 4 <u>7</u> <u>9</u>

The scenario in (6a) has a single underlined number, which is not the required number. As indicated, the sentence is true in (6a). The existence and uniqueness of the actual set is also given in (6b), where Liz underlines the number she is supposed to underline, which makes the sentence false. In (6c), the actual set is empty. According to Schwarz, the sentence is not interpretable in this scenario, which suggests a failure of an existence presupposition. In (6d), more than one number is underlined. Again, Schwarz argues that the sentence cannot be interpreted in such a constellation, which points to a violation of a uniqueness presupposition.

According to Schwarz (2020), the required set comes with an existence and a uniqueness presupposition as well. The scenario in (7a) has an empty required set. Indeed, the sentence is odd in this scenario, which points to a violation of an existence presupposition. In (7b), Liz is supposed to underline two numbers but underlines one number that is not required. According to my intuition, the sentence is true in (7b).

(7) Liz underlined the wrong number. (P)

a. # 3 4 7 9 | 3 <u>4</u> 7 9
b. ✓3 4 <u>7</u> <u>9</u> | 3 <u>4</u> 7 9

This intuition carries over to more natural cases such as (8).[1] I characterize the relevant scenario in square brackets. The police is supposed to arrest more than one person (i.e., there is a non-singleton required set), but a single person is arrested that is not in this set. The sentence is perceived as true in such a scenario.

(8) [We know that three criminals robbed a bank and should be arrested, but the police arrested a single, innocent bystander instead.] The police arrested the wrong person. (P)

[1] I am grateful to the CSSP 2019 audience for providing the example scenario in (8).

My judgements of (7b) and (8) point to an asymmetry between the presuppositions of the actual and the required set. Schwarz (2020) calls the combination of an existence and uniqueness presupposition a *definiteness presupposition*. Given my judgments, only the actual set has a definiteness presupposition under the P-reading.

We can now apply the same method to the B-reading. For this reading, the excluded set is relevant rather than the required set. It is important to note that there is no required set in the B-reading. In other words, Bluebeard's wife is not forced to open any of the doors at all in (4), she is just not allowed to open a particular one. For B-reading scenarios, I will use the same marking as above, but indicate what is allowed rather than what is obligatory. To avoid confusion, I will use "||" to separate the *permitted* set from the actual set. The scenarios relevant for testing existence and uniqueness for the actual set under the B-reading are given in (9).

(9) Liz underlined the wrong number. (B)

a. ✓4 7 9 || 4 7 9
b. ✗4 7 9 || 4 7 9
c. # 4 7 9 || 4 7 9
d. ✓4 7 9 || 4 7 9

In (9a), Liz underlined exactly the one number she was not allowed to underline. The actual set in (9b) is a singleton as well. As it does not contain the forbidden number it is judged as false. In (9c), Liz did not underline any number. The sentence is perceived as odd in this context, which indicates a violation of an existence presupposition. In (9d), Liz underlined two numbers, one of which being the forbidden number. According to my intuition, the sentence is true in this scenario under a B-reading.

What about the excluded set? Scenarios in which existence or uniqueness of the excluded sets is not given are shown in (10). In both cases, the sentences are odd under a B-reading.

(10) Liz underlined the wrong number. (B)

a. # 4 7 9 || 4 7 9 (no excluded number)
b. # 4 7 9 || 4 7 9 (more than one excluded number)

Given these observations, only the excluded set comes with a definiteness presupposition in the B-reading. The actual set is required to be non-empty, and the required set is irrelevant.

2.2 The referent of *the wrong N*

In this subsection, I will argue that the differences in definiteness presuppositions between the P- and the B-readings also have a reflex in the referent of noun phrases of the form *the wrong N* (*tw*NP). To show this, I will look at the anaphoric potential of such noun phrases

In (11), I use different continuations for sentence (1) above. The sentence could potentially provide two antecedents for a pronoun: the actually opened bottle (indexed as *a*) or the bottle (or bottles) that should have been opened (indexed as *b*).

(11) Alex opened the wrong bottle. (P)

a: the bottle that Alex opened

b: the bottle(s) that Alex should have opened.

a. Unfortunately, its$_a$ cork broke.

b. #Unfortunately, Alex didn't find it$_b$/them$_b$ in the cellar.

This shows that reference to the actually opened bottle is possible, see (11a). Reference to the required set is excluded, independently of whether this is done via a singular or a plural pronoun, see (11b).

For the B-reading, the pronominalization shows an analogous asymmetry, now between the excluded set and the actual set. This is a bit more difficult to show, as the excluded set is part of the actual set. Therefore, we can only find a difference in cases in which the actual set is non-unique.

I provide such a scenario for sentence (2) in (12). The continuations show that it is possible to refer to the forbidden number, (12a), but not to the underlined numbers, see (12b).

(12) Liz underlined the wrong number. (B)

4 7 9 || 4 7 9

a: the number Liz was forbidden to underline, i.e. *7*

b: the number(s) Liz actually underlined, i.e. *7* and *9*

a. It$_a$ was a prime number.

b. #It$_b$ was an odd number. / # They$_b$ were odd numbers.

These observations show that a *tw*NP refers to the entity on which it imposes a definiteness presupposition, i.e., the actual set in the P-reading and the excluded set in the B-reading.

We can relate this to a general property of adjectives. At least in English, adjectives cannot introduce an antecedent. This is illustrated by the contrast in (13). The noun phrases in subject position both refer to an envoy and require the existence of a president. However, only in (13a) can the president be used as the antecedent of a pronoun in the sentence. If the president is merely introduced inside the adjective *presidential*, such a coreference is impossible, see (13b).

(13) a. [The envoy of the president$_i$] ...
b. *[The president$_i$-ial envoy] ...
informed him$_i$ about the state of the negotiations.

The data in (13) are important for our discussion of *wrong*. If we assume that the set that lacks a uniqueness presupposition is contributed sub-lexically by *wrong*, the observed contrast in pronominalization follows by the generalization that is independently needed for adjectives.

2.3 The type of definiteness

I have shown in §2.1 that both readings come with a uniqueness and an existence presupposition on one element, i.e. both have one element that can be considered definite in the sense of Schwarz (2020). It is known that there are various types of definite: at least *strong* and *weak* (Schwarz 2009; Löbner 2011; Am-David 2014; Ortmann 2014). Simplifying, strong definites (or pragmatic uniques) typically refer to entities that have been introduced in the previous discourse. Weak definites (or semantic uniques) refer to entities that are given in the background.

English does not distinguish formally between strong and weak definites. In order to determine the type of definiteness that we find in *tw*NPs, I will use data from Papiamentu, a language in which such a distinction is found. Papiamentu is a Portuguese/Spanish/Dutch-based creole language spoken on Aruba, Bonaire and Curaçao (Maurer 2013). Papiamentu has both a definite article, *e*, and an indefinite article, *un*. The definite ar-

ticle is used for strong definites, i.e., primarily in anaphoric contexts, see (14). No article is used, however, with weak definites like *solo* 'sun' in (15).

(14) *Mi a kumpra un bolo. *(E) bolo a wòrdu kome den 10 minüt.*
I PERF buy a cake the cake PERF PASS eat in 10 minutes
'I bought a cake. The cake was eaten in 10 minutes.'
(Kester & Schmitt 2007: 119)

(15) *(*E) solo ta kima sin miserikordia.*
the sun PRES burning without mercy
'The sun is burning without mercy.' (Kester & Schmitt 2007: 113)

Papiamentu does not use the definite article with its equivalent of English *wrong* in the relevant uses. This is illustrated in (16) and (17).

(16) *Polis a arestá hende robes pa Interpol.*
police PERF arrest person wrong for Interpol
'The police has arrested the wrong person for Interpol.'[2]

(17) *Ta duel mi. Señor a yama number robes.*
PRES hurts me Mister PERF call number wrong
'I am sorry. You have the wrong number, Sir.'[3]

This shows that the noun phrases *hende robes* and *number robes* are weak definites. Consequently, their uniqueness and existence should be treated on par with that of other weak definites, i.e., it is presupposed in the context and need not be introduced explicitly. Therefore, I conclude *robes* 'wrong' has the effect of turning a noun into a weak definite.

I think it is legitimate to generalize this to English *wrong*. In English, weak definites are marked with the definite article. We, thus, have an explanation why non-local *wrong* requires the definite article in English. I will come back to this point in §4.

[2] `https://extra.cw/polis_a_aresta_hende_robes_pa_interpol/`, 2020/04/08.
[3] *Papiamento extended phrasebook*, Learningonlinexyz Inc. Accessed via googlebooks, 2020/04/08.

2.4 Modality

Every paraphrase of a sentence with *wrong* contains a modal expression. I will adopt the three dimensions of modality from Kratzer (1977; 1991) (modal force, modal base, and ordering source) to discuss the type of modality found with *wrong*. As far as the modal force is concerned, the terms *required* and *excluded* set in Schwarz (2020) indicate that we are dealing with a necessity modality in all cases. I will show that *wrong* is compatible with circumstantial (or root) modality, but not with epistemic modality.

Wrong can occur felicitously with circumstantial necessity. In (18a), the obligation is imposed by laws or regulations. In (18b), it comes from moral, ethical or other considerations rather than from strict rules.

(18) a. [The university obliges us to use a particular cloud service, but Alex is using a different one.]
So, Alex is using the wrong cloud service. (P)

b. [We all know that we should reduce CO_2 emission and move away from coal power generation. Nonetheless, the government has just approved a new coal-fired power station.]
So, the government supports the wrong type of energy. (B)

This contrasts with epistemic necessity. In (19), I provide contexts for a potential use of *wrong* with epistemic necessity. Below each example, I indicate the intended reading. As shown by the marking "#" such readings are not possible – neither for the P-reading nor for the B-reading.

(19) a. [From what I know about Alex and Kim, Alex must be on vacation now, but, in reality, Kim is.]
So, the wrong person seems to be on vacation. (P)
Intended: 'The person that is on vacation is not among the people who must be on vacation according to what I know.'

b. [I was sure that Alex would not pass the biology test. However, Alex did fairly well in it.]
So, Alex passed the wrong test. (B)
Intended: 'The test that Alex must fail according to my knowledge is among the tests that Alex passed.'

The examples of circumstantial modality above contained some possible ordering sources. While they suggest that the obligation is imposed on the grammatical subject of the sentence, this is not necessarily the case. We find the same readings when the logical subject is implicit, as in passives without overt *by*-phrase, see (20).

(20) a. [(18a)] So, the wrong cloud service is used.
b. [(18b)] So, the wrong type of energy is supported.

In bouletic modality, the person whose wishes or desires are at stake need not be overtly expressed. Nonetheless, we find cases of *wrong*, as in (21). In these examples, *wrong* relates to the wishes of Alex in a particular lottery, even though Alex is not mentioned explicitly in the sentence.

(21) a. [Alex would win in the lottery if the number *4* was drawn, but the number *7* was drawn.]
So, clearly, the wrong number was drawn. (P)
b. [Alex would win in the lottery unless the number *4* was drawn. However, this number was drawn.]
So, clearly, the wrong number was drawn. (B)

This preliminary discussion shows that *wrong* comes with (possibly all types of) circumstantial necessity, but not with epistemic modality.

2.5 Semantic representations

Taking together the observations from this section, we arrive at the semantic representations given in (22) and (23) for slightly simplified versions of sentences (3) and (4), respectively.

(22) Lestrade arrested the wrong person. (P)
$\neg((\iota_s x : \mathbf{pers}(x) \wedge \mathbf{arr}(\mathbf{l}, x)) \in \{x | \mathbf{pers}(x) \wedge \mathbf{OBL}(y, {}^{\wedge}\mathbf{arr}(\mathbf{l}, x))\}_{\geq 1})$

(23) Anne opened the wrong door. (B)
$(\iota_s x : \mathbf{door}(x) \wedge \mathbf{OBL}(y, {}^{\wedge}\neg\mathbf{op}(\mathbf{a}, x))) \in \{x | \mathbf{door}(x) \wedge \mathbf{op}(\mathbf{a}, x)\}_{\geq 1}$

I will first go through the representation of the P-reading in (22). The highest operator is the negation of the membership relation. There is an ι-expression that refers to the person that Lestrade arrested. Assuming a standard semantics for this operator, this reflects the existence and

uniqueness presupposition on the actual set discussed in §2.1 and provides a potential antecedent for pronominal reference in the discourse, see §2.2. I use the subscript "*s*" on the ι-operator to indicate semantic uniqueness (§2.3). The relevant set is the set of people that Lestrade should arrest. The choice of set-membership rather than identity expresses the non-uniqueness of the required set (§2.1). The subscript "≥ 1" abbreviates the presupposition that this set is not empty.

The modal operator **OBL** expresses a circumstantial modality (§2.4). Consequently, it has a propositional argument and an individual argument expressing whose obligation is considered. This individual argument need not be linked to an overt element in the sentence, but is determined contextually. This is expressed with a free variable, *y*, in (22).

The representation for the B-reading in (23) consists of the same ingredients as the one for the P-reading, but they are arranged differently. In the P-reading, the negation has wide scope over the set-membership, and the modal operator **OBL** occurs inside the set. In the B-reading, the set-membership is the highest operator. The modal operator does not occur inside the set and has scope over the negation. The ι-expression refers to the door that was required not to be opened, i.e., to the unique member of the excluded set. This entity is said to be a member of the actual set, which is not necessarily a singleton, though its non-emptyness is presupposed. I will leave out the subscripts "*s*" and "≥ 1" in the following.

The proposed semantic representations capture the data on the two readings of *wrong*. In the next section, I will demonstrate that both readings qualify as *non-local* in the sense of Schwarz (2020).

3 Non-locality of the readings

The P- and the B-readings of *wrong* can be considered non-local as the meaning of the verb appears embedded inside a meaning contribution of the adjective. This criterion has been generally applied in the literature, including Morzycki (2016). Schwarz (2020) proposes two entailment tests for non-local adjectival modifiers: *extensionality* and *monotonicity*. A non-local adjective allows for neither an extensional nor a monotone entailment in the way to be described below. I will first illustrate the two tests for the P-reading and, then, apply them to the B-reading.

I will use the two tests from Schwarz (2020), but modify some of his ex-

amples. I will contrast the local adjective *red* with non-local *wrong*. Let us assume a situation in which whenever someone underlines a number they also put a circle around it, i.e., in the current situation s, $[\![\textbf{underline}]\!]^s = [\![\textbf{circle}]\!]^s$. In this situation, sentence (24a) entails sentence (24b).

(24) a. Liz underlined the red number.
b. $\models$ Liz circled the red number.

Imagine a situation in which Liz was supposed to underline the number *4* but circle the number *7*, and in which she underlined (and circled) *7*. Then sentence (25a) is true, but (25b) is false. This shows that the P-reading is not extensional.

(25) 4 (7) 9 | 4 (7) 9 (P)
a. Liz underlined the wrong number.
b. $\not\models$ Liz circled the wrong number.

The second test looks at monotonicity. Underlining a number is a way of marking it. Consequently, sentence (26a) entails (26b).

(26) a. Liz underlined the red number.
b. $\models$ Liz marked the red number.

In (27), we replace the local adjective *red* with *wrong*. The entailment from (27a) to (27b) does not hold in general. In the scenario in (27), Liz was supposed to underline *4* and to cross out *7*. She underlined *7*. Consequently, sentence (27a) is true, but sentence (27b) is false – because *7* was among the numbers to be marked.

(27) 4 7 9 | 4 7 9 (P)
a. Liz underlined the wrong number.
b. $\not\models$ Liz marked the wrong number.

This shows that a *tw*NP in the P-reading does not conserve the monotonicity of a definite noun phrase of the form *the N* or *the red N*. From the results of these two tests, Schwarz (2020) concludes that *wrong* is a non-local adjective in the P-reading, which is the only one he discusses.

In the next step, I will show that the B-reading is equally non-local ac-

cording to these tests. The extensionality test is exemplified in (28). Let us assume, again, that underlining and circling have the same extension in our situation. Furthermore, Liz was only forbidden to underline the number *4*, but allowed to circle whichever numbers she likes. She actually underlined and circled *4*. Then (28a) is true, but (28b) is false.

(28) ④ ⑦ ⑨ ‖ ④ 7 9 (B)

a. Liz underlined the wrong number.
b. ⊭ Liz circled the wrong number.

For monotonicity, consider a scenario in which Liz was allowed to underline any number except for *4* and to cross out any number except for *7*. She underlined *4*. Then, (29a) is true. Sentence (29b), however, is undefined as there is no unique number that must not be marked.

(29) 4̸ 7 9̸ ‖ 4 7 9 (B)

a. Liz underlined the wrong number.
b. ⊭ #Liz marked the wrong number.

We can use a plural noun phrase to overcome this problem. The scenario in (30) is chosen in such a way that the excluded set for underlining consists of *4*, *7*, and *9*. The excluded set for marking only consists of *7* and *9*. Liz underlined exactly the numbers that she must not underline, which makes sentence (30a) true under the B-reading. However, sentence (30b) is false, as the number *4* is not in the excluded set for marking.

(30) 4̸ 7 8̸ 9 ‖ 4 7 8 9 (B)

a. Liz underlined the wrong numbers.
b. ⊭ Liz marked the wrong numbers.

This shows that both the P-reading and the B-reading as defined in §2 pass the tests for non-local adjectives in Schwarz (2020).

4 Challenges for previous approaches

An important issue in previous discussions of non-local *wrong* is the observation that we always find a definite article. This is even more puzzling under the classical paraphrase in (1) – repeated in (31) – as this paraphrase

does not contain a definite noun phrase.

(31) I opened the wrong bottle of wine.
'I opened a bottle of wine that it was wrong for me to open.'

Consequently, Abbott (2001) and Schwarz (2006) assume that a *tw*NP is semantically indefinite. In support of this, Abbott (2001: 12) provides example (32). While definites cannot occur in existential *there*-clauses, this seems to be possible for *tw*NPs.

(32) There was the wrong address on the envelope.

This argument is not fully conclusive. First, there are occurrences of definites with existential *there*, as in (33) from COCA (Davies 2008–).

(33) There was my wife in the living room. (COCA)

Second, uses of *tw*NPs in *there*-clauses is far from common, if existing at all. For instance, there is no relevant hit parallel to (32) in COCA for the query '`there BE the wrong _nn*`' (2020/04/08).

In Schwarz (2006), *the wrong* is simply treated as one lexical item. It is, however, possible to find examples in which there is material between the article and the adjective, see (34).

(34) Archaeologists, who have spent decades digging at the apparently wrong location, will soon be moving to the new site.[4]

Morzycki (2016) shows that many non-local adjectives require a definite article, such as *average* in (35). He suggests that this is due to a kind reading of nouns with non-local adjectives. For kinds, the use of a definite article is to be expected.

(35) The average American has 2 children.

While this is plausible for *average*, Morzycki himself states that this explanation cannot be applied to non-local *wrong*. For example, sentence (2) is about a concrete number, not about kind of number, and (3) is about

[4] `https://tinyurl.com/y5fta3qw`, accessed 2020/09/07

a concrete person, not a kind of person, etc.

Instead, Morzycki (2016) claims that there is no definiteness requirement and that there are non-local readings of *a wrong N* as well. I will argue in §6 that such uses, indeed, are local readings.

Larson (2000) suggests that the definiteness requirement of *wrong* can be captured by assigning it a superlative semantics. As English superlatives come with a definite article, it should not be a surprise that the same holds for non-local *wrong*. However, Larson does not provide such a superlative-like semantics.

Based on Papiamentu data, I argued in §2.3 that *wrong*-noun phrases behave like unique nouns such as *sun*. This seems to be a cross-linguistically robust generalization, as in English, unique nouns require a definite article, and so do noun phrases with non-local *wrong*. Thus, Larson (2000) is correct in pointing out that superlatives and non-local *wrong* share an important semantic property which is responsible for the parallelism in definiteness marking. However, the relevant property is not the superlative semantics but the semantic uniqueness they both express.

Let us finally look at the proposal in Schwarz (2020). Schwarz works within a framework such as Heim & Kratzer (1998), in which functional application is the central device for computing the meaning of a complex expression. Consequently, non-local readings are a serious challenge as the adjective cannot be interpreted directly in its surface position. For this reason, syntactic operations are postulated to adjust the syntactic structure to the semantics. Schwarz eventually favours what he calls a *main functor analysis*, under which the adjective acts as the highest semantic functor in the clause. Following Morzycki (2016), there are two instances of Quantifier Raising to arrive at an interpretable syntactic structure, see (36). First, the noun phrase *the wrong number* is fronted. Then, the combination *wrong number* is fronted further, stranding the determiner.

(36) Liz underlined the wrong number.
QR1: [the [wrong number]] λ_2 [Liz underlined t_2]
QR2: [wrong number] λ_1 [[the t_1] λ_2 [Liz underlined t_2]]

While the movements in (36) are required within the particular framework, I think that there is little independent motivation for them. The ad-

jective *wrong* is in the surface position in which we find attributive adjectives, and determiners cannot be stranded in English, see (37). The Morzycki/Schwarz analysis might be a solution within their framework – still, I think that an approach stressing the parallels between *wrong* and "ordinary" adjectives is conceptually more attractive.

(37) *[Wrong number]$_1$ Liz underlined [the t_1]

In this section, I have mentioned challenges of previous approaches and, in part, already specified how they can be resolved in the analysis proposed in §2. In particular, the questions of whether non-local *wrong* is definite at all and is always definite have both received a positive answer in this paper and an explanation in terms of treating *wrong* as an adjective that creates semantic uniques. In the next section, I will formalize my analysis in a constraint-based syntax-semantics interface.

5 Underspecified semantics of *wrong*

In this section, I will present an integration of my analysis of non-local *wrong* into a formal framework of the syntax-semantics interface, *Lexical Resource Semantics* (LRS). I will present the basic ideas of this framework in §5.1 and develop my analysis in §5.2.

5.1 Framework: Lexical Resource Semantics

LRS is a formal system of the syntax-semantics interface. It is representational in the sense that it assumes that linguistic expressions have (at least) a syntactic and a semantic representation. Semantic representations can be expressions of any standard semantic representation language. In the present paper, I will use the ones from the previous sections.

LRS is a system of constraint-based underspecified semantic combinatorics. This means that words and phrases add constraints on what the eventual semantic representation of an utterance should be. It is underspecified in the sense that these constraints need not fix exactly one semantic representation but could be compatible with several readings. As such, LRS is in the tradition of underspecified semantic systems as characterized in Pinkal (1999) and Egg (2010).

LRS has mainly been used in Head-driven Phrase Structure Grammar (HPSG, Pollard & Sag 1994), as HPSG is a constraint-based grammar frame-

work that integrates all modules of grammar within one formalism. There are, however, LRS analyses that are independent of a particular grammar framework, including Sailer (2004a) and the present paper. A general introduction to LRS is given in Richter & Sailer (2004).[5] I will use a version of the compact notation introduced in Penn & Richter (2004).

LRS is lexical in the sense that only lexical items, i.e. words or phrasal lexical units, determine which constants, variables, and operators may occur in the semantic representation of an utterance (*contribution constraints*). Non-lexical items can only constrain how these should be combined to arrive at the overall semantic representation (*embedding constraints*). I will illustrate this with simple example sentences, such as (38). I indicate the constraints contributed by the words below the sentence.

(38) [S: Everyone [VP: didn't call]].

a. *call*: **call**(x)
b. *didn't*: $\neg\alpha$
c. *everyone*: $\forall x(\textbf{person}(x) \rightarrow \beta[x])$

The constraint in (38a) is to be read in the following way: Whenever the word *call* occurs in an utterance, the semantic representation must contain an occurrence of the formula **call**(x).

Following Bos (1996), I assume a semantic meta-language. I use lowercase Greek letters for meta-variables ($\alpha, \beta, \ldots$). The constraint contributed by the negated auxiliary *didn't* in (38b) restricts its use to utterances whose semantic representation contain a negation. The scope of the negation is marked with the meta-variable α.

The word *everyone* has the most complex constraint. It specifies that whenever it occurs, there will be a universal quantifier that binds a variable and has an implication in its scope. The antecedent of this implication is of the form **person**(x), and its consequent is marked with a meta-variable, β, which means that it is not fully constrained by *everyone*. The notation $\beta[x]$ expresses that whatever expression β will be interpreted as, it must have an occurrence of x in it.

These words combine syntactically in the way indicated by the bracketing in (38). The VP *didn't call* collects the constraints from (38a) and (38b).

[5] See `https://www.lexical-resource-semantics.de` for further material on LRS.

In addition, a constraint is added that the meaning of *call* occurs in the scope of the negation contributed by *didn't*. This is written as $\alpha[\mathbf{call}(x)]$ All these constraints are collected in (39). There, a new meta-variable, γ, is introduced, which must be some formula that satisfies the three constraints given inside the square brackets.

(39) VP: $\gamma[\mathbf{call}(x), \neg\alpha, \alpha[\mathbf{call}(x)]]$

The S-node collects the constraints from the subject and the VP node. In addition, it adds the requirement that the meaning of *call* must occur in the consequent of the implication contributed by *everyone*, which can be expressed as $\beta[\mathbf{call}(x)]$. The overall constraint is given in (40), where I introduce a new meta-variable, δ.

(40) S: $\delta[\forall x(\mathbf{person}(x) \rightarrow \beta[x]), \gamma[\mathbf{call}(x), \neg\alpha, \alpha[\mathbf{call}(x)]], \beta[\mathbf{call}(x)]]$

Once all constraints are gathered, there is a closure constraint saying that semantic representation of an utterance can only contain the constants, variables, and operators that occur in the constraints contributed by lexical items and that it must respect all constraints contributed by the lexical and non-lexical items contained in the utterance.

We can arrive at the overall semantic representation of an utterance by assigning each contributed meta-variable some expression in such a way that all constraints are satisfied. Such a meta-variable assignment is called a *plugging* (Bos 1996). For ambiguous sentences, there should be more than one plugging, which then leads to more than one possible semantic representation. In our example, there are two possible pluggings that respect all constraints. These are given in (41). In (41a), the negation has narrow scope. In (41b), it has wide scope over the universal quantifier.

(41) Possible pluggings

a. $\alpha \equiv \mathbf{call}(x)$; $\beta \equiv \gamma \equiv \neg\alpha$; $\delta \equiv \forall x(\mathbf{person}(x) \rightarrow \beta)$
Reading 1: $\forall x(\mathbf{person}(x) \rightarrow \neg\mathbf{call}(x))$

b. $\alpha \equiv \gamma \equiv \delta \equiv \forall x(\mathbf{person}(x) \rightarrow \beta)$; $\beta \equiv \mathbf{call}(x)$
Reading 2: $\neg\forall x(\mathbf{person}(x) \rightarrow \mathbf{call}(x))$

This example illustrated how the lexical specifications determine the

resulting readings together with the additional constraints added at the phrases. In order to formulate the phrase-level constraints in a systematic way, we flag certain contributions. For the present paper, the relevant contributions are the *internal content* and the *external content*. The internal content, which I will mark as $\lceil\alpha\rceil$, signals the scopally lowest contribution in a phrase. The external content, indicated as $\underline{\alpha}$, is the representation associated with the overall phrase. In (42), I repeat the lexical entries from (38), augmented by the indication of internal and external content.[6]

(42) a. *call*: $\lceil\mathbf{call}(x)\rceil$
b. *didn't*: $\neg\alpha[\lceil\alpha'\rceil]$
c. *everyone*: $\underline{\forall x(\lceil\mathbf{person}(x)\rceil \rightarrow \beta[x])}$

With these two auxiliary notions, we can define the constraints that I had used in the first run through example (38). First, in every headed phrase, the internal content and the external content both percolate from the head daughter to the mother. Second, when a raising verb, such as the auxiliary *didn't*, combines with its verbal complement, the auxiliary inherits its complement's internal content. Consequently, the internal content of the VP *didn't call* is $\mathbf{call}(x)$, which leads to the above-mentioned constraint $\alpha[\mathbf{call}(x)]$. The third general constraint applies when a quantifier is the non-head in a phrase. In this constellation, the quantifier takes scope over the internal content of the head. This constraint has the effect that $\mathbf{call}(x)$ must be a component of β, i.e., $\beta[\mathbf{call}(x)]$.

As a second example, I will discuss (43), which contains an attributive adjective and the definite article.[7]

(43) [S: Alex [VP: opened [NP: the [N′: red bottle]]]].
a. *bottle*: $\lceil\mathbf{bottle}(x)\rceil$
b. *red*: $\underline{(\alpha[x] \wedge \beta[\lceil\mathbf{red}(x)\rceil])}$
c. *the*: $\underline{(\lceil\iota\rceil x : \phi[x])}$

[6]The specifications in (42) are simplifications. All all signs have an internal and an external content, but I have left out some of the required meta-variables.

[7]Given HPSG's lexical approach to argument linking, the discourse referents of the arguments of **open** are constrained to occur inside the appropriate argument slots. For the subject, I simplify this to **alex**. I am more explicit for the complement: the second argument of **open** is an expression χ containing the complement's discourse referent x.

d. *opened*: $[\mathbf{open}(\mathbf{alex}, \chi[x])^{]}$
e. *Alex*: $\underline{[\mathbf{alex}^{]}}$

First, the noun *bottle* combines with the adjective *red*. The external content of the adjective is a conjunction. The second conjunct contains the adjective's internal content, $\mathbf{red}(x)$. The combination of these two words is subject to the constraint in (44), introduced in Sailer (2004b).

(44) In a head-modifier combination, if the external content of the modifier is of the form $\alpha \wedge \beta$, the head's internal content is a subexpression of α and the modifier's external content is a subexpression of the head's external content.

In our example, this has the effect that the constraint $\alpha[\mathbf{bottle}(x)]$ is added, which leads to the overall constraint in (45).

(45) N′: $\alpha[^{[}\mathbf{bottle}(x)^{]}] \wedge \beta[\mathbf{red}(x)]$

Following the HPSG tradition, I assume that the determiner is the nonhead in the next combination. The constraint in (46) is relevant here, which embeds the noun's internal content inside the determiner's restrictor.

(46) When a determiner combines with a nominal head, the determiner and the head have the same external content, and the head's internal content is embedded in the determiner's restrictor.

Given this constraint, we arrive at (47) for the noun phrase.

(47) NP: $\underline{(\iota x : \phi[\alpha[^{[}\mathbf{bottle}(x)^{]}] \wedge \beta[\mathbf{red}(x)]])}$

The rest of the sentence does not require any new principles of grammar. As the definite noun phrase is not qunantificational, no embedding constraint will be added at the VP level. This leads to the overall constraint in (48).

(48) VP: $\gamma[^{[}\mathbf{open}(\mathbf{alex}, \chi[x])^{]}, (\iota x : \phi[\alpha[\mathbf{bottle}(x)] \wedge \beta[\mathbf{red}(x)]])]$

Finally, the subject is added. As the sentence is not ambiguous, there is only one plugging satisfying all constraints. This plugging is given in (49).

(49) $\alpha \equiv \mathbf{bottle}(x)$; $\beta \equiv \mathbf{red}(x)$; $\phi \equiv (\alpha \wedge \beta)$; $\chi \equiv (\iota x : \phi)$; $\gamma \equiv \mathbf{open}(\mathbf{alex}, \chi)$
Reading: $\mathbf{open}(\mathbf{alex}, (\iota x : \mathbf{bottle}(x) \wedge \mathbf{red}(x)))$

I showed with example (43) how intersective, attributive adjectives and definite noun phrases are handled in LRS.

Finally, I want illustrate the analysis of a semantically unique noun, based on Sailer & Am-David (2016). As a unique noun, *sun* introduces an ι-operator lexically, see (50).

(50) *sun*: $\underline{(\iota x : [\mathbf{sun}(x)])}$

All English singular count nouns require an overt determiner syntactically, therefore the noun *sun* selects a determiner. In particular, it can combine with the definite article, whose constraint we saw in (43c). This combination is subject to (46), i.e., the noun and the article have the same external content and the noun's internal content must be in the restrictor of the determiner, which is the body of the ι-expression. The result is given in (51). Assuming $\phi \equiv \mathbf{sun}(x)$, this reduces to the expression $(\iota x : \mathbf{sun}(x))$.

(51) [NP: the sun]: $\underline{(\iota x : \phi[\mathbf{sun}(x)])}$

It is important to note that both the noun and the determiner constrain the overall semantic representation to contain an ι-expression. There is nothing requiring, however, that there need to be two ι-expressions. This potential of *redundant semantic contributions* is one of the key properties of LRS and has been exploited in analyses of negative concord, tense marking, and others (Richter & Sailer 2006; Sailer 2004a).

All constraints used in this subsection were proposed in previous LRS papers. I will show below, that non-local *wrong* combines with the head noun in exactly the same way as *red*, and result in a semantically unique noun such as *sun*.

5.2 Analysis of non-local attributive *wrong*

We can now turn to the LRS formalization of the analysis of non-local *wrong* developped in §2. The constraint associated with the lexical entry of *wrong* is given in (52).

(52) Lexical constraints of attributive *wrong*:

$\zeta[\alpha[x] \wedge \beta[x]$, (i)
$\overline{(\iota x : \alpha \wedge \gamma') \in \{x|(\alpha \wedge \gamma'')\}}$, (ii)
$\neg\delta[\beta]$, (iii)
$\alpha \wedge \mathbf{OBL}(y, {}^{\wedge}\epsilon[\beta])]$ (iv)
and neither $\gamma' \equiv \neg\delta$ nor $\gamma''[\neg\delta]$

This constraint introduces a meta-variable, ζ, for which four components are specified. First, (i), the external content of the adjective is a conjunction just like the external content of an ordinary intersective adjective like *red* in (43b). The first conjunct, α, will eventually be the content of the head noun, the second conjunct, β, will be the content of the clause in which the noun phrase occurs. Second, (ii), there is a set-membership expression whose first argument is an ι-expression and whose second argument is a set. The body of the ι-expression is, again, a conjunction with α as its first conjunct and some expression γ' as its second conjunct. Similarly, the body of the set is of the form $\alpha \wedge \gamma''$. Third, (iii), there is a negation that takes scope over β. Fourth, (iv), there is another conjunction with α as its first conjunct and an obligation modal in its second conjunct which also takes scope over β

Finally, there is a further condition in (52) which determines that only pluggings are acceptable in which neither γ' is the negation $\neg\delta$ nor does γ'' contain the negation.

The constraint in (52) does not fully specify the relative scope of the four mentioned components of ζ. The membership relation specifies two conjunctions with α as their first conjunct, and there cannot be more occurrences of such conjunctions in ζ. Consequently, either $\beta \equiv \gamma'$ or $\beta \equiv \gamma''$. In the first case, we know that $\gamma'' \equiv \mathbf{OBL}(y, \epsilon)$, in the second case, $\gamma' \equiv \mathbf{OBL}(y, \epsilon)$. The negation is also restricted: it must have scope over β, but it may not be the second conjunct in the body of the ι-expression nor may it occur in side the body of the set. This leaves only two options: it can have wide scope over the set-membership expression, $\zeta \equiv \neg\delta$, or it can have narrow scope inside the second conjunct of the body of the ι-expression if there is another operator above it – which can only be the modal operator in our case. Taking these considerations together, we end up with exactly two possibilities, which are given schematically in (53). As indicated, these correspond to the P- and the B-reading.

(53) a. $\neg((\iota x : \alpha \wedge \beta) \in \{x|\alpha \wedge \mathbf{OBL}(y, {}^{\wedge}\beta)\})$ (P)
$\beta \equiv \gamma' \equiv \epsilon; \quad \gamma'' \equiv \mathbf{OBL}(y, \epsilon);$
$\delta \equiv (\iota x : \ldots) \in \{x|\ldots\}; \quad \zeta \equiv \neg\delta$

b. $(\iota x : \alpha \wedge \mathbf{OBL}(y, {}^{\wedge}\neg\beta)) \in \{x|\alpha \wedge \beta\}$ (B)
$\beta \equiv \gamma' \equiv \mathbf{OBL}(y, \epsilon); \quad \gamma'' \equiv (\alpha \wedge \beta);$
$\delta \equiv \beta; \quad \epsilon \equiv \neg\delta; \quad \zeta \equiv (\iota x : \ldots) \in \{x|\ldots\}$

I can now show how the two P- and the B-readings of an example similar to (1) can be derived. I will use the lexical entries from (43). As indicated in (54), I assume the same syntactic structure as in the case with a local intersective adjective.

(54) [S: Alex [VP: opened [NP: the [N′: wrong bottle]]]].

When *wrong* combines with the head noun *bottle*, the constraint in (44) applies. In other words, we combine the constraints from (43a) and (52) and add the constraint that the internal content of *bottle* be inside the first conjunct of the external content of the adjective, i.e., $\alpha[\mathbf{bottle}(x)]$.

I described the lexical entry of *the* in (43c) and the effect of it combining with a noun in (51). In the present example, the constraint in (55) is added.

(55) $\underline{\iota x : \phi[{}^{\lfloor}\mathbf{bottle}(x)^{\rfloor}]}$

When we combine the constraints of the three words occurring in the noun phrase and the ones added by the phrases, we arrive at (56).

(56) Accumulated constraints for *the wrong bottle*:
$\zeta[(\iota x : {}^{\lfloor}\mathbf{bottle}(x)^{\rfloor} \wedge \gamma'[\beta[x]]) \in \{x|\mathbf{bottle}(x) \wedge \gamma''[\beta]\},$
$\quad \underline{\mathbf{bottle}(x) \wedge \beta},$
$\quad \neg\delta[\beta]],$
$\quad \mathbf{bottle}(x) \wedge \mathbf{OBL}(y, {}^{\wedge}\epsilon[\beta]))$

Since both the adjective and the determiner contribute an ι-expression, we are in exactly the same situation as with unique nouns in (50). Both expressions constrain the overall semantic representation to contain an ι-expression, which is compatible with there being just one such expression in the overall representation.

In the next step, the noun phrase combines with the verb *opened*. Being

a definite noun phrase, it is not quantificational. Consequently, there is no new scopal constraint. We just add the lexical constraint of the verb from (43d), $\mathbf{open}(\mathbf{a}, \chi[x])$. To complete the sentence, there is only one meta-variable in the overall constraint that can be equated to this formula: β.

(57) Accumulated constraints for sentence (54):
$\zeta[(\iota x : \mathbf{bottle}(x) \wedge \gamma'[\lfloor \mathbf{open}(\mathbf{a}, x) \rfloor]) \in \{x | \mathbf{bottle}(x) \wedge \gamma''[\mathbf{open}(\mathbf{a}, x)]\},$
$\mathbf{bottle}(x) \wedge \mathbf{open}(\mathbf{a}, x),$
$\neg\delta[\mathbf{open}(\mathbf{a}, x)]],$
$\mathbf{bottle}(x) \wedge \mathbf{OBL}(y, {}^\wedge\epsilon[\mathbf{open}(\mathbf{a}, x)]))$

We saw in the abstract discussion of which pluggings are compatible with the lexical specification of non-local *wrong* that there are only two possibilities. These are given in (58) and (59) together with the resulting semantic representations, the P- and the B-reading respectively.

(58) $\neg((\iota x : \mathbf{bottle}(x) \wedge \mathbf{open}(\mathbf{a}, x)) \in \{x | \mathbf{bottle}(x) \wedge \mathbf{OBL}(y, {}^\wedge\mathbf{open}(\mathbf{a}, x))\}$
$\gamma' \equiv \epsilon \equiv \mathbf{open}(\mathbf{a}, x); \gamma'' \equiv \mathbf{OBL}(y, {}^\wedge\epsilon); \delta \equiv ((\iota x : \ldots) \in \{x | \ldots\}); \chi \equiv x$

(59) $(\iota x : \mathbf{bottle}(x) \wedge \mathbf{OBL}(y, {}^\wedge\neg\mathbf{open}(\mathbf{a}, x))) \in \{x | \mathbf{bottle}(x) \wedge \mathbf{open}(\mathbf{a}, x)\}$
$\gamma' \equiv \mathbf{OBL}(y, {}^\wedge\epsilon); \gamma'' \equiv \delta \equiv \mathbf{open}(\mathbf{a}, x); \epsilon \equiv \neg\delta; \chi \equiv x$

Note that throughout the derivation, the individual argument of the modal operator **OBL** is a free variable. As it is not a meta-variable, it will not be resolved by a plugging, but needs to be resolved in the context.

There are some important features of this analysis. First, there is no combinatorial difference – neither syntactically nor semantically – between the adjectives *red* and *wrong*. Both have an intersective core semantic contribution, their external content, and are subject to the constraint in (44) when combining with a noun.

In addition to this core semantic contribution, *wrong* also contributes other constraints. This is the material responsible for the non-local reading. It is scopally restricted within the lexical entry of the adjective, but irrelevant for the constraints contributed by the syntactic combinatorics.

Just as other definite noun phrases in English, a *tw*NP refers to an individual with an existence and uniqueness presupposition. In English, the definite article is semantically redundant for unique nouns, but fulfills this role in such noun phrases just as in any other definite noun phrase.

Consequently, we capture the parallelism to unique nouns in English and other languages. Finally, there is a single lexical constraint on non-local *wrong* which is underspecified with respect to the P- and the B-readings.

6 Local reading of attributive *wrong*

Previous studies of attributive *wrong* are restricted to non-local readings and have largely ignored local readings. However, I would like to stress the connection between local and non-local readings of attributive *wrong*.

Larson (2000) and Schwarz (2006) observe that *a wrong N* does not allow for non-local readings. This is also my impression, based on a cursory inspection of COCA hits for the query *a wrong N*. In (60), the speaker will answer if a number has been dialed that does not exist, i.e., if it is an unassigned number, not whenever it is not the number that one was supposed or intended to dial. Similarly, in (61), the decision is the one that should not have been made, not the one that should not be overturned.

(60) I also do other intercept messages, when you dial a wrong telephone number, or dial a number that's been disconnected, or you need to deposit 25 cents before making a call. That's me.
⇒ a telephone number that (necessarily) doesn't exists

(61) "Institutional integrity" turns out to mean the Court must not overturn a wrong decision if there has been angry opposition to it.
⇒ a decision that should not have been made

The basic idea pursued here is that, in the local reading, we infer what should have been done to the referent of the noun phrase. In (61), for example, an inferrable property should not hold of the decision. This is sketched in the simplified semantic representation in (62), which states that there is a decision such that it is not among the ones that should be made, $P(x)$, and that this decision is overturned.

(62) $\exists x((\mathbf{decision}(x) \wedge (x \in \{x|\mathbf{decision}(x) \wedge \mathbf{OBL}(y, {}^{\wedge}\neg P(x))\}))$
$\wedge \mathbf{overturn}(\mathbf{Court}, x))$, where P can be inferred

In (63), I provide the lexical semantic contribution of local attributive *wrong*. As indicated by the underlining, the entire expression is the external content. This makes it a *local* adjective in the sense that there is no

contribution of the adjective that contains its own external content.

(63) Lexical constraints of local *wrong*:
$\underline{(\alpha \wedge \beta[\mathbf{OBL}(\hat{}\epsilon[P(x)]), \neg\delta[P(x)]])}$, where P can be inferred.

With this lexical specification, local *wrong* is compatible with all determiners, including both the definite and the indefinite article, while non-local *wrong* only allowed for a redundant definite article. Nonetheless, the representation of the local readings shares major parts with that of the non-local readings: negation, and a modal operator. The constraint-based view of LRS allows us to say that the constraints of non-local and local *wrong* overlap to a large extend.

The inferred predicate P is an important difference between local and non-local *wrong*. P is not necessarily identical with the main predicate of the clause. In the non-local readings, the predicate in the scope of the modal operator **OBL** needs to be the same as the predicate in the clause. This latter identity is what gave rise to the impression that *wrong N* takes scope over a VP, as encoded syntactically in the analysis in Haïk (1985). This contrast follows from the fact that P is an object-level predicate in (63), whereas I used a meta-variable in the constraint on non-local *wrong* in (52). The meta-variable needs to be resolved in the plugging as an expression that occurs in the semantic representation, the object-level variable on the other hand will be assigned a value in context.[8]

Schwarz (2020) follows Morzycki (2016) in assuming that *a wrong N* can have a non-local reading, providing example (64). This sentence is true, for example, when Liz underlined the one number that she was supposed to underline and at least one forbidden number.

(64) Liz underlined a wrong number.
⇒ a number that Liz should not underline

However, I think that (64) shows a local reading. As the local readings depend on an inferred predicate, there is nothing wrong with inferring that predicate from the clause itself. Consequently, there would be a po-

[8]This is parallel to the treatment of the individual argument of **OBL** as a free object-level variable in the lexical specification of non-local *wrong* in (52).

tential overlap between the local and the non-local readings.

However, I think that indefinite attributive *wrong* is always local. To show this, I need to construct a context in which the non-local indefinite reading would be true – i.e. the reading in which the clause-mate predicate is what should not have happened – but the local indefinite reading is not. If, in such a context, the use of the indefinite *wrong* noun phrase is considered to make the sentence false, this shows that the non-local reading is not available.

In the scenario in (65), both the definite and the indefinite versions of a *wrong*-sentence are true.

(65) 4 ~~7~~ 9 | 4 7 9

a. ✓Liz underlined the wrong number. (P)

b. ✓Liz underlined a wrong number.

Once we make salient that it is just about marking a number and not about the difference between underlining and crossing out, the version with the definite article, (66a), may sound a bit off-topic. The version with the indefinite article in (66b), however, can no longer be considered true.

(66) A: Liz' task was very complex, so let's just say, we are happy if she marked the numbers *4* and *7* in whichever way.

a. B: ? Still, she underlined the wrong number.

b. B: ✗Still, she underlined a wrong number.

If a non-local reading of *wrong* was available in (66b), the sentence should be true or, at worst, be as off-topic as its definite counterpart. However, the sentence is false, which indicates that a non-local reading is not available.

This shows that an apparent indefinite use of non-local *wrong* as in (64) should be considered a local use in which the salient property *P* is provided by the information structure of the sentence itself.

I have argued in this section that local attributive *wrong* is not restricted to a particular determiner, whereas non-local *wrong* is, but by virtue of creating semantically unique nouns. Nonetheless, the two forms of attributive *wrong* share a number of semantic contributions, such as negation and the modal operator. However, a detailed study of other adjectives with local and non-local readings would be required to determine if we

can relate one reading to the other in a general and systematic way.

7 Conclusion

In this paper, I refined the semantics of non-local *wrong* in several respects. I identified the so-far unnoticed *Bluebeard*-reading, looked at the potential of *wrong* noun phrases as antecedents in discourse, and considered the types of necessity modality associated with *wrong*. I provided additional support for the claim that non-local *wrong* is restricted to definite noun phrases and derived this by specifying the type of definiteness as semantic uniqueness. This allowed me to explain the obligatoriness of the definite article in English and its absence in Papiamentu.

The formalization of my analysis in *Lexical Resource Semantics* has the advantages that a single lexical constraint for non-local *wrong* can capture both readings. Furthermore, a surface-oriented syntactic analysis is possible and the semantic combination of the head noun with non-local *wrong* is subject to exactly the same constraints as the combination of local, intersective adjectives.

Finally, I showed that local *wrong* differs from non-local *wrong* in looking for a contextually salient rather than for a clause-mate predicate in the scope of the necessity operator. As local *wrong* is not restricted to any particular determiner, there is a potential ambiguity when a definite determiner is used. However, I argued that, contrary to recent claims in the literature, only local *wrong* can combine with an indefinite article.

Acknowledgments I thank the reviewers and the audience of CSSP 2019 for comments and suggestions. I am grateful to Cécile Meier and Ahmad Al-Bitar for discussion and for pointing me to Schwarz (2020). All errors are mine.

References

Abbott, Barbara. 2001. Definiteness and identification in English. In Enikő T. Németh (ed.), *Pragmatics in 2000: Selected papers from the 7th International Pragmatics Conference*, vol. 2. 1–15.

Am-David, Assif. 2014. *A semantic typology of definiteness. The crosslinguistic distribution of definiteness markers.* Saarbrücken: Lambert.

Bos, Johan. 1996. Predicate logic unplugged. In Paul Dekker & Martin Stokhof (eds.), *Proceedings of the 10th Amsterdam Colloquium*. 133–143. ILLC/Department of Philosophy, University of Amsterdam.

Davies, Mark. 2008–. The Corpus of Contemporary American English (COCA). https://www.english-corpora.org/coca/.

Egg, Markus. 2010. Semantic underspecification. *Language and Linguistics Compass* 4(3). 166–181.

Haïk, Isabelle. 1985. *The syntax of operators.* Cambridge, MA: MIT dissertation.

Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar.* Malden, MA: Blackwell.

Kester, Ellen-Petra & Cristina Schmitt. 2007. Papiamentu and Brasilian Portuguese: A comparative study of bare nominals. In Marlyse Baptista & Jacqueline Guéron (eds.), *Noun phrases in creole languages. A multi-faceted approach*, vol. 31 (Creole Language Library), 107–143. Amsterdam, Philadelphia: Benjamins.

Kratzer, Angelika. 1977. What 'must' and 'can' must and can mean. *Linguistics and Philosophy* 1. 337–355.

Kratzer, Angelika. 1991. Modality. In Arnim von Stechow & Dieter Wunderlich (eds.), *Semantik: Ein internationales Handbuch zeitgenössischer Forschung*, 639–650. Berlin: Mouton de Gruyter.

Larson, Richard K. 2000. ACD in AP? Paper presented at the 19th West Coast Conference of Formal Linguistics (WCCFL 19), Los Angeles, CA. http://semlab5.sbs.sunysb.edu/~rlarson/wccfl19.pdf.

Löbner, Sebastian. 2011. Concept types and determination. *Journal of Semantics* 28. 279–333. doi:10.1093/jos/ffq022.

Maurer, Philippe. 2013. Papiamentu. In Susanne Maria Michaelis, Philippe Maurer, Martin Haspelmath & Magnus Huber (eds.), *The survey of pidgin and creole languages*, vol. II, Portuguese-based, Spanish-based, and French-based Languages, 163–181. Oxford: Oxford University Press.

Morzycki, Marcin. 2016. Toward a general theory of nonlocal readings of adjectives. In Nadine Bade, Berezovskaya Polina & Anthea Scholler (eds.), *Proceedings of Sinn und Bedeutung 20*. 512–532. https://ojs.ub.uni-konstanz.de/sub/index.php/sub/article/view/278.

Ortmann, Albert. 2014. Definite article asymmetries and concept types: Semantic and pragmatic uniqueness. In Thomas Gammerschlag, Doris Gerland, Rainer Osswald & Wiebke Petersen (eds.), *Frames and concept types: Applications in language and philosophy*, 293–321. Berlin, Heidelberg: Springer.

Penn, Gerald & Frank Richter. 2004. Lexical resource semantics: From theory to implementation. In Stefan Müller (ed.), *Proceedings of the 11th international conference on HPSG 2004, Leuven*. 423–443. Stanford, CA: CSLI Publications. cslipublications.stanford.edu/HPSG/5/penn-richter.pdf.

Pinkal, Manfred. 1999. On semantic underspecification. In Harry Bunt & Reinhard Muskens (eds.), *Computing meaning* (Studies in Linguistics and Philosophy), 33–55. Dordrecht: Kluwer Academic Publishers.

Pollard, Carl & Ivan A. Sag. 1994. *Head-Driven Phrase Structure Grammar*. Chicago: The University of Chicago Press.

Richter, Frank & Manfred Sailer. 2004. Basic concepts of Lexical Resource Semantics. In Arne Beckmann & Norbert Preining (eds.), *ESSLLI 2003 – Course Material I*, vol. 5 (Collegium Logicum), 87–143. Vienna: Kurt Gödel Society Wien.

Richter, Frank & Manfred Sailer. 2006. Modeling typological markedness in semantics: The case of negative concord. In Stefan Müller (ed.), *Proceedings of the 13th International Conference on Head-Driven Phrase Structure Grammar, Varna 2006*. 305–325. Stanford, CA: CSLI Publications. cslipublications.stanford.edu/HPSG/7/richter-sailer.pdf.

Sailer, Manfred. 2004a. Past tense marking in Afrikaans. In Cécille Meier & Matthias Weisgerber (eds.), *Proceedings of Sinn und Bedeutung 8*. 233–248. Konstanz. www.ub.uni-konstanz.de/kops/volltexte/2004/1383.

Sailer, Manfred. 2004b. Propositional relative clauses in German. In Stefan Müller (ed.), *Proceedings of the 11th International Conference on Head-Driven Phrase Structure Grammar, Leuven 2004*. 223–243. Stanford: CSLI Publications. cslipublications.stanford.edu/HPSG/5/sailer.pdf.

Sailer, Manfred & Assif Am-David. 2016. Definite meaning and definite marking. In Doug Arnold, Miriam Butt, Berthold Crysmann, Tracy Holloway King & Stefan Müller (eds.), *Proceedings of the Joint 2016 Conference on Head-driven Phrase Structure Grammar and Lexical Functional Grammar, Polish Academy of Sciences, Warsaw, Poland*. 641–661. Stanford, CA: CSLI Publications. http://cslipublications.stanford.edu/HPSG/2016/headlex2016-sailer-am-david.pdf.

Schwarz, Bernhard. 2006. Attributive *wrong*. In Donald Baumer, David Montero & Michael Scanlon (eds.), *Proceedings of the 25th West Coast Conference on Formal Linguistics*. 362–370. http://www.lingref.com/cpp/wccfl/25/paper1469.pdf.

Schwarz, Bernhard. 2020. Non-local adjectival modification: *the wrong number*. In Daniel Gutzmann, Lisa Matthewson, Cécile Meier, Hotze Rullmann & Ede Zimmermann (eds.), *The Wiley Blackwell companion to semantics*, 1–27. Wiley. doi:10.1002/9781118788516.sem114.

Schwarz, Florian. 2009. *Two types of definites in natural language*. Amherst, MA: University of Massachusetts dissertation.

Distributional profiling and the semantics of modifier classes

Martin Schäfer

Abstract The versatility of modifiers makes their analysis a challenge. Their interpretation depends on their own lexical semantics, their syntactic position, and the semantics of the material they combine with. This paper explores what distributional profiling reveals about four modifiers that are traditionally taken to correspond to different semantic classes. The profiles of four adjectives are presented and the results are qualitatively interpreted. While the results show that assumptions made in the literature are partly reflected in the distributional data, they also show many patterns that require further explanation.

Keywords modification · distributional semantics · English · adjective · adverb · adverbial

M. Schäfer, Eberhard Karls Universität Tübingen, https://www.martinschaefer.info/

In Christopher Pinon & Laurent Roussarie (eds.), Empirical Issues in Syntax and Semantics 13, 139–166. Paris: CSSP. http://www.cssp.cnrs.fr/eiss13/

1 Introduction

Pairs like English *quick/quickly*, *slow/slowly*, *wise/wisely*, and *lucky/luckily* can occur in a variety of positions. This syntactic versatility alone makes fully understanding their behavior a challenge. On top of this, the lexical semantics interacts with their combinatorial potential in intricate ways. Even within semantic classes there is considerable variation in behavior, raising the question of the extent to which investigations of individual items in individual constructions yield usable generalizations. This paper explores the possibility of using distributional semantics across the usages of these items in order to come to interpretable patterns of the behavior of different adjective/adverb pairs.

The pairs were selected because they fulfill the following syntactic and semantic criteria: in their base form, they occur in attributive as well as predicative position. In predicative position, they allow the combination with *to*-INF(initival). In their *-ly* form, they occur in sentence-initial, pre-

verbal and postverbal position, and exhibit a number of different, position-dependent, readings. Further, the occurrence in the *to*-INF construction has been described as expressing the same meaning as their occurrence as *-ly* adverbs.

This paper focuses on these patterns, three adjectival and three adverbial patterns, for each of the pairs. They are illustrated for *wise/ly* in (1) and (2), with examples from the British National Corpus (BNC).

(1) a. attributive adjective [attrib]
"A most wise precaution," Karl said. [A7A 3043]
b. *to*-infinitival [INF]
Perhaps Mrs Nicholson had been wise to leave. [AT4 892]
c. predicative [pred]
If that is so, he is wise. [AKY 830]

(2) a. sentence-initial adverb [advSI]
Wisely, Bright has included biographical entries of dead linguists only. [J7K 33]
b. preverbal adverb [advV]
The CO wisely decided not to notice this particular instance of it.
[ACE 2163]
c. postverbal adverb [Vadv]
We help you choose wisely.
[A65 1983]

Note that they are all treated as mutually exclusive, e.g. *predicative* translates to 'predicative but not followed by *to*-INF'.

The paper is structured as follows: §2 presents a short overview of work on modifier classes as relevant to this investigation. §3 gives an overview of the occurrence frequencies for the four pairs in the respective pattern in the ukWaC, a 2 billion word corpus of English. §4 presents the distributional analysis. The results are discussed in §5, and §6 concludes.

2 Lexeme type, modifier classes, and available readings

One of the main motivations for the exploration of distributional profiles for these four adjectives is the difficulty of getting to grips with them with

other approaches. In particular, it is difficult to reliably establish classes without appealing to intuition, and even if one has established plausible classes, it is not easy to show that they are also linguistically and/or cognitively relevant.

At the level of lexical semantics, the four adjectives, *quick*, *slow*, *wise*, and *lucky*, have been discussed to various degrees. A predicate like *quick* is often seen as an event predicate, cf. Pustejovsky (1995) for *fast* and Bücking & Maienborn (2019) for several of its German translation equivalents. For these authors, the main question is how such an event predicate can successfully be combined with non-event predicates, e.g. artefacts (*fast car*, etc.). A similar judgement should hold for its antonym *slow*. Geuder (2002: 10) categorizes *slow* as an external property connected to movement/change, and thus expected to be "primarily a property of dynamic entities – i.e. events, not individuals".

Geuder (2002) also discusses *wise* and *lucky*. Similar to his view (p. 10) on *intelligent* as connected to psychological conditions, he says that *wise* is a disposition of an individual. The disposition describes a capability (p. 113). *Wise* belongs to the group of agentive adjectives (p. 113), contrasting with evaluative adjectives, a class which contains *lucky*. On his analysis, neither are predicates of events. But there is a link to events for agentive adjectives (in his terms, they make covert reference to an event). If Geuder is right, what are the linguistic reflexes of these different adjective classes?

As evidenced by the availability of the four pairs in all six patterns under consideration (three for the adjectives and three for the related adverbs), there are no striking restrictions on their overall combinatorial possibilities (cf. attributive-only or never-attributive adjectives like *main* or *alone*). One example for a combinatorial constraint can be found in the standard predicative usage: only *wise* allows for an *of*-phrase indicating the agent (cf. Oshima 2009), as in (3).

(3) "That's wise of you, miss." [FR6 1043]

Of-phrases often occur together with *it*-extraposition, as in (4).

(4) . . . , it would surely be wise of Althusser to show how he proposes to do it. [CMN 628]

Geuder (2002: 112), who also discusses *of*-phrases in distinguishing between agentive and evaluative adjectives, further points to a corresponding inability of agentive adjectives to occur with *for*-phrases, as in (5).

(5) That was clever of John/??for John. (= Geuder's (19))

However, at least for *wise*, the situation is not so clear-cut, as seen in examples like (6).

(6) Under those circumstances, I ask him whether it would be wise for the House to proceed with the Bill tonight. [HHX 17055]

In the BNC, the *wise for*-construction, with 30 occurrences, is distinctly more frequent than *wise of*, which occurs only ten times.

In their attributive usages, there is on the surface little difference between the four adjectives. For example, all four adjectives occur with heads referring to more than two different ontological types, all including at least events and physical entities, as in (7).

(7) a. quick decision
b. quick antidote
c. slow progress
d. slow boat
e. lucky draw
f. lucky winner
g. wise counselling
h. wise dragon

Showing that there is a quantitative pattern behind this requires annotation of the ontological types of the heads. In Schäfer (2020a), I show that the majority of heads in the top 120 collocations for both *quick* and *slow* fall in the event category, which is in line with the assumption that these are event predicates.

Larson (1998: 18) and Bücking & Maienborn (2019: 35–36) remark that *quick* and *slow* are more restricted in the available readings in predicative position than in attributive position. Differences between *wise* and the other three adjectives should follow from the expectation that dispositions are better suited for attributive position than predicative position (cf.

Cruse 2004: 301), although I don't know any study showing this explicitly. Geuder (2002: 139–147) links fine-grained differences in the properties of clausal complements of his agentive and evaluative classes to his analysis of them as non-event predicates, and of the covert event reference of the former.

For the adverbial usages of the adjectives, the four adjectives show three patterns in alignment with their different lexical semantics. Of the four, *wisely* is the only one that participates in the standard high-low pattern, where a sentential reading contrasts with a manner reading (e.g. Ernst 2002). While the immediately preverbal position allows both readings, the sentence-initial position is linked to the high reading, and postverbal position to the low reading. In its high reading, *wisely* is a subject or rather agent-oriented sentence adverbial (cf. Maienborn & Schäfer 2011), and in its low reading, it is a manner adverbial. The semantic difference to manner usages is, among other things, revealed through paraphrases:

(8) a. Wisely, Bright has included biographical entries of dead linguists only. [J7K 33]
b. It was wise of Bright to have included biographical entries of dead linguists only.

In contrast, if manner usages are available, paraphrases like "in an ADJ manner/the way in which . . . BE ADJ" are more apt:

(9) a. The old lady nodded wisely: "I thought so . . . scientists would have tried it out on rats first." [A57 17]
b. The old lady nodded in a wise manner/The way in which the old lady nodded was wise.

Quickly and *slowly*, aspect-manner adverbs in the terminology of Ernst (2002), also occur with different readings depending on position, as in (10).

(10) a. Lynn quickly raced down the hallway. (= Ernst's (2.149a))
b. Lynn raced down the hallway quickly. (= Ernst's (2.149b))

While the clarity of the readings depends on the specific verb, one can distinguish between inceptive, holistic, and true rate usages of *quickly*,

with inceptive and holistic readings aligned with high positions, and the true rate reading with the low position. An inceptive usage is one where *quickly* indicates that the time up until an event is short, a holistic usage one where the whole event referred to by the verb phrase took only a short amount of time, and a true rate usage one where the rate of an internal movement inherent to the event is targeted. While they are not always independent of each other (cf. §5.3), an inceptive reading of (10a) is one where *quickly* indicates that the racing-down-the-hallway event takes place shorty after some other, contextually supplied event. The holistic reading is one where the racing-down-the-hallway event itself takes only a short time. Finally, the true rate reading, associated with (10b), indicates that the running itself consists of quick movements, i.e. short subevents. Several terms have been used to describe these differences. Ernst (2002) refers to the inceptive readings as clausal readings of aspect-manner adverbs and takes them to be special cases of a manner reading. Cinque (1999) discusses the preverbal and postverbal occurrences of *quickly* in terms of two different aspectual projections, celerative aspect I and celerative aspect II. While the former involves quantification over an event, the latter involves quantification over a process. For Schäfer (2013) it's a matter of event-related vs. verb-related modification. The latter assumes that the modifier does not directly predicate over the event referred to by the verbal predicate but is connected to the event via some appropriate relation, as in the proposals regarding speed and manner in Dik (1975) and Piñón (2008), respectively.

Luckily, a sentential adverb from the evaluational subclass of speech-act adverbials (cf. Maienborn & Schäfer 2011), can be paraphrased as in (11b).

(11) a. Luckily the flies had gone by now. [A0N 2400]
b. It was lucky that the flies had gone by now.

According to Ernst (2002: 78), *luckily* is a pure evaluative and does not come with a manner reading (though it can occur as a verb modifier for some verbs, with a resultative-like interpretation, as in his (2.132a) *The performance turned out pretty luckily, considering the troubles we'd had*).

While these paraphrases already show that the same content can be

expressed by using either the adjective or the adverb of each pair, both *wisely* and *luckily* have been explicitly linked to the *to*-INF pattern. Oshima (2009: 364) points to the following two sentences with *wise/ly* as being "roughly synonymous" (an observation also made in earlier works):

(12) a. Wisely, John left early. (= Oshima's (1a))
b. John was wise to leave early. (= Oshima's (1b))

The corresponding two sentences with *lucky/ily* show the same pattern:

(13) a. Luckily, John passed the exam. (= Oshima's (6a))
b. John was lucky to pass the exam. (= Oshima's (6b))

As shown in Schäfer (2020a), *to*-INF is equally frequent with *quick* and *slow*, and they also seem very close to their adverbial counterparts, as in (14).

(14) a. Therefore, they are slow to respond to market changes as reflected by movements in relative prices. [HXL 133]
b. Acne responds slowly and drugs need time to work. [CDR 1954]

For all four cases, the additional challenge is how to further tackle the differences within the spectrum of "rough synonymy". Oshima (2009: 372–373), building on Wilkinson (1970) and Barker (2002), sees an assertion/presupposition reversal between the adverbial and the *to*-INF for both *wise* and *lucky*: e.g. that John left early is asserted in (15a) but presupposed in (15b).

(15) a. Wisely, John left early.
b. John was wise to leave early.

Karttunen (2013) argues that *lucky-to*-INF does not presuppose its complement, but is two-way implicative, that is, it yields a positive entailment in positive contexts, but a negative entailment in negative contexts. An example for the latter is (16), where the entailment is negative ("I did not get a table on this trip") in the negative context provided by ordinary negation.

(16) Anyway, I was not lucky to get a table on this trip. Maybe next time. (= Karttunen's (3a))

In Schäfer (2020b), I point out a further difference between *lucky-to*-INF and *luckily*: whereas for the standard evaluative usage of *luckily* there is no restriction on the recipient of the luck (lucky for who?), *lucky-to*-INF is always restricted to a subject-oriented interpretation.

For *quick/slow-to*-INF and the adverbial usages, I (in Schäfer 2020a and Schäfer 2020b) try to establish a reliable difference. For both adjectives, there are minority usages with non-ordinary subjects patterning with the *tough*-construction, as in (17).

(17) The following recipes are quick to prepare and very low in calories. [CDR 220]

In the majority pattern, where *quick/slow-to*-INF take ordinary sujects, as in (18), *quick/slow-to*-INF share with *wise* the alignment with the high adverbial readings, that is, for these two, a true rate reading is excluded.

(18) a. And he is quick to point out that it was a joint decision to make a serious bid. [G39 1207]
b. But foreign governments have been slow to respond with aid. [B7N 203]

Other than that, there are only tendencies, with *quick-to*-INF on the whole showing a more consistent pattern: There is more overlap in the semantic classes of verbs and a preference for inceptive readings for its top verbal collocates. The subjects typically refer to humans or institutions. With *quick to point out*, it also has a very dominant most frequent member whose inceptive reading might serve as an analogical model for the other verbs occuring in this construction.

Note that the only distributional difference discussed in all of the above that does not require careful semantic annotation is the contrast involving the *of*-phrases in the simple predicative pattern. Little is known about the distribution of these four pairs across the patterns in actual corpora, and one main goal of this paper is to explore distributional semantics as a means to close in on the specifics of the different items, ideally being able to link the results of the distributional analysis to ideas discussed in the theoretical literature. As a side effect, this might also shed some light on an unresolved issue regarding the relationship between the base and the

-ly forms: is it derivation or inflection? Bauer et al. (2013: 536) write "that the evidence is inconclusive".

3 The target patterns in the ukWaC

The ukWaC is a 2 billion word corpus of English (Baroni et al. 2009). It is web-derived from only the .uk domain, therefore likely to be more representative of British English than any other variety. The version I used was part-of-speech-tagged and lemmatized with TreeTagger. Each adjective and the corresponding *-ly* form are treated as separate lemmata. All eight lemmata are high-frequency items, as seen in Table 1, and their distribution across the target patterns is shown in Table 2.

pair	base form	*-ly* form	total
quick/quickly	135705	187725	323430
slow/slowly	84391	59272	143663
wise/wisely	37950	6730	44680
lucky/luckily	51252	13911	65163

Table 1 Raw frequencies adjective vs. adverb overview

	quick	slow	wise	lucky
attrib	72607	33791	11551	12309
INF	8035	4733	5280	6034
pred	12307	15884	4563	20220
advSI	1433	2475	236	8950
advV	64149	22600	1848	513
Vadv	36178	15262	2405	64

Table 2 Distribution across target patterns

Within the adjective forms of *quick*, *slow*, and *wise*, the attributive usage is the most frequent usage. For *lucky*, the attributive usage comes second after the standard predicative usage. For *wise*, the *to*-INF pattern is actually slightly more frequent than the standard predicative usage. Closer inspection reveals that the *wise-to*-INF pattern is special in that there is a high proportion of instances of subject extraposition in the data, as illustrated in (19).

(19) a. So, if you have one then it is wise to send it so we can add it to the site.
b. But it is wise to be on your guard against these abuses.

Excluding all instances of *it* followed by any form of *be* from the *to*-INF pattern halves the instances for *wise*, reduces the instances of *quick* by 20% and leaves *slow* and *lucky* by and large unchanged, as seen in Table 3.

	quick	slow	wise	lucky
INF	8035	4733	5280	6034
INF (excluding *it*)	6308	4556	2650	6004

Table 3 The ADJ-*to*-INF pattern including and excluding *it* in subject position

While the exact reason for the preponderance of this pattern for *wise* and the effect of excluding this type on the distributional analysis must await further investigation, it is clear that the corresponding sentences are not paraphrasable by sentences with sentence-initial *wisely*. I therefore excluded this subpattern from further analysis.

For the adverb forms, it is noticable that both *quickly* and *slowly* most often occur preverbally, followed by their postverbal usage. The sentence-initial usage is the least frequent. In contrast, *wisely* occurs more often postverbally than preverbally. *Luckily* is special in being the only item occurring most frequently in sentence-initial position, and rarely preverbally and even more rarely postverbally. This behavior of *luckily* is in line with the observation in the theoretical literature that it does not allow a low reading, which is syntactically associated with preverbal and postverbal position.

4 Distributional analysis

The distributional analysis compares the similarities of the four adjective/adverb pairs across three adjectival and three adverbial patterns. The three adjectival patterns are (a) the adjective in attributive position, (b) the adjective in predicative position followed by *to*-INF (excluding the subpattern in which the auxiliary is immediately preceded by *it*, following the discussion above), and (c) the adjective in predicative position not followed by *to*-INF. The three adverbial patterns are (a) the *-ly* form in

sentence-initial position, (b) the *-ly* form immediately preceding a main verb, and (c) the *-ly* form in postverbal position.

4.1 Preliminaries

To compare the different usages of each adjective, I used distributional semantics. The main idea behind this approach is to represent words exclusively via their distribution (for a comprehensive introduction and overview, see Sahlgren 2006). There are many different ways this can be done, and I proceeded as follows:

4.1.1 Initial steps

1. I first collected cooccurrence counts for each adjective, distinguishing between the three adjectival and the three adverbial patterns in the tagged ukWaC corpus.
2. The cooccurrence counts were collected for the top 10,000 content words (nouns, verbs, adjectives, and adverbs).
3. Only cooccurrences of at least 5 were used in the further calculations.

4.1.2 Further parameter setting and validation

Lapesa & Evert (2014: 542) point out that the three parameters score, transformation, and distance metric consistently play a crucial role in the performance of distributional semantic models, while the parameter of window size is influential but more task dependent. To find a useful setup for the task at hand, I first compared the performance of different settings against the human ratings on the adjective subset of the SimLex-999 dataset described in Hill et al. (2014). This dataset provides two human judgements of interest: a similarity score and and an association score. For the similarity score, raters where instructed to rate synonyms and near-synonyms high, and to not confuse similarity with relatedness. For example, the pair *glasses/spectacles* was given as a reference for a pair with very similar meanings, whereas pairs like *car/tyre*, *car/crash*, and *car/motorway* were used to exemplify the difference between similarity and semantic relatedness. The ratings were given by setting a slider on an integer scale ranging from 0 to 6 with the low values indicating "less similar" and the high values indicating "more similar". In the dataset itself, these values are linearly mapped from the [0,6] range to the [0,10] range. The association

score gives the strength of free association from one word in a pair to its partner. Values are taken from the University of South Florida Free Association Dataset (Nelson et al. 1998). In particular, they use the cue-to-target strength, arrived at by dividing the number of participants producing the target by the number of participants seeing the cue. In the SimLex dataset, this value is multiplied by the factor 10.

While I only used cosine similarity as a distance metric, I used a Python script to vary the scoring and transformation, exploring pointwise mutual information and log-likelihood, both with and without logarithm, and the window size, considering complete sentences, and ranges of two to four words to the left and to the right of the target word. The best performing settings for the similarity scores was a window of three words to the left and right of the target and pointwise mutual information without logarithm. The best performing setting for the association scores was a window of two words, and pointwise mutual information (pmi). Both are shown below, with the best correlation score for each of the two measures in boldface, and the best performing settings for the sentence window at the bottom.

		similarity		association	
window	score and trans	cor	p value	cor	p value
3	pmi w/out log	**0.518**	5.6e-09	0.219	0.021
2	pmi	0.452	6.5e-07	**0.359**	1.2e-04
sentence	pmi w/out log	0.331	3.9e-04	0.304	0.001

Table 4 Best performing settings overall and best performing setting with a sentence window

As Table 4 shows, both narrow window versions clearly outperform the sentence window on the similarity task, and a two word window version also outperforms the sentence window version on the association task. So far, only base forms of different adjectives were compared. For many adjectives, these forms share a preference for attributive position. In contrast, the adjective and the adverb of each pair typically never occur in the same position, and I further distinguish three distinct syntactic environments for each form. To further explore which setting to use for the comparison of these pairs across the syntactic environments, I calculated

the similarity score between the adjective and the adverb of each of the four pairs using the three settings. This reveals a stark difference between the window sizes, as shown in Table 5.

pair	sentence	three words	two words/pmi
quick/quickly	0.63	0.13	0.24
slow/slowly	0.63	0.16	0.29
wise/wisely	0.28	0.04	0.09
lucky/luckily	0.46	0.08	0.20

Table 5 Cosine similarities between adjective and adverb of each pair, using the best settings per window size

A crucial factor responsible for these differences is the apparent inability of the narrow word windows to meaningfully represent the similarity between the attributive usage and the usages of the *-ly* forms. For example, the similarities between attributive *quick* and all usages of *quickly* are only at 0.07 and 0.06 for the two word and three word windows in the above settings, and not higher than 0.04 for the three adverbial positions considered here individually. In contrast, the similarity between attributive *quick* and all *-ly* forms is 0.52 when using the sentential window in its best performing settings.

One factor behind this becomes clear when considering textbook equivalences of adjective and adverb usages illustrated in (20).

(20) a. The quick runner . . .
b. The man runs quickly.

While the head noun *runner* in (20a) is an important and straightforward cue for distributional systems that compare different adjectives, it is not as straightforwardly helpful in establishing a similarity across the two sentences, because the system is blind to the relationship between the verbal base *run* and its nominalization *runner*. The larger contexts used by the sentence window seems to be able to circumvent this problem.

Since the sentence-based window is better suited to comparisons across constructions and forms, it is used when comparing pairs across usages. When forms are compared within a pattern, I will use the setting perform-

ing best for the similarity task, that is, a 3-word window with pmi without log. Note that this method does not reduce the effects of absolute frequency in the corpus completely. Given the considerable differences in absolute frequencies between the six patterns and the four pairs, this factor might well play a role. However, of all 24 pattern/form combinations, only postverbal *luckily* is a clear low frequency outlier.

4.2 Results

The resulting similarities are shown for each of the four adjectives in Tables 6–9. A cosine similarity of 1 indicates perfect similarity (the vectors point in the same direction). The closer the value gets to 1, the more similar two vectors are. A cosine of 0, corresponding to a 90 degree angle, indicates unrelated scores. Negative cosine values are not possible in the two setups selected in §4.1.2 , as both use pmi without log on count data which cannot result in negative values. NA in the last column of Table 9 results from the rarity of postverbal *luckily*; the corresponding similarities could not be meaningfully calculated.

pattern	INF	pred	advSI	advV	Vadv
attrib	0.20	0.33	0.09	0.49	0.41
INF		0.16	0.10	0.33	0.26
pred			0.09	0.33	0.31
advSI				0.16	0.12
advV					0.68

Table 6 Cosine similarities between *quick*-usages.

pattern	INF	pred	advSI	advV	Vadv
attrib	0.20	0.57	0.14	0.49	0.48
INF		0.28	0.08	0.19	0.19
pred			0.12	0.34	0.36
advSI				0.45	0.29
advV					0.73

Table 7 Cosine similarities between *slow*-usages.

Looking at the patterns across the four items, the tables show the fol-

pattern	INF	pred	advSI	advV	Vadv
attrib	0.17	0.51	0.02	0.17	0.13
INF		0.20	0.09	0.17	0.13
pred			0.01	0.16	0.13
advSI				0.13	0.03
advV					0.17

Table 8 Cosine similarities between *wise*-usages.

pattern	INF	pred	advSI	advV	Vadv
attrib	0.21	0.33	0.19	0.06	NA
INF		0.41	0.34	0.22	NA
pred			0.40	0.14	NA
advSI				0.18	NA
advV					NA

Table 9 Cosine similarities between *lucky*-usages.

lowing:

1. The attributive usage across all four adjectives shows little similarity to *to*-INF, with all similarities in the narrow band between 0.17 and 0.21. The attributive usage is more similar to the standard predicative usages, with a clear difference between *slow* and *wise* (0.57/0.51), on the one hand, and *quick* and *lucky*, on the other hand (0.33/0.33). The highest similarity to the sentence-initial adverbial pattern is only 0.19, for *lucky*. The similarities to the pre- and postverbal adverbs show a clear split between relatively high values for *quick* and *slow*, on the one hand, and the other two adverbs, which show little or no similarity, on the other hand.
2. The *to*-INF pattern varies a lot in its similarity to the other usages across the four adjectives. With the exception of *lucky*, it is less similar to the predicative pattern than the attributive is to the predicative pattern. *Lucky* also stands out when comparing the similarity between the *to*-INF pattern and the sentence-initial adverb: the cosine value is 0.34, as opposed to no or hardly any similarity for the other three items. The similarity between *to*-INF and the preverbal adverbs is at the same level for *slow*/*wise*, and *lucky*, with a higher

value for *quick* (0.33). *Quick* has also the highest similarity between INF and the postverbal adverb (0.26), with only 0.19 and 0.13 for *slow* and *wise*, respectively.

3. The similarity between the predicative usage and the sentence-initial adverbs is relatively high for *lucky* (0.41), low for *quick*/*slow*, and non-existing for *wise*. It is moderately high for preverbal and postverbal *quick* and *slow*, and low for *wise* and preverbal *lucky*.
4. The sentence-initial pattern is always more similar to the preverbal usages than the postverbal ones, with moderate similarity to the preverbal and postverbal adverb for *slow*, and lower similarities with the preverbal adverb for *quick*/*wise*/*lucky*, and little or not at all similar for postverbal *quickly* and *wisely*.
5. The preverbal and postverbal adverbs are highly similar for *quickly*/*slowly* (0.68/0.73), showing little similarity (0.17) for *wise*. As mentioned above, for *lucky* this contrast does not apply.

5 Discussion

This section cannot meaningfully discuss all the nuances of the distributional data across the four adjective/adverb pairs. Instead, I will focus on contrasts that are of special interest in view of the discussions and classifications in the previous literature. Therefore, I will discuss the following points more closely:

1. Event and non-event predicates
2. The *to*-INF pattern and the adverbial usages
3. The preverbal and postverbal adverb positions
4. The relationship between forms with *-ly* and without *-ly*

5.1 Event and non-event predicates

In the literature, *quick* and *slow* are held to be event predicates, as opposed to *wise* and *lucky*. *Wise* is claimed to make covert reference to events. Of the adverbial positions, both the preverbal and the postverbal position allow low readings, which directly relate to events. The preverbal position also allows high readings, which, for *quick* and *slow* at least, are both linked to the event encoded by the verbal predicate (see §5.3 for more discussion). Since *quick* and *slow* are event predicates, they can simply be used in the same way there as in attributive and predicative position. All

four usages should be similar to each other. In contrast, for *wise*, the low adverbial reading should be quite different from its attributive and predicative usage, because now it must be brought in a direct relation with an event. Table 10 shows the similarities for the relevant patterns.

	attrib		predicative	
adjective	advV	Vadv	advV	Vadv
quick	0.49	0.41	0.33	0.31
slow	0.49	0.48	0.34	0.36
wise	0.17	0.13	0.16	0.13
lucky	0.06	NA	0.14	NA

Table 10 Cosine similarities between the attributive and the predicative patterns and the preverbal and postverbal patterns for all four pairs

The observed similarities are by and large in line with the expectations: both *quick* and *slow* have relatively high similarity values for both preverbal and postverbal adverbial usages, while for *wise*, the values are markedly lower. While it is unclear how many of the preverbal usages correspond to the high reading, the similarity value here is still not very high for *wise*, showing that it is more distinct across its usages overall.

Note that the values for *lucky* are also low. I do not have an explanation for this. It might have to do with the preverbal usage of *luckily* being relatively rare and perhaps restricted to more idiosyncratic combinations.

Another quantitative effect of *wise* behaving more like an event predicate in the two adverbial usages can be seen when looking at the similarities between *wise* and *quick*/*slow* across the respective usages. Table 11 shows the similarities between the relevant four different usages of *wise* to the corresponding usages of *quick* and *slow* (because I am now comparing the same usage, the reported similarities are from the best performing similarity setting, that is, using a 3-word window and pmi without log).

While there is no or almost no similarity across the adjective usages, similarity slightly increases in the preverbal adverbial usage. There is a marked increase in similarity for the postverbal usage, that is, the position restricted to low readings.

	across-item similarity: *wise*			
adjective	attributive	predicative	preverbal	postverbal
quick	0.00	0.01	0.02	0.05
slow	0.00	0.00	0.01	0.12

Table 11 Cosine similarities between four *wise*-patterns and the corresponding *quick/slow* patterns.

5.2 The *to*-INF pattern and the adverbial usages

It is unclear what kind of similarities to expect for the *to*-INF pattern in comparison to the attributive and predicative usage. For *quick*, I argue in Schäfer (2020a) that its occurrence in the *to*-INF pattern is maximally different from its usage in the attributive position. Among other things, it is restricted to intentionally acting subjects, typically, humans or institutions. This would lead one to expect a low similarity between these two usages, and a similar point could be made for *slow*. More interesting are the expectations for the similarities to the adverbial readings, especially those described as roughly synonymous. Since the *to*-INF construction is aligned with the higher readings, I expect it to be more similar to the sentence-initial and preverbal occurrences of *quick*, *slow*, and *wise* than to the postverbal one. For *luckily*, there is only one reading, but the restriction to subject-oriented interpretations might make it less similar than the correspondences to the high readings for the other adverbs. I have no clear idea what influence the presupposition/assertion reversal for *wise* should have; see below for more discussion.

pattern	attrib	pred	advSI	advV	Vadv
quick-to-INF	0.20	0.16	0.10	0.33	0.26
slow-to-INF	0.20	0.28	0.08	0.19	0.19
wise-to-INF	0.17	0.20	0.09	0.17	0.13
lucky-to-INF	0.21	0.41	0.34	0.22	NA

Table 12 Cosine similarities between the *to*-INF pattern and all other patterns for each pair

Among the similarities of the *to*-INF pattern to the other usages, as seen in Table 12, the relation to the attributive usages stands out as being the

most consistent in magnitude. For all four adjectives, it is also clearly lower than the corresponding attrib-pred values (0.33, 0.57, 0.51, 0.33). *Lucky-to*-INF/*luckily* differ from the three corresponding pairings due to their high value for the sentence-initial usage. The values for the similarity to the preverbal usage are relatively consistent for all four adjectives, with *quick* an outlier in being markedly more similar. The three values for the postverbal usage are all different by at least 0.06, and there is a clear drop in similarity for *quick* and *wise*, but not for *slow*.

All this only partially matches the expectations. As expected, *quick-to* and *slow-to* are less similar to their corresponding attributive usages than the attributive usage is to the predicative usage. However, this also holds for *lucky* and *wise*. As for the adverbial usages, the high value for sentence-initial *luckily* is expected, the consistently low values for the other three adverbs are unexpected. The most likely explanation for these low values is that the sentence-initial position is only a very marginal option for these adverbs, used when regular options are exhausted. This is in line with the observation that it is, for all three adverbs, the least frequent usage, less frequent than their preverbal and postverbal usages by at least a factor of 6. In contrast, for *luckily*, the sentence-initial pattern is the dominant pattern, with the preverbal and postverbal usages less frequent by more than a factor of 10 and of 100, respectively. That the similarity for *luckily* is also the highest overall is unexpected giving that the usages are restricted to subject-oriented interpretations (more on this in §5.2.1). The drop in similarity between preverbal and postverbal adverbs observed for *quick* and *wise* is expected. Why does it not obtain for *slow*? It could be related to the observation in Schäfer (2020b) that *quick-to*-INF is internally more coherent, but further research is needed here.

5.2.1 The subject-oriented reading of luckily

The narrowing of the interpretation of *lucky* to an evaluation relative to the subject seems not to be something that is picked up by the distributional analysis. In fact, the sentence-initial similarity value is the highest across all four adjectives. Why might this be? I think that the analysis of *luckily* as a speaker-oriented evaluative is correct, that is, it is the speaker's evaluation of a fact. For who this is lucky can be made explicit by a *for*-phrase, cf. (21).

(21) Luckily for her, she had a clever lawyer at her trial, and was never punished for the murder. [FPU 1866]

Without an explicit *for*-phrase, the tendency seems to be to assume that it is lucky for the speaker, as in (22), where the context does indeed confirm that it was an instance of luck for the speaker, and what it meant for the referent of *she* is irrelevant.

(22) Luckily she left the school! [KA1 2120]

My impression is, though, that, when combined with sentences with human subjects, the speaker's evaluation often plausibly coincides with the assessment the speaker would give when taking the viewpoint of the subject, as in a typical example for *luckily* (from Ernst 2009) in (23).

(23) a. Luckily, Aaron did not fall off his bicycle. (= Ernst's (1c))
b. It is lucky that Aaron did not fall off his bicycle. (= Ernst's (2b))

While Ernst gives (23b) as a paraphrase, the corresponding *to*-INF sentence will in most instances also be consistent with the situation described in (23a), as shown in (24).

(24) Aaron was lucky to not fall off his bicycle.

Note that these subject-oriented readings are still different from the standard examples of subject-oriented adverbials in not describing a mental attitude, nor are they agent-oriented in the sense that the agent has control over the action (the notion of agent orientation in Ernst 2002: 55). That is, the readings are best described as still corresponding to speaker-oriented adverbials, with the target of the luck made explicit, similar to the adverbial usage with a *for*-phrase. If this is correct, then the difference is indeed a minor one, and perhaps expected to not impact much on the distributional similarity measure.

5.2.2 The assertion/presupposition reversal for wisely

The correspondence between *wisely* and *wise-to* INF is on average the lowest. One reason could be that the distributional analysis is able to pick up on the assertion/presupposition reversal. However, this is impossible

to ascertain when one looks at other work on entailments in distributional semantics.

Distributional semantics is in principle able to deal with entailments, and there are numerous studies dealing with lexical entailment (cf. Turney & Mohammad 2014 for a comparison of different approaches). Studies on entailment above the word level also exist, but the domain for which entailment in a strict formal semantic sense is explored is typically still very small (cf. Baroni et al. 2012, who report two experiments on adj-noun to noun and quantifier phrase to quantifier phrase entailments). At the sentence level, the distinction between entailment and presupposition seems to be largely irrelevant for the NLP community, where this task falls into the domain of recognizing textual entailment: "In RTE [recognizing textual entailment], the gold standard for entailment is established by common sense, rather than formal logic" (Turney & Mohammad 2014: 2). No differentiation is usually made between the role of linguistic and world knowledge in this. In general, the inferences are not seen as absolute, and a well-known approach states "We say that T entails H if, typically, a human reading T would infer that H is most probably true" (Dagan et al. 2009: iv). That is, there is also no specific distinction between entailment and implicature. However, on this level, all four adjectives pattern together, as opposed to adjectives like *eager* in this construction.

(25) It was a huge budget and the top agencies were eager to get their hands on it. [ADK 667]

This lack of interest in the distinction between entailments, implicatures, and presuppositions seems to be driven mostly by the assumption that this level of detail is irrelevant for applied tasks. I don't know of any study that uses training data distinguishing between presuppositions and implicatures. Either way, I have no bottom line to compare the possible effect of *wise*/*wisely* to, since, in my data, *wisely* is the only adverb that has the standard subject-oriented adverbial reading. Its presuppositional behavior is thus not its only difference to the three other adverbs.

5.3 The preverbal and postverbal adverb positions

The literature assumes, where available, pairs of high and low readings for the adverbs. Of the positions tested, the sentence-initial position is re-

served for high readings, and the postverbal position for low readings. The position immediately preverbal is ambiguous between high and low readings when no auxiliaries are present, whose presence was not controlled for in this study. For *quickly* and *slowly*, this means that postverbally only true rate readings are predicted to occur, with the immediately preverbal position in addition allowing inceptive and holistic readings. For *wisely*, only the manner reading is expected to occur postverbally, while the high reading is restricted to preverbal position. Leaving aside the issue of how many readings in the preverbal position are actual high readings, the two variants are more closely related for *quick* and *slow*. This not only holds on the theoretical level, but also conceptually: the different readings often stand in implicative relationships. This can already be seen when looking at (10) from §2, repeated here for convenience:

(26) a. Lynn quickly raced down the hallway.
b. Lynn raced down the hallway quickly.

For Ernst, (26a) "can be interpreted as saying that Lynn's beginning the action of racing-down-the-hallway occurred quickly after some other event (perhaps a command to go fetch something), while (2.149b)[(26b)] is a description of the speed of her movement" Ernst (2002: 85). I suspect that in real world situations, the inceptive interpretation of (26a) typically is taken to imply the true rate reading of (26b). And given that the hallway is of a limited length, the true rate reading for (26b) implies the holistic interpretation that the whole action only took a short amount of time. In this particular example, since *to race* already is something connected to high speed, the true rate reading might be even more likely to imply the holistic interpretation. This requires more investigation, as it is also clear that these implicative relationships do not obtain for all verb types (for example, *quickly* in combination with stative verbs cannot receive a true rate reading and typically receives an inceptive reading). In contrast, the high and low reading of subject-oriented adverbs like *wisely* are conceptually clearly distinct, and, while not very common, opposites can be used with the respective readings, leading to much-discussed examples like Parson's (1972) *John painstakingly wrote illegibly*. Thus, we can expect higher similarity values for *quickly* and *slowly* between the preverbal and postverbal

usages, and lower ones for *wisely*. This is borne out by the data; in fact, the preverbal and postverbal usages of *quickly* and *slowly* lead to the highest similarity values observed in this study, 0.68 and 0.73, respectively. In contrast, the corresponding value for *wisely* is a mere 0.17.

5.4 The relationship between forms with *-ly* and without *-ly*

The relationship between the adjectival base form and the *-ly* form is often discussed in English linguistics. The issue is whether it should best be seen as a derivational suffix or an inflectional suffix: see Payne et al. (2010) and Giegerich (2012) as representatives of the two positions.

Criteria that can be approached by distributional semantics should concern meaning. Plag (2003: 195–196) argues that two meaning-related aspects have a bearing on this issue: *-ly* does not encode lexical meaning, which would be expected from derivation, and *-ly* is always semantically transparent (the latter with only a few exceptions), which would be expected from inflection.

Bonami & Paperno (2018) discuss a list of five criteria from Stump (1998). These also includes Plag's two points (instead of semantic transparency, Stump speaks of semantic regularity: "inflection is semantically more regular than derivation"). They point out that these criteria are "formulated in terms of high-level morphological notions that are not easy to operationalize" and instead provide an operationalization of the semantic regularity criterion in terms of stability of contrasts, as given in (27) (= Bonami & Paperno's (2)):

(27) **Stability of contrast:** The morphosyntactic and semantic contrasts between pairs of words related by the same *inflectional* relation are more similar to one another than the contrasts between pairs of words related by the same *derivational* relation.

Bonami & Paperno explored this criterion by looking at sets of triplets of <pivot, inflectionally related form, derivationally related form> in French. For example, one such triplet was <verb INF, verb PST.IPFV.3SG, SG *-eur* noun>, e.g. <*baigner* 'to bathe', *baignait* 'bathed', *baigneur* 'bather'>. They found that overall the contrasts between inflectionally related forms were more stable (measured in terms of dispersion around the average vector offset for a specific paradigmatic system).

While their approach was word-form based, the comparison across patterns here used lemmatized forms, and I only looked at four word pairs. However, even these four pairs are interesting because they suggested that a combination of word form and syntactic pattern can yield more insight into this issue. Recall that the overall similarity between the two forms across all usages were 0.63, 0.63, 0.28, and 0.46, respectively, leading to an overall standard deviation of 0.17. Taking the standard deviation as a measure of dispersion, we can now compare the standard deviations across all usages, as seen in Table 13.

pattern	INF	pred	advSI	advV	Vadv
attrib	0.02	0.12	0.07	0.22	0.23
INF		0.11	0.13	0.07	0.11
pred			0.17	0.10	0.16
advSI				0.15	0.13
advV					0.36

Table 13 Standard deviations of the four similarity values for the four pairs in each pairing

As the table shows, the dispersions for the similarities within the adjective usages are all clearly below 0.17, with 0.12 the highest value. Within the adverbial usages, the standard deviation for the similarities between preverbal and postverbal usage is 0.36, the highest standard deviation in the table. The two other values, while higher than any of the adjectival values, fall below the threshold of 0.17. Looking at the values for similarities involving one adjectival and one adverbial usage, we see a wide range, with only two values higher than 0.17. Only four out of nine values fall in the range of the adjective-only values: the standard deviations for the similarities between attributive adjective and sentence intitial adverb, *to*-INF and preverbal adverb, *to*-INF and postverbal adverb, and simple predicative usage and preverbal adverb. Following the general logic of Bonami & Paperno (2018), these values support the assumption that *-ly* is derivational: low dispersion around different usages not involving derivation, higher values between derived forms, and on average higher values for pairs of forms related by derivation.

6 Summary and outlook

This paper presented a pilot study using distributional semantics to compare four adjective/adverb pairs across six different usages. The four pairs were meant to represent a variety of different lexical types that are discussed in the literature to be reflected in their linguistic behavior. The general idea was to explore whether qualitative observations made about these items and their corresponding classes show up in their distributional characteristics. After presenting the distributional analysis of *quick*, *slow*, *wise*, and *lucky* across three adjectival and three adverb patterns, four areas were discussed in more detail:

1. The assumption that *quick* and *slow* are event predicates is in line with the high similarity across their standard adjectival and standard adverbial usages, contrasting with *wise*, which shows consistently lower similarities.
2. The comparison between the *to*-INF pattern and the adverbial patterns only partially reflected expectations. In general, *lucky* behaved differently from the three other adjectives, which might in part be due to its already very different distribution of absolute frequencies. Also, all three other adjectives were expected to be more similar to the preverbal usages, but this was only the case for *quick* and *wise*, and not for *slow*. Further, it remained unclear whether anything in the distributional data was sensitive to the subject-oriented interpretation of *lucky* in the INF construction and the presupposition/assertion reversal observed for the *wise-to*-INF construction.
3. Of the three adverbs that occured frequently enough in both pre- and postverbal position, *quickly* and *slowly* were relatively similar, in contrast to *wisely*. This is in line with the observation that the high and low readings for *wisely* are clearly conceptually distinct. For *quickly* and *slowly*, the respective readings in many cases imply at least one of the other readings.
4. The data presented is in line with the assumption that *-ly* is a derivational affix if one assumes that derivation is linked with less stability of contrasts, here operationalized by higher standard deviations across the similarity values for a given comparison.

Overall, this paper has shown that the combination of distributional se-

mantics of different modifier usages with qualitative analysis is a promising step forward in the analysis of the semantics of adjectives and adverbs and their interrelation. Many findings that remained unresolvable here might be resolvable when more pairs are taken into account. Consideration of a larger number of pairs, and also exploration of further distributional setups, would also allow one to clarify the possible influence of the raw frequencies of the forms across the six patterns on the results reported here.

Acknowledgments Many thanks to the original anonymous reviewers of my abstract for CSSP and to the audience in Paris for their precise and helpful comments. Heartfelt thanks also to one anonymous reviewer and Chris Piñón for their careful reading and comments on the final version of this paper. Work on this paper was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 833 – Project ID 75650358.

References

Barker, Chris. 2002. The dynamics of vagueness. *Linguistics and Philosophy* 25(1). 1–36. doi:10.1023/A:1014346114955.

Baroni, Marco, Raffaella Bernardi, Ngoc-Quynh Do & Chung-chieh Shan. 2012. Entailment above the word level in distributional semantics. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*. 23–32. Avignon, France: Association for Computational Linguistics. `https://www.aclweb.org/anthology/E12-1004`.

Baroni, Marco, Silvia Bernardini, Adriano Ferraresi & Eros Zanchetta. 2009. The wacky wide web: a collection of very large linguistically processed web-crawled corpora. *Language Resources and Evaluation* 43(3). 209–226. doi:10.1007/s10579-009-9081-4.

Bauer, Laurie, Rochelle Lieber & Ingo Plag. 2013. *The Oxford reference guide to English morphology*. Oxford: Oxford University Press.

Bonami, Olivier & Denis Paperno. 2018. Inflection vs. derivation in a distributional vector space. *Lingue e Linguaggio* 17(2). 173–195.

Bücking, Sebastian & Claudia Maienborn. 2019. Coercion by modification. The adaptive capacities of event-sensitive adnominal modifiers. *Semantics and Pragmatics* 12(9). doi:10.3765/sp.12.9.

Cinque, Guglielmo. 1999. *Adverbs and functional heads: A cross-linguistic perspective*. Oxford: Oxford University Press.

Cruse, Alan. 2004. *Meaning in language: An introduction to semantics and pragmatics*. Oxford: Oxford University Press.

Dagan, Ido, Bill Dolan, Bernardo Magnini & Dan Roth. 2009. Recognizing textual entailment: Rational, evaluation and approaches. *Natural Language Engineering* 15(4). i–xvii. doi:10.1017/S1351324909990209.

Dik, Simon C. 1975. The semantic representation of manner adverbials. In A. Kraak (ed.), *Linguistics in the Netherlands 1972–1973*, 96–121. Assen: Van Gorcum.

Ernst, Thomas. 2002. *The syntax of adjuncts*. Cambridge: Cambridge University Press.

Ernst, Thomas. 2009. Speaker-oriented adverbs. *Natural Language and Linguistic Theory* 27. 497–544.

Geuder, Wilhelm. 2002. *Oriented adverbs: Issues in the lexical semantics of event adverbs*. Tübingen: Universität Tübingen dissertation.

Giegerich, Heinz J. 2012. The morphology of *-ly* and the categorial status of 'adverbs' in English. *English Language and Linguistics* 16(3). 341–359. doi:10.1017/S1360674312000147.

Hill, Felix, Roi Reichart & Anna Korhonen. 2014. Simlex-999: Evaluating semantic models with (genuine) similarity estimation.

Karttunen, Lauri. 2013. You will be lucky to break even. In Tracy Holloway King & Valeria de Paivaw (eds.), *From quirky case to representing space: Papers in Honor of Annie Zaenen*, 167–180. Stanford, CA: CSLI Publications.

Lapesa, Gabriella & Stefan Evert. 2014. A large scale evaluation of distributional semantic models: Parameters, interactions and model selection. *Transactions of the Association for Computational Linguistics* 2. 531–545.

Larson, Richard K. 1998. Events and modification in nominals. In D. Strolovitch & A. Lawson (eds.), *Proceedings from Semantics and Linguistic Theory 8*. Ithaca, NY: Cornell University.

Maienborn, Claudia & Martin Schäfer. 2011. Adverbials and adverbs. In *Semantics: An international handbook of natural language meaning*, vol. 2, 1390–1420. Berlin: Mouton de Gruyter.

Nelson, D. L., C. L. McEvoy & T. A. Schreiber. 1998. The University of South Florida word association, rhyme, and word fragment norms. `http://www.usf.edu/FreeAssociation/`.

Oshima, David Y. 2009. Between being wise and acting wise: A hidden conditional in some constructions with propensity adjectives. *Journal of Linguistics* 45(2). 363–393. doi:10.1017/S0022226709005684.

Parsons, Terence. 1972. Some problems concerning the logic of grammatical

modifiers. In Donald Davidson & Gilbert Harman (eds.), *Semantics of natural language* (Synthese Library). 127–141. Dordrecht: D. Reidel.

Payne, John, Rodney Huddleston & Geoffrey K. Pullum. 2010. The distribution and category status of adjectives and adverbs. *Word Structure* 3(1). 31–81.

Piñón, Christopher. 2008. From properties to manners: A historical line of thought about manner adverbs. *Studies van de Belgische Kring voor Linguïstiek – Travaux du Cercle Belge de Linguistique – Papers of the Linguistic Society of Belgium* 3.

Plag, Ingo. 2003. *Word-formation in English.* Cambridge: Cambridge University Press.

Pustejovsky, James. 1995. *The generative lexicon.* The MIT Press.

Sahlgren, Magnus. 2006. *The word-space model: Using distributional analysis to represent syntagmatic and paradigmatic relations between words in high-dimensional vector spaces.* Stockholm: Stockholm University dissertation.

Schäfer, Martin. 2013. *Positions and interpretations: German adverbial adjectives at the syntax-semantics interface.* Berlin: De Gruyter Mouton.

Schäfer, Martin. 2020a. From *quick* to *quick-to*-INFINITIVAL: On what is lexeme specific across paradigmatic and syntagmatic distributions. *English Language and Linguistics* 1–31. doi:10.1017/S1360674320000167.

Schäfer, Martin. 2020b. Modifier semantics and constructional idiosyncrasies in the English ADJ-*to*-INF construction. Submitted to an edited volume on Modifiers and Complex predicates.

Stump, Gregory T. 1998. Inflection. In A. Spencer & A. Zwicky (eds.), *The handbook of morphology*, 13–43. London: Blackwell.

Turney, P. D. & S. M. Mohammad. 2014. Experiments with three approaches to recognizing lexical entailment. *Natural Language Engineering* 21(3). 437–476. doi:10.1017/s1351324913000387.

Wilkinson, Robert W. 1970. Factive complements and action complements. In *Proceedings of the Chicago Linguistic Society. Papers from the Sixth Regional Meeting.* 425–44. Chicago: Chicago Linguistic Society.