

No means No!
Speech Acts in Conflict

Lwenn Bussière-Caraes

The research for/publication of this doctoral thesis received financial assistance from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 758540)

No means No!
Speech Acts in Conflict

ILLC Dissertation Series DS-2022-05



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

For further information about ILLC-publications, please contact

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Science Park 107
1098 XG Amsterdam
phone: +31-20-525 6051
e-mail: illc@uva.nl
homepage: <http://www.illc.uva.nl/>

Copyright © 2022 by Lwenn Bussière-Caraes

Cover design by Nuno M. Santos.
Printed and bound by Ipskamp Printing.

ISBN: 978-94-6421-915-9

No means No!
Speech acts in conflict

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. P.P.C.C. Verbeek
ten overstaan van een door het College voor Promoties ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op vrijdag 25 november 2022, te 16.00 uur

door Leila Natacha Bussière
geboren te Longjumeau

Promotiecommissie

Promotores:

dr. L. Incurvati
prof. dr. A. Betti

Universiteit van Amsterdam
Universiteit van Amsterdam

Overige leden:

prof. dr. M. Mikkola
prof. dr. M.J.B. Stokhof
prof. dr. I. Dickie
prof. dr. A. Tanesini
dr. K.H. Krzyzanowska

Universiteit van Amsterdam
Universiteit van Amsterdam
University of Toronto
Cardiff University
Universiteit van Amsterdam

Faculteit der Geesteswetenschappen

La pensée avant d'être œuvre est trajet. N'aie pas honte de devoir passer par des lieux fâcheux, indignes, apparemment pas faits pour toi. Celui qui pour garder sa « noblesse » les évitera, son savoir aura toujours l'air d'être resté à mi-distance.

HENRI MICHAUX

Contents

Acknowledgments	xi
1 Introduction	1
1.1 The speech acts of disagreement	3
1.1.1 Rejection	3
1.1.2 Adversation	3
1.1.3 Weak assertion	4
1.1.4 Retraction	4
1.2 Theories of conversation	5
1.2.1 Interpreting silence	5
1.2.2 Interpreting communication	5
1.3 Note on chapters material	5
2 Rejection	7
2.1 Introduction	7
2.1.1 Rejection as a distinct activity	7
2.1.2 Norm-based approaches to speech acts	9
2.1.3 Outline	9
2.2 The messiness of rejection	10
2.2.1 Classifying theories of speech acts	10
2.2.2 Effect-based models of conversation	13
2.2.3 Modelling rejection in effect-based models	15
2.2.4 Explaining the messiness of rejection	16
2.3 From assertion to rejection	17
2.3.1 The norms of assertion	18
2.3.2 Normative duality	20
2.4 Norms of Rejection	23
2.4.1 The Truth Norm.	23
2.4.2 The Belief Norm	25

2.4.3	The Knowledge Norm	25
2.4.4	The Belief ⁺ Norms	27
2.5	Discussion	29
2.5.1	Which norm for rejection?	29
2.5.2	Consequences of the Knowledge norm of rejection	31
2.6	Conclusion	32
3	Yes, but No	35
3.1	Empirical profile of the Adversative Marker <i>but</i>	35
3.1.1	<i>But</i> 's projection behaviour	36
3.1.2	The multiple uses of <i>but</i>	41
3.2	Extant accounts for the adversative marker	43
3.2.1	Formal contrast approaches	43
3.2.2	Inferential contrast approaches	45
3.2.3	QUD approach	46
3.2.4	Intermediate conclusions	47
3.3	Challenging data	48
3.3.1	Iterations	48
3.3.2	Negationless corrections	49
3.3.3	Combination with other speech acts	50
3.4	Denial condition as a speech act	51
3.4.1	Expanding the denial condition	51
3.4.2	Denial condition: an indirect rejection	53
3.4.3	Weak Rejection and <i>but</i> -sentences	54
3.4.4	Pragmatic rejection and <i>but</i> -sentences	56
3.5	Conclusion	57
4	<i>Perhaps, might</i> and weak assertion	59
4.1	Weak assertion	59
4.2	Assertion and weak assertion, the theoretical picture	60
4.2.1	The theoretical argument for weak assertion	60
4.2.2	Weak assertion and Common Ground	62
4.2.3	The weak assertion account of <i>perhaps</i> and <i>might</i>	63
4.3	What a weak assertion picture supposes	64
4.3.1	Expectations regarding <i>perhaps</i> embedding behaviour	64
4.3.2	Predictions regarding <i>perhaps</i> and <i>might</i>	66
4.4	Gathering examples	67
4.4.1	Conditional antecedents	69
4.4.2	Attitude reports	72
4.4.3	The case of indirect discourse	73
4.4.4	Think that perhaps: a superhedge	74
4.4.5	Redundancy	75
4.5	Conclusion	77

5	Retraction and Common Ground	81
5.1	The Retraction Dilemma	81
5.2	What Retractions undo	84
5.2.1	Retractions and assertions	84
5.2.2	Retractions as a speech act	85
5.2.3	Typology of undoing	87
5.2.4	The unique effect of retractions	88
5.3	Retraction as a proposal	90
5.3.1	Retractions failures	90
5.3.2	Rejected retractions	92
5.3.3	Retractions and Power Dynamics	94
5.4	Retraction and Common Ground	96
5.5	Problematic Retractions	102
5.5.1	Retractions and subsequent utterances	102
5.5.2	When retractions go wrong	104
5.6	Conclusion	105
6	Silence, Dissent, and Common Ground	107
6.1	Introduction	107
6.2	Silence and Common Ground	108
6.3	Silent Dissent	110
6.4	Silence and default attitudes	112
6.4.1	Dialogue cooperativeness	112
6.4.2	Identifying dialogue situations	114
6.4.3	From cooperativeness to interpretation of silence	116
6.5	Silence and updates	117
6.6	Silencing without assent	119
6.7	Conclusion	121
7	Communicating with colourings	123
7.1	Introduction	123
7.2	Colourings and Conventional Implicatures	124
7.2.1	Colourings	124
7.2.2	Conventional Implicatures	127
7.2.3	Distinguishing colourings and conventional implicatures	129
7.3	Notions of communication	129
7.4	Communicating with colourings	136
7.5	Conclusion	140
8	Conclusion	141
8.1	What we learned	141
8.2	Further Prospects	143

Bibliography	145
List of Figures	155
List of Tables	157
Abstract	159
Samenvatting	163

Acknowledgments

This is the part where I write about feelings, about the people who made me into the researcher and the person I am today. I cannot possibly mention everyone who brought me there. And, for the ones I do mention, I cannot possibly do them justice. Let's try anyway.

First and foremost, I want to thank my supervisors. Luca Incurvati guided me every step of the way, and in many respects, I couldn't have written this thesis without him. Someone (I am not naming names) once joked that the harshest indictment they ever heard from you was to call a book 'a little odd'. That sums up well your constant kindness and curiosity. From the first day I discussed a rough intuition with you to the last thesis meetings, you taught me what it means to be a philosopher – intelligent, kind, detailed, constantly questioning, organised, open-minded. I am far from there, but you remained unwavering in your support and thought-provoking comments. Thank you for trusting me to develop my ideas, always challenging me, and believing in me when I didn't believe in myself. Arianna Betti's pointed questions were exactly what I needed to bring up new ideas and problems to light. Thank you for taking time out of your summer to give me one last push, and for reminding me that done is better than perfect.

I am beyond honoured that researchers whose character and work I admire accepted to be part of my committee. Mari, Martin, Imogen, Alessandra and Karolina, I am grateful that you took the time to read this manuscript, anxious to hear what you thought of it, and more than a little impressed.

Giorgio and Julian, I had a blast working with you during the EXPRESS project. Because of you, I discovered many topics I would have neglected otherwise, and was challenged to grow as a researcher and as a person.

For the past five years, the ILLC was both a place and a community. I want to thank some of the many people who made it so, colleagues and friends. Peter Hawke, Frederik Möllerström, Robert van Rooij, Paul Dekker and Martin Stokhof, I learned so much about teaching by working with you. Special thanks also go to Paul Dekker and Marianne de Heer Kloots for helping my useless-at-Dutch self

with the summary of this dissertation. Caitlin, Dean, Giovanni, Jenny, Raquel, and Marco, it was a privilege to participate in the Diversity Committee with you. Anna, Ece, Jan, and Robert, you made PhD Council meetings feel more fun than thankless. Maria and Katrin, thank you for being an inspiration, and a bit of a role-model.

Anna, you were my friend almost from the beginning. You took time for me when you had little, you listened to me and made it clear there was someone on my side. Thank you for being there, and being you. Arianna, Sirin, you brought back magic during the pandemic. I don't believe in love at first sight, but Alina and Simon made me believe in friendship at first sight. I couldn't imagine better people to be my paranympths. There is no PhD without procrastination, and Levin, Daira and Tobias helped a lot with that, as well as coffee, and occasional food. Patricia, you told me once that you wished to give back some of the mentoring and support you received during your PhD. I don't know if you meant me, but you sure were a positive presence, and the best company for the past few Christmases. Let's brunch soon. I complained a lot about sharing office with *mathematicians*; but I hope you know, Jan, Guillermo, and Johannes, that you made F1.13 the best room in the building.

I think it's Paul who told me that *Amsterdam was not the worst place to study*. I didn't imagine it would become a home, and yet. The friends I made there made me wish to stay. Hráfn, you were a fantastic flatmate. Every expat needs a Dutch cultural ambassador and translator (literally); I can't recommend Effie enough. Wojciech, you helped me stay sane in my last year of PhD, reached out when I neglected you, and let me eat all your food. My loud neighbour, Zach, turned out to be a fast and solid friend. Americans, I guess. Nikki, you brought a whole lot of poetry in my life, and you have the best book recs. Arthur, Sebastian, and Sacheen, you showed me that love takes many shapes. It's wonderful having you in my life.

I left a piece of my heart in Paris. Thank you, Elarnon, Corbeau, Miltøn, a3nm, Tobast, Evarin, Nath, Orphée, Tito, Noémie, for keeping it safe despite the changes. Basile, I am honoured to be your friend, and I am so proud of you. See you soon. Guillain, you reminded me that there is a big beautiful world out there. Thank you for guiding me on mountain paths. Growing up also comes with new friends. Neha, Orphée, Mélodie and Ner, thank you for helping me figure it out.

Nuno, you entered my life by chance, and I hope you never leave. The more I get to know you, the more you amaze me with your smarts and your kindness. I love you.

Amsterdam
August, 2022.

Lwenn Bussière-Caraes

Chapter 1

Introduction

My first ever philosophy class was *not* about a chicken. Our teacher let us sit down, briefly introduced herself, and went to the blackboard. She drew (rather badly) the side view of a chicken, and asked us what it was – *A chicken. The idea of a chicken. A drawing of a chicken. Chalk lines on a board arranged to represent a chicken.* She smirked. *What this is, she said, is half a chicken. And your job, as philosophy students, is to go around this half-chicken to see if there is another half.*

The first half of the chicken. Speakers use language to perform actions (Austin, 1975). With speech acts, speakers bring new information, make promises, give orders, etc. Under the assumption that conversations were a cooperative endeavour where audiences were sympathetic to the speaker's goals (Grice, 1975), we could study assertions, promises, orders, questions, as they relate to speakers intentions. For example, speech acts could be characterised in terms of the unique effect they had on conversations when accepted by the participants (Stalnaker, 1978).

Speakers use language to describe the world. Among many speech acts at their disposal, they use assertions as one speech act intimately connected with truth (Weiner, 2005). Assertions allow them to introduce information in the form of asserted statements. The meaning of statements can be determined by relating them to propositions that have truth-conditions (Frege, 1918a, 1956). In turn, assertions, as they relate discourse to facts, got the lion's share of attention. Assertions, famously difficult to characterise (MacFarlane, 2011b), stood as a missing link between speech acts and truth. Other speech acts, like rejection or retraction, were mentioned in relation with assertion. But the idea was that, in a cooperative picture of conversation, once we got assertions we would get other speech acts.

The other half of the chicken. Speakers do not only use language to describe the world. Meanings are not only a matter of truth-conditions, but also of the uses speakers make of language (Wittgenstein, 1953). *Inferentialism* is the view that the meaning of lexical items is determined by their inferential role (Carnap, 1934; Dummett, 1981, 1991; Gentzen, 1934; Prawitz, 1965, 1971). Instead of accounting for meanings in terms of truth conditions, we can account for the meaning of a word by taking only into account its deductive introduction and elimination rules. For example, the meaning of the connective *and* (\wedge) is determined by its introduction and elimination rules:

$$\wedge\text{-intro} \frac{A \quad B}{A \wedge B} \quad \frac{A \wedge B}{A \quad B} \wedge\text{-elim}$$

Expressivism is the view that the meaning of (some) lexical items is determined by their expressive role (Blackburn, 1984; Brandom, 1998; Gibbard, 2003). Instead of accounting for meanings in terms of truth conditions, we can account for the meaning of certain words by taking them to express a speaker’s attitude. For example, when a speaker utters *Cannibalism is wrong*, the meaning of the word *wrong* in this utterance is to express the speaker’s disapproval. The EXPRESS project, led by Luca Incurvati, worked on combining inferentialism and expressivism. With *inferential expressivism*, we can account for the meaning of some lexical items as being given by their deductive role in terms of speakers’ attitudes (Incurvati and Schlöder, 2021).

Speakers use language to dissent and disagree. Some speech acts serve to prevent conversational moves, remove information. These speech acts have been taken as complementary to assertion, and side-stepped. However, in an inferentialist and expressivist perspective, these speech acts are as relevant to determining the information at stake in a conversation as assertions (Rumfitt, 2000; Smiley, 1996). They deserve being treated as important in their own right (Dickie, 2010; Price, 1983).

My work, as summarised in sections 1.1 and 1.2, focuses on the speech acts speakers use to express disagreement. I don’t characterise them in relation to assertion, but for their own sake. With weak rejections and weak assertions, speakers prevent speech acts they disagree with from being made (chapter 2; chapter 4). With retractions, they undo the effect of previous utterances they no longer agree with (chapter 5). When studying these speech acts in their own right, we can use them, in inferential expressivist fashion, to explain the meaning of specific lexical items. I use weak rejection to explain the contrast that the adversative marker *but* contributes in sentences (chapter 3). I use weak assertion to explain the behaviour of *perhaps* and *might* (chapter 4). A study of the speech acts that speakers use in disagreement naturally leads to considering conflictual settings of conversation. Chapter 6 and chapter 7 concern what it means for a theory of communication to take into account speakers conflicts and disagreements. In them, I develop a non-ideal picture of conversation and

of communication by taking into account not only speakers' attitudes, but also audience attitudes.

With this dissertation, I approach disagreement, and the speech acts that occur within, as significant for their own sake. Thus, I challenge an idealised picture of conversation, and instead put conflict first.

1.1 The speech acts of disagreement

1.1.1 Rejection

The first type of speech acts that occur in contexts of disagreement are *rejections*: the speech act by which speakers prevent utterances from being accepted in the conversation. Chapter 2 is the result of recurring investigations on the speech act of rejection. Rejection has been characterised as the opposite side of assertion. Where an assertion of P amounts to answering the question *Is it the case that P ?* – *Yes*, a rejection of P is the *No* answer. But, in conversation, we don't only say *No* to assertions: we also reject orders, questions, proposals. And we don't only say *No* when we don't believe P is the case: we reject statements that are true, but we do not wish to integrate in the conversation. In chapter 2, I first set out to investigate the messiness of rejection. It comes, I argue, from characterising rejection in virtue of its effect, as a speech act by which one prevents the effect of a speech act from entering the conversation. In the case of disagreeing with a statement, a rejection prevents the information of this statement from becoming part of what participants agree upon. But this effect-based characterisation makes no difference between rejecting statements we disagree with on epistemic grounds, and rejecting statements just because we don't want them in the conversation. Then, I turn to the conditions of rejections, when they concern information participants must agree or disagree upon. What attitude must a speaker hold with respect to a statement to sincerely reject a statement? I aim to characterise rejection, as the speech act that may block an assertion, in norm-based accounts of speech acts. To do so, I build possible norms for rejection, based on candidate norms for assertion, and assess whether they account for speakers practices of rejection.

1.1.2 Adversation

Inferential expressivism holds that the meaning of some lexical items is determined by inferences involving speakers' attitudes. For example: we can understand the meaning of *not* in virtue of inferences involving rejection (Incurvati and Schlöder, 2021). But this is not the only use of rejection. After characterising the speech act of rejection, we can apply it to the interpretation of other lexical items that convey speakers dissent.

Chapter 3 started as a project to clarify the interpretation of the adversative marker *but*. Debates surround the interpretation of the contrast it contributes. Is it a contrast explainable only in semantic, context-independent terms? Is it an inferential contrast, on the pragmatic side, highly context-dependent? Because *but* is a versatile marker, contributing a variety of contrasts in different contexts, none of these ‘unifying’ approaches had a ready-made explanation for all cases. They appealed to pragmatic processing to cover all cases. I did not find this appeal very satisfying: if there is pragmatic processing, there must be pragmatic rules at stake? Since the different approaches to the adversative marker all mentioned a *yes-no* polarity, I interpreted this polarity in terms of attitudes expressed by indirect speech acts of assertion and rejection. This yields the necessary explanation for the pragmatic processing at stake in accounting for *but*-sentences.

1.1.3 Weak assertion

In chapter 2, I followed Incurvati and Schlöder (2017) in characterising weak rejection as a speech act that blocks a statement *P* from being added in the conversation. But since there are speech acts for blocking assertions, we can also postulate speech acts that block negative assertions. While weak rejection is a speech act sufficient to block an assertion without committing to its negation, *weak assertion* is the speech act sufficient to block a strong rejection (Incurvati and Schlöder, 2019). Chapter 4 focuses on the speech act of weak assertion. Not a full-fledged assertion, weak assertion leaves open a possibility in the conversation without committing to it. In chapter 4, I gather linguistic evidence, both qualitative and quantitative, that *perhaps* and *might* are lexical markers of weak assertion. *Perhaps* embedding behaviour corresponds to the behaviour of a force modifier, while the behaviour of the epistemic modal *might* licenses its interpretation as a modifier corresponding to weak assertion.

This chapter is a first attempt at gathering empirical data. It provides new examples to assess the use of *perhaps* and *might* taken from two corpora of English. With an associated annotated list of results, it also paves the way for further quantitative studies.

1.1.4 Retraction

Rejection and weak assertion are self-standing speech acts that preemptively block adding certain information to a conversation. But blocking is not the only way to work with negative information. *Retraction* is a second-order speech act that undoes the effects of a previous utterance. Chapter 5 focuses on characterising the speech act of retraction. Retraction is the speech act by which a speaker takes back a previous utterance. It cancels the illocutionary effects of the target utterance. In previous works (Caponetto, 2020; MacFarlane, 2014), retractions are studied in relation to assertion, and assumed to immediately cancel the effect of

the previous utterance. Instead, in chapter 5, we study retractions for their own sake, not only as a counterpoint to assertions. By doing so, we paint a picture of retractions as a proposals on a conversation. After delineating a picture of retractions as proposals, we set to investigate the effect of retractions on utterances that occurred between the target utterance and the retraction.

1.2 Theories of conversation

1.2.1 Interpreting silence

The work on speech acts that forms the first part of this dissertation is informed by a stalnakerian perspective on speech acts. According to Stalnaker (1978; 2002), we can represent the effect of speech acts as updates on a set of possibilities compatible with what the participants in a conversation have agreed upon. But, in this perspective, conversational updates are accepted unless they are rejected. This idealised picture of conversation leads to an interpretation of silence as expressing a default assent. In chapter 6, I challenge this picture of silence as giving assent by default. Instead, I propose a fine-grained analysis of silence. Crucially, a speaker can assess the cooperativeness of a conversation, based on priors and pragmatic cues. Then, they ascribe a default attitude to the audience, which supplies an interpretation for silence.

1.2.2 Interpreting communication

Studying the contribution of the adversative marker *but* in chapter 3, and in particular wondering whether it occurred at a semantic or pragmatic level, led me to the topic of Fregean colourings. In their Fregean characterisation, colourings are lexical items or grammatical constructions that do not affect the truth of a sentence, but may convey additional contents. The main question I tackle in chapter 7 is the status of the content that is conveyed by certain colourings: is the content of colourings merely conveyed, or is it part of what a speaker communicates with their utterance? This depends on the notion of communication at play. I defend a notion of communication that is intersubjective, where a speaker also communicates what their audience infers them to intend to communicate.

In chapter 8, I summarise the most important findings of this collection, and open further directions for the study of speech acts in settings of disagreement.

1.3 Note on chapters material

With the exception of chapter 5, I am the sole author of the material in this dissertation.

The material in chapter 2 was presented, in different forms and under different titles, to audiences at the ICSO V Workshop: *Perspectives on Content* (Buenos Aires, 2018), the EXPRESS Workshop: *Bilateral Approaches to Meaning* (Amsterdam, 2019), and the OZSW conference (Amsterdam, 2019). I benefitted immensely from these conversations. The material in chapter 3 was presented at the MLC seminar in Amsterdam (2020). I also thank Grégoire Winterstein for his very detailed comments on a previous version of this chapter.

The general idea of treating retractions as a proposal originated from reading and discussing the work of Caponetto (2020) and Marques (2018) during the EXPRESS reading group. The general idea of treating retractions as a proposal originated from discussions during the EXPRESS reading group. As such, the material in chapter 5 is based on joint work with Luca Incurvati, Giorgio Sbardolini and Julian Schlöder. We decided together on an interpretation of retractions as proposals on Common Ground, which is the central idea of this chapter. The content of the chapter was further developed over the course of multiple meetings and conversations. In chapter 5, I wrote down the results of these joint discussions and added explicit Common Ground models of the effects of retractions and detailed studies of retraction failures. Since this is joint work, I use first-person plural throughout. The material in chapter 5 was also presented at the EXPRESS Workshop on *Non-Assertoric Speech Acts* (online, 2021), and Dan Zeman's workshop *Theoretical and empirical challenges to Retraction* (online, 2021).

The material in chapter 6 is based on published work (Bussière-Caraes, 2021). The material in chapter 7 is based on published work (Bussière-Caraes, 2022).

2.1 Introduction

(1) No, it is not raining.

When she utters (1), the speaker *rejects* the statement *It is raining*. With this speech act, she expresses her dissent from the statement. Rejection is the dual of *assertion*, the speech act by which one expresses that they agree with a statement.

I treat assertion and rejection as distinct activities on all fours with one another. Where assertion expresses assent to a proposition, rejection expresses dissent from it, in each case by means of a sentence whose sense is the proposition in question. (Smiley, 1996, p. 1)

When a speaker rejects a statement, she disagrees with it.

A rejection, as an expression of dissent from a statement, is self-standing. While it can, it needs not reply to an assertion, or another speech act. But ‘treating assertion and rejection as distinct activities on all fours with one another’ requires investigating the conditions and norms that govern a speaker expressing dissent from a proposition. In this chapter, I analyse rejection in relation with assertion, and assess candidate norms for rejecting a proposition.

2.1.1 Rejection as a distinct activity

Early on, Frege (1918b) dismisses rejection. He takes a sentential question to express a content of possible judgment. With an assertion, the speaker associates an affirmative judgment to the thought expressed in a sentential question. Then, Frege asks whether a negative judgment, which would correspond for us to the attitude speakers express with rejection, should be considered a separate act of the speaker.

Are there two different ways of judging, of which one is used for the affirmative, and the other for the negative, answer to a question? Or is judging the same act in both cases? Does negating go along with judging? (Frege, 1918b, p. 153, translated in Beaney, 1997, p. 356)

According to Frege, *the assumption of two different ways of judging must be rejected* (Frege, 1918b, p. 154, Beaney, 1997, p. 357).

Thus for every thought there is a contradictory thought; we acknowledge the falsity of a thought by admitting the truth of its contradictory. The sentence that expresses the contradictory thought is formed from the expression of the original thought by means of a negative word. (Frege, 1918b, p. 154, translated in Beaney, 1997, p. 358)

This leads to a traditional picture according to which defining the speech act of rejection is superfluous, as a rejection is naught but the assertion of a negation.

A sentential question admits of two possible answers, ‘Yes’ and ‘No’ ... the answer ‘No’ is precisely tantamount to, and is best analyzed as, an assertion of the negation of the sentence uttered in interrogative form. (Dummett, 1981, p. 425)

But, contrary to this Fregean *doxa*, the speech act of rejection is not equivalent to the assertion of a negation (Incurvati and Schlöder, 2017; Price, 1983). That a speaker rejects *it is raining* does not mean that they wish to assert its negation, and thus add the information that *it is not raining* to the conversation.

(2) There might not be a seminar tomorrow.

(3) B rejects the statement that A asserts:

A: Adèle or Bernard will win the election.

B: No, Charles is also a candidate.

Examples (2), (3) are cases where the speaker rejects a statement, but does not wish to assert its negation. These correspond to Incurvati and Schlöder’s *weak rejection*.

The notion of *weak rejection* complicates the picture of rejection (section 2.2). It is grounded in a characterisation of rejection from its effect on a conversation: rejection is the speech act by which a speaker can prevent the participants in a conversation from adding a statement to their shared information.

A full-fledged treatment of rejection and assertion *as distinct activities on all fours with one another* brings up additional questions: can we characterise rejection positively, and find out under which conditions speakers may perform it?

2.1.2 Norm-based approaches to speech acts

Norm-based approaches seek to uniquely characterise speech acts by providing a picture of the norms that govern them. With the expression *norm-based approaches*, I refer specifically to the family of approaches according to which speech acts, such as assertion, obey one constitutive norm that uniquely characterises them. That is, according to *norm-based* approaches, to perform a certain speech act, the speaker needs to meet a certain condition – otherwise, the performance can be denounced as an abuse of the speech act. However, it should be noted that other ways of characterising speech acts, such as Common Ground (Stalnaker, 1999) and Commitment (Brandom, 1983, 1998) analyses of speech acts are *also* normative pictures of speech acts. The distinction between *norm-based approaches* and Common Ground and Commitment approaches lies in the fact that norm-based approaches associate *upstream* norms with speech acts, instead of *downstream* norms (MacFarlane, 2011b).

But, while the constitutive rules approach looks at “upstream” norms — norms for *making* assertions — the commitment approach looks at “downstream” norms — the normative *effects* of making assertions. (MacFarlane, 2011b)

A stalnakerian analysis provides a picture of the *normal effects* of speech acts: what *should* occur on the information structure of a conversation when a speaker performs a speech act, if everything goes according to plan (more on that in section 2.2.2). Conversely, norm-based approaches look for *necessary conditions* that need to be met for the performance of a speech act.

Within norm-based approaches to speech acts, a lot of attention has been given to defining the constitutive norm of assertion (J. Brown and Cappelen, 2011; Goldberg, 2015, 2020c; MacFarlane, 2011b). Norm-based accounts try to find the norm that uniquely defines a speech act, following a normative schema such as:

[Speech act] *P* only if φ *P*.

In this picture, φ is the *necessary* condition that enables a speech act to have its unique illocutionary effect (Austin, 1975). Various norms have been suggested for assertion: the Truth Norm (Weiner, 2005), the Belief Norm (Bach and Harnish, 1979; Mandelkern and Dorst, 2021), the Reasonable or Justified Belief Norm (Lackey, 2007; McKinnon, 2013), the Knowledge Norm (DeRose, 2002; Engel, 2008; Williamson, 1996).

2.1.3 Outline

In this chapter, my goal is to investigate possible constitutive norms for rejection. First, I account for the *messiness* of rejection. Rejection is *messy* because, in a

downstream characterisation of speech acts, rejection applies both to statements and utterances (section 2.2). But, in an *upstream* account of speech acts, rejection is characterised with respect to a content.

A conditions-based picture of rejection is much less messy, as it focuses on speakers' disagreement with a statement. To investigate possible norms of rejection, I start from the stalnakerian notion that a rejection is sufficient to block an assertion, and work on characterising and evaluating candidate norms for rejection, dual to candidate norms for assertion. Because of the interdefinition of norms for assertion and rejection, a discussion of the norms for rejection also sheds light on the debate surrounding norms of assertion. In section 2.3, I defend normative duality. The candidate norms for rejection complement candidate norms for assertion. In section 2.4, I establish the possible norms for rejection given four proposals of norms for assertion. Given normative duality, the way to build a norm for rejection is to flip the parameter given by the corresponding norm of assertion. However, not all norms fare equally with respect to weak rejection – the cases where a speaker rejects a statement, while being unwilling to assert its negation. In section 2.5, I discuss the divergence of the epistemic norms, and advocate for a Knowledge Norm of assertion and rejection.

2.2 The messiness of rejection

The goal of this section is to address a worry that comes up in the discussion of rejection in speech acts theory: that rejection is a speech act too messy to work with (Dickie, 2010). If rejection is, fundamentally, what happens when speakers say *No!* to something, then we run into a problem. Speakers do not only reject potential assertions. They also reject promises, questions, proposals, etc. Speakers do not only reject on epistemic grounds. They also reject irrelevant, impolite, incomplete, interventions. However, I think this worry comes from confusing competing perspectives on speech acts, that have different foci. Once we know which perspective we work with, rejection becomes much less messy.

2.2.1 Classifying theories of speech acts

There are many different ways to approach the analysis of speech acts. I offer a principled classification. Let's start from the example of assertion. There are four families of accounts of assertions. They, and some of the main arguments against each of them, are summarised by MacFarlane (2011b).

- i. To assert is to express an attitude.
- ii. To assert is to make a move defined by its constitutive rules.
- iii. To assert is to propose to add information to the conversational Common Ground.

iv. To assert is to undertake a commitment.

These four accounts of assertion correspond to four different perspectives on speech acts. MacFarlane distinguishes between characterisations of assertion that define *upstream* and *downstream* norms. *Upstream* norms characterise conditions that a speaker needs to meet to perform a speech act. For example, in order to make a promise sincerely, a speaker needs to intend to fulfil it. *Downstream* norms characterise results that normally obtain after a speaker performs a speech act. For example, after making a promise, a speaker is held responsible to fulfil their promise.

In addition, I divide theories of speech acts according to their focus, either on the informational aspects or the social aspects of a speech act. This new typology proves useful in understanding different perspectives on speech acts. A theory that focuses on the effects of speech acts on an information structure (Stalnaker, 1978, 2002) will attempt to characterise speech acts by finding their unique effect on the information structure.

This classification shows that the four most common contemporary theories of assertion exhaust the possible combinations between two distinctions that seem relevant to speech acts theory.

Conditions vs. Effects. When saying that to assert is to express an attitude, or to assert is to make a move defined by constitutive norms, one characterises assertion with respect to the *conditions* that have to be fulfilled for an assertion to be made. These approaches focus on the preliminary conditions of assertions. The assertion as expression of an attitude approach posits that, when a speaker has a certain attitude with respect to the content of their statement, the statement is asserted. In the norm-based approach, a statement must comply with certain norms to be properly asserted. Thus, both those accounts hold that these preliminary conditions are what individuates the speech act of assertion.

By contrast, the Common Ground approach and the commitment approach do not define the category of assertion from the conditions the speaker complies with. Instead, they focus on the *effect* assertions have on the audience. In the Common Ground approach, an assertion is seen as a proposal to update the information participants share. Therefore, if accepted by the audience, we model its effect as modifying the Common Ground of information, by deleting from it possible worlds that are incompatible with the propositional content of the assertion. Suppose an assertion that *it is raining* is accepted. That means the participants agree to update the Common Ground: they remove all possible worlds where it is not raining from the previous state of information. In a similar way, the commitment approach claims that the making of an assertion produces a certain effect on what can be called the ‘social situation’ of the conversation. The commitment approach characterises assertions as utterances from which follows a commitment of the speaker to the truth of their statement. This commitment

is also recognised by the audience, which can challenge the speaker and ask them to provide justifications for their statement.

Information states vs. Social states. Both the assertion as expression of an attitude approach and the Common Ground approach take into account *information states*. In the first case, the attitude of the speaker is the content expressed by the assertion. In the second case, the effect of an assertion is defined in epistemic terms. An update to the Common Ground models a change in the information state shared by the participants. In this sense, the Common Ground account of assertion focuses on the information state of the participants, and analyses how this state gets updated once an assertion is accepted.

But one may focus on *social* norms and effects governing assertions. The norm-based approach takes external norms to govern the making of an assertion, and the commitment approach analyses the effect of assertions in social terms. In the norm-based account, certain conditions stated by the norms governing assertions are assumed to be fulfilled when a speaker makes an assertion. In the commitment approach, the assertion made by a speaker has a certain normative result – the speaker is held responsible for the truth of an assertion and can be compelled to undertake a justificatory procedure to show that they fulfil their commitment.

For example: the assertion as expression of an attitude view qualifies assertion in virtue of the epistemic state of the speaker, and the epistemic state the speaker wants to obtain in their audience. The speaker believes the content of their assertion, they want their audience to believe the content of their assertion. In norm-based approaches, assertion is qualified by a norm, let's say *Only assert what you know*. This norm is *socially* enforced in conversation, by challenging speakers to verify that they know the content of their assertions, and requiring them to retract their assertions if they do not know it (MacFarlane, 2014). Common Ground models qualify assertions in virtue of their information update behaviour. Assertions are the speech acts that *narrows down* the set of possibilities compatible with what speakers agree upon. In this sense, Common Ground characterisations of assertion focus on its effect on information states. Commitment approaches qualify assertions in virtue of the commitments that follow them. A speaker who made an assertion is socially and publicly committed to the truth of their assertion, held responsible to provide justification, etc.

By applying these two distinctions, we obtain the typology summarised in table 2.1.

This typology shows that the competing theories individuate assertions, or other speech acts, according to different features, which come from different aspects of the phenomena. While the features that each theory picks for assertion are usually presented as the only individuating feature of assertion, there is nothing, in principle, to prevent someone from thinking that assertions both have

	Focus on information states	Social norms focus
Condition-based approaches	<p>Assertion as Expression of an attitude</p> <p>The speaker has an attitude with respect to the content of the utterance</p> <p>The attitude supports the assertion</p>	<p>Assertion as a move defined by norms</p> <p>There is a norm for assertions</p> <p>The speaker respects the norm</p>
Effect-based approaches	<p>Assertion as a proposal to update to the Common Ground</p> <p>The assertion has an effect</p> <p>The audience updates the Common Ground with the content of the assertion</p>	<p>Assertion as a commitment</p> <p>The assertion has an effect</p> <p>The speaker is held responsible for providing a justificatory procedure for their utterance; they issue a warrant to their audience to be committed to the content of their utterance.</p>

Table 2.1: Classification of theories of Assertion

unique conditions and unique effects that distinguish them from other speech acts. In particular, given that the different theories of assertion characterise the phenomenon according to distinct aspects – informational and social – one might well think that whereas the *conditions* of assertion are defined by a unifying social norm, its *effects* are best characterised within a Common Ground approach, that models the effect of assertions on the information states of the participants.

2.2.2 Effect-based models of conversation

In effect-based theories of conversation, speakers agreement, or disagreement, with a claim is crucial to keep track of the information introduced in a conversation. Common Ground pictures of conversation (Stalnaker, 1978, 1999, 2002) characterise speech acts according to their effect on the pool of information that the participants share. Stalnaker represents conversational information as a set of possibilities (modelled as *possible worlds*) compatible with what participants have agreed upon. For my purposes, it is sufficient to think of possible worlds as useful modelling entities, which make propositions true or false. The reader may choose to interpret them as full-blown metaphysical entities of the same sort as our world (Lewis, 1986), or to declare that there are no such things as possible worlds (Quine, 1948). What is relevant is that possible worlds are *modelling tools*

(Stalnaker, 1984) that can keep track of actions or information that are compatible or incompatible with what the participants in a conversation have publicly agreed upon. In this type of models, the set of possible worlds compatible with what the participants have agreed upon is the *context set*. The context set, plus a representation of how participants can introduce new information constitutes a Common Ground model. The structure of the Common Ground can be more or less complexified to account for the effect of different speech acts. Farkas and Bruce (2010) implement the effect of polar questions in a Common Ground like model, while Portner (2004, 2007) represents imperatives as proposals to update participants to-do lists.

Alternatively, a conversation can be represented as a score that the participants keep track of (Lewis, 1979, 1980). Participants in a conversation are aware of which utterances have been made so far, and how the information they agreed upon has been modified. Lewis represents this awareness by means of a *conversational record*, developed by analogy between the game of language and the game of baseball, that keeps track of the different moves that have been performed in conversation. States of the conversational record have a normative effect on the conversation: at a given state of the record, certain moves are available, others are forbidden.

Most models of conversation as an information exchange, where speakers update an information state, rely on some combination of a Common Ground and scorekeeping picture. In a Common Ground-like picture, the speech acts performed by the participants update the information they share, by reducing the number of possibilities compatible with what they agree upon, to try to get closer to an exact information state. This is the principle behind Veltman's Update Semantics (Veltman, 1990, 1996). In a scorekeeping picture, the participants keep track of discourse referents that have been introduced, moves that have been made in a conversation, etc. This is the principle underlying, for example, Segmented Discourse Representation Theory (Asher and Lascarides, 2003), but also Dynamic Predicate Logic (Groenendijk and Stokhof, 1991). These two approaches to information can be combined. Multiple pictures of the Common Ground integrate a scorekeeping device (Camp, 2018).

When speakers make moves in conversation, they undergo certain responsibilities (Austin, 1975; Brandom, 1983). Certain conversational moves have preliminary conditions to be performed; participants are responsible to ensure they meet these conditions (Austin, 1975). From certain conversational moves, epistemic or deontic commitments follow for participants (Brandom, 1983, 1998). As seen in table 2.1, when a speaker performs an assertion, they undertake a responsibility to provide a justificatory procedure if challenged. I also take these speakers' commitments to be part of the information publicly exchanged in a conversation.

For my purposes, the choice of model is not really relevant – at least until chapter 5. What matters is to represent a conversation as an information ex-

change, where the participants bring new information. Hence, I adopt a idealised Common Ground model where possibilities compatible with what participants have agreed upon are represented as possible worlds. I also assume that this model of conversation keeps track of the conversational score (Lewis, 1979), and of the commitments that participants have publicly endorsed (Brandon, 1983). In this idealised Common Ground model, all the participants in the conversation are aware of what possibilities are compatible with what they have agreed upon, and they are all aware that they have access to this information.¹

2.2.3 Modelling rejection in effect-based models

Stalnaker models the effect of speech acts as update operations on the Common Ground. This characterises speech acts in virtue of their *information change potential*. The paradigmatic example that Stalnaker uses to model speech acts as update proposals is assertion.

[...] how does the *content* of an assertion alter the context? My suggestion is a very simple one: to make an assertion is to reduce the context set in a particular way, provided that there are no objections from the other participants in the conversation. The particular way in which the context set is reduced is that all of the possible situations incompatible with what is said are eliminated. (Stalnaker, 1999, p. 86)

He represents the effects of an utterance of *It is raining* with assertive force as a proposal to update the Common Ground by adding the information that *It is raining* to the conversation. With this new information, epistemic possibilities are ruled out: they are not compatible anymore with what the participants have agreed upon. On the Common Ground model, the effect of an assertion that *it is raining* (*R*) is modelled as an update that removes all *non-raining* worlds from the set of possible worlds (fig. 2.1).

On the other hand, the rejection of *It is raining* (1) prevents adding the information that *it is raining* to the conversation. Or, it prevents ruling out the possibility that *it is not raining*. In a Common Ground model, the result of a rejection is to block an update of the Common Ground:

... the essential effect of an assertion is to change the presuppositions of the participants in the conversation by adding the content of what is asserted to what is presupposed. This effect is avoided only if the assertion is rejected. (Stalnaker, 1978, p. 86)

¹I simplify the Common Ground picture a lot here. In particular: I replace a Common Knowledge requirement, that Lederman (2018) argues is neither possible, nor necessary to ground a conversation, with an awareness requirement.

In other words: when a participant rejects *it is raining*, they prevent removing *non-raining* worlds from the possibilities compatible with the common beliefs of the participants.

The Common Ground approach provides a *downstream* and *negative* characterisation of rejection (MacFarlane, 2011b). Downstream because it characterises rejection according to its effects on conversation, not its conditions. Negative because it characterises rejection according to a phenomenon it prevents from happening. Rejection is the speech act whose essential effect is to block an update to the Common Ground.

2.2.4 Explaining the messiness of rejection

By modelling the effect of rejection as *blocking* an update, a Common Ground approach to rejection represents Price (1983) and Incurvati and Schlöder (2017)’s understanding of rejection more than Frege (1918b) and Dummett (1981)’s.

In his article, Price (1983) highlights the fact that a rejection reveals the incorrectness of a target utterance (actual or potential). This incorrectness may be due to the content of the utterance or its performance:

To resolve this new confusion, we must distinguish two ways in which the description ‘incorrect’ may be used in reply to an assertion that S. On the one hand is ‘It is incorrect to assert that S’, which (as it were) is a statement of semantic theory, providing information as to S’s assertion conditions (on a par, for example, with “S” is assertible’). While on the other is ‘That’s incorrect’, said in response to an assertion or suggestion that S – which is on a par with the assertion itself. With this distinction, even if – in the intuitionist style – we take the latter statement to be assertible when (it is recognisable that) S’s assertion conditions do not hold, we shall not have to say that the statement’s content is that this is so. (Price, 1983, p. 164)

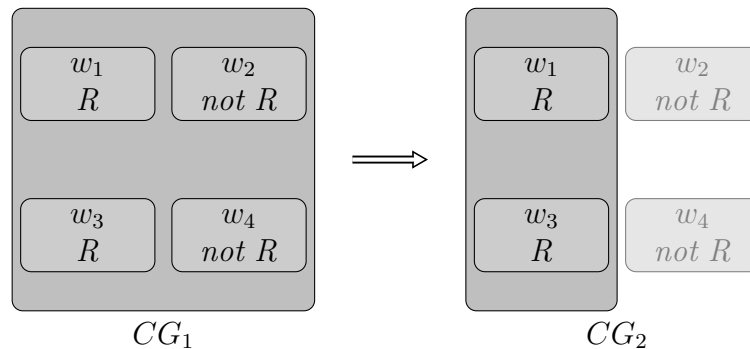


Figure 2.1: Updating the Common Ground with an assertion that *R*

Rejections are ambiguous in what they target. Some address propositional content:

- (4) a. It is raining.
b. (No!) It's not raining.

The utterance in (4b) need not *respond* to a previous speech act. In fact, it can be a standalone speech act. This corresponds to Price's first notion of incorrectness, where the statement is deemed unassertible in virtue of its semantic content.

The second type of rejections address the performance of a speech act itself:

- (5) You can't say that your boss is an idiot in a work meeting!

In (5), the speaker does not target a propositional content, but the performance of a speech act. The statement is unassertible in virtue of other norms surrounding assertions: social, epistemic, conversational, etc. The Gricean maxims of communication come to mind when assessing the performance of speech acts (1975).

This is how I cash out the distinction that Price makes between two types of incorrectness. An utterance can be rejected as incorrect in virtue of its content, but also in virtue of the illocutionary force associated with it, because it doesn't respect certain norms of conversation, etc. In the second case, rejections do not only target assertions, but also requests, promises, proposals, questions, etc. Moreover, rejections can have different grounds: falsity, but also irrelevance, impoliteness, etc. One could reject a question as irrelevant in a certain conversation, an order because the person issuing it does not have the authority to do so, a promise as they do not believe it, ...

Effect-based accounts of speech acts, such as Common Ground or commitments approaches, or the one I adopt in section 2.2.2, characterise rejection in virtue of its effects on a conversation. Whether the rejection is on the basis of an utterance's content or performance does not matter. What matters is that the speaker who rejects refuses to commit to the utterance, and refuses that the target utterance (actual or potential) fulfils its intended effect over the conversation. For these reasons, effect-based (*downstream*) characterisations of rejection group together the different types of rejection that Price (1983) characterises.

By contrast, condition-based accounts of rejection focus on rejections that are performed in virtue of the content of an utterance. These are cases where rejection is a speech act that targets an information. A complete account of rejection requires elucidating the norms of rejection, which I now set to do.

2.3 From assertion to rejection

While there is little consensus on what properly characterises a rejection, one thing scholars agree upon is that the rejection of a statement is incompatible with

its assertion (Dummett, 1991; Frege, 1983; Ripley, 2015, *inter alia*). To focus on a norm-based account of rejection as a speech act that blocks assertions, I analyse the incompatibility between rejection and assertion in terms of *unassertibility*. The norms that have been proposed to characterise assertion qualify a necessary condition for *assertibility* (section 2.3.1). A statement that does not meet the condition cannot be sincerely asserted. When a speaker rejects a statement, they deem this statement unassertible under their criteria for assertion. The speaker needs not make explicit this inference from rejection to unassertibility. In fact, they rarely do. However, the inference licenses a defence of normative duality: according to which the norm for rejection is complementary to the norm for assertion (section 2.3.2). Normative duality underlays the construction of norms for rejection given different norms for assertion (section 2.4).

2.3.1 The norms of assertion

The notion of language game, put forth by Wittgenstein (1953, 2009), analyses discursive practices by analogy (*family resemblance*) with games humans play. In games, there are rules that constrain players moves, and determine the result of their actions. In language, we can similarly (*pace* Maitra (2011)) qualify speakers' utterances according to the rules they obey, and their result on a conversation (Lewis, 1980, 2002). Of course, rules vary across language games and social situations. Lyotard (1984) gives the example of the unsolvable conflict between the discourse of the Shoah survivor and the discourse of the negationist historian (n° 2). The discourse of the negationist historian obeys rules about justification where the assertion *There were gas chambers* can only be justified by direct testimonial justification: a survivor having experienced a gas chamber. The discourse of the survivor (and most historians) holds that noone who has experienced a gas chamber survived. Other examples abound: the norms that govern a scientific discourse are different from the norms of a work of fiction. The norms governing linguistic and conceptual practices in late Medieval Europe are radically upended by the Enlightenment change of scientific paradigm (Foucault, 1966, 1972; Koyré, 1957).

Norm-based approaches to speech acts claim that, while some norms that govern speech acts vary across language games, other norms are invariant (Sbisà, 2018). In some language games, speakers may break conversational maxims that advise on optimal speech acts performance (Grice, 1975). In most language games, some speech acts have material requirements to be performed (Austin, 1975). For example, in order to correctly issue a command, the speaker needs to have relevant authority over the addressee. But, in addition, speech acts have *constitutive* rules, upon which their performance depends. And the constitutive rule of a speech act states a unique, necessary condition to perform it.

[Speech act] P only if φP .

Where φ is the constitutive rule for the speech act.

The conversational maxims help participants determine whether performing the speech act is a good idea with respect to common conversational goals. The material requirements help participants determine whether the speech act is correct in the situation to which it relates. By contrast, the performance of a speech act depends on the constitutive rule. For example, if the constitutive rule for a command is that the speaker intends the addressee to conform to the command, the speaker cannot sincerely issue a command if they do not intend the addressee to conform to it.

If the constitutive rule is violated, the utterance constitutes an *abuse* of the speech act (Austin, 1975, pp. 16–18). The speech act might have its intended illocutionary effect, but it shouldn't. For example, suppose that intending the addressee to conform to its content is a constitutive rule of commands. When a speaker issues a command that she does not intend her audience to execute, her command might be successful, in that the audience executes it. But she is abusing the speech act, as she performed the speech acts under false pretenses.²

The constitutive norm of assertion is famously difficult to establish (Cappelen, 2011; Montminy, 2013; Pagin, 2016). Candidate explanations are of the form:

Norms of Assertion.

Where φ qualifies the truth, or epistemic status of P .

One must: assert P only if φP .³

Four norms have been suggested for assertion (MacFarlane, 2011b; Pagin and Marsili, 2021)⁴

The *Belief norm of assertion* is a weak norm of assertion, according to which the speaker must believe their statement to assert it (Bach and Harnish, 1979; Mandelkern and Dorst, 2021). Bach and Harnish take it to be the only fundamental rule that we blame a speaker for breaking.

²In that sense, violating the constitutive norm is different from violating other preliminary norms. If other preliminary norms are violated, the speech act constitutes a *misfire*, and does not produce its intended illocutionary effect. For example: an order given by someone who has no authority to give an order, and is known to have no such authority, does not take effect (Austin, 1975, pp. 16–18).

³Like other constitutive norms, the norms for assertion follow a *only if* schema. If a speaker follows a Truth Norm for assertion, it does not mean that she commits to asserting *all* true statements. It does mean that she commits to the truth of the statements she asserts.

⁴In addition to Cappelen (2011) and Maitra (2011) position of global distrust towards norm-based accounts of assertion, each of these proposed norms of assertion encounters local criticism (see MacFarlane, 2011b; Pagin and Marsili, 2021, for an overview). These objections are not the topic of this chapter. I will however make some note of them when discussing these norms in relation to rejection.

At the other end of the spectrum, the *Truth norm of assertion* enjoins the speaker to strive to only assert true statements (Weiner, 2005). That we often ask speakers to retract assertions that turn out false, even if they had no way to know, supports that truth is the goal, or requirement of assertions (MacFarlane, 2014).

The *Knowledge Norm of assertion* holds speakers to higher epistemic standards than the Belief Norm of assertion (DeRose, 2002; Engel, 2008; Williamson, 1996). Knowledge is what distinguishes assertions from guesses, or conjectures.

The *Justified Belief Norm of assertion* grounds assertions in justification, thus relaxing the epistemic standard compared to the Knowledge Norm of assertion (McKinnon, 2013).

The *Reasonable Belief Norm of assertion* stems from the observation that, sometimes, speakers don't know, or even believe what they assert (Lackey, 2007). It suggests relaxing the dependence of sincere assertions on the speaker's own epistemic attitude, and favouring a notion of justification. Operating under a *Reasonable Belief Norm of assertion* would license a speaker to assert a statement that she is not certain of, but that she has grounds to believe is the most reasonable to assert. Lackey (2007) gives the example of a religious teacher who, while she personally believes (on grounds of faith) that a divinity created humans from scratch, holds the theory of evolution to be more reasonable to believe (on grounds of scientific evidence). Therefore, the religious teacher can sincerely assert the theory of evolution, while not knowing it for certain, or even fully believing it.

The question is, for each norm, whether they can provide a dual account of rejection. Rejection points out the unassertibility of a statement; can we positively characterise this unassertibility as a converse to the norm for assertion?

2.3.2 Normative duality

But why would the norm for rejection be dual to the norm for assertion? I start from the incompatibility of assertion and rejection, that I take as a primitive. A speaker cannot assert P and reject P in the same context lest they appear irrational. In this picture, rejection is no less primitive than assertion: they stand on equal grounds as incompatible moves.

Incompatibility, taken as a primitive, triggers two inferences. If a speaker asserts P , it follows that they are not willing to reject P . If a speaker rejects P , it follows that they are not willing to assert P . This is the minimum by which we define rejection in stalnakerian terms as the move that blocks an assertion.

We can then gloss rejection in terms of *unassertibility*. By rejecting a statement, the speaker indicates that they take it to be *unassertible*. This analysis does not threaten an equal treatment of assertion and rejection. The inference from rejection to unassertibility mirrors an inference from assertion to unrejectability. In this sense, the speech acts remain *on all fours with one another*.

Unassertibility explains away part of the messiness of rejection. The reasons for a speaker to not assert a statement are numerous. They can refuse to assert a statement because they don't want to commit to it, because they are not in the right epistemic position with respect to its content. But they could also refuse to make an assertion that would violate other norms of conversation (Grice, 1975). Maybe this assertion would be impolite, unnecessarily verbose, or straight up irrelevant. However, the reasons for a speaker to deem a statement *unassertible* in virtue of its content run deeper than that. For a statement to be unassertible, its assertion has to constitute a potential *abuse* of the procedure of assertion – i.e. to violate the essential norm for assertion (Austin, 1975). In this sense, the norm for rejection should cover cases of unassertibility under the norm for assertion: cases where, if the speaker asserted the statement they reject, they would be abusing the procedure of assertion.

(6) No, it is not raining.

(7) There might not be a seminar tomorrow.

Rejecting a statement indicates that a speaker deems it unassertible. I explain this in terms of commitments. When the speaker utters the rejections in (6), (7), they express a refusal to commit to a sentence (Brandom, 1998). This explicit refusal to commit is a commitment to not committing. It constitutes a dual move to assertion.

In a commitment picture of assertion, when a speaker asserts P , they commit to P (Brandom, 1983). This commitment is double. First, the speaker can be *held responsible* for their assertion of P . If challenged, they should provide a justificatory procedure highlighting their reasons to stand by the content of P . Second, the speaker issues a *warrant* for participants in the conversation to use P in their deductions (McKinnon, 2015). They introduce P as a statement they hold true. They also license their audience to *repeat* P , and they set themselves as a guarantor for P , where the audience can refer back to them for justification of P .

When the speaker rejects a statement, they refuse to be held liable for the statement they reject, they refuse to be required to provide a justificatory procedure for the statement they reject, and they refuse to warrant its use as an asserted premise in deduction. This refusal to commit can have different grounds. In (6), the speaker refuses to commit to the statement *It is raining* because they are willing to commit to its negation. In (7), the speaker refuses to commit to the statement *There will be a seminar tomorrow* on grounds of ignorance. What unifies these examples is that the speaker deems unassertible the sentence they refuse to commit to.

I can now articulate speaker commitments, unassertibility, and norms of assertion. If assertion is a move governed by a constitutive norm, a speaker is

committed to respecting this norm. When making sincere assertions, they are also committing to their assertions respecting the norm of assertion. Since the purported norms of assertion are epistemic, the speaker thus commits to being in a certain epistemic relation with the content of their assertion. For example, they commit to the sentence they assert being true, or to knowing it, etc. Conversely, when a speaker rejects *P*, they explicitly refuse to commit to it. They commit to not committing to *P*. Doing this, they commit to *P* being unassertible. Since the norms for speech acts underlay the speaker's commitment structure, that means that these norms rule out the assertion of *P*. So if there is a norm for rejection, it compels the speaker to reject *P* only when *P* is unassertible.

An example. Marianne says *No, Paul is not coming to the party*. She rejects the potential statement *P: Paul is coming to the party*. From this rejection follows a commitment. Marianne commits to *P* being unassertible. If the Truth Norm is the right norm for assertion, then unassertible statements are false. So Marianne should only reject *P* if it is false. Otherwise, she would represent as unassertible a statement that is, in fact, assertible under the Truth Norm of assertion.⁵

Even in conversational situations where the impact of assertion and rejection are skewed, it does not make sense for a participant to use different epistemic standards for assertion and rejection. Take Dutilh Novaes' second Prover-Skeptic scenario (2020, p. 55sq). The speaker wants her assertions to go through no matter what (as this counts as a win for her). The audience, by contrast, wants 'good' assertions to go through and 'bad' ones to be rejected. In such situations, it might make sense for the speaker to defend a weak epistemic norm for assertions: where she gets to assert a statement that she only believes; and a strong epistemic norm for rejection: where the audience can reject a statement only if it does not reach a low epistemic standard. But crucially: a low epistemic standard for assertion (only assert what you believe, for example), corresponds to a high epistemic standard for rejection (only reject what you don't believe), which leads to accepting a lot of statements. On the other hand, a lower epistemic standard for rejection (reject what you don't know) would enable the audience to reject *more* statements, which would not be beneficial for the speaker.

- (8) Suppose a scenario where Speaker wins when her assertion is accepted. Audience wants to reject ungrounded statements. Speaker gets to set the norm for assertion and rejection. If Speaker chooses a Belief Norm for both assertion and rejection, this allows her to make a lot of assertions (anything she believes) while Audience can only reject few statements (things they do not believe).

The Belief Norm, which is the weakest for assertion, is also the strongest

⁵This does not exclude that speakers may refuse to assert – and explicitly so – statements that they deem assertible. If Marianne wanted to express that she does not wish to assert *P* for other reasons than *P* being unassertible, she should use another discourse move than rejection. If it were, for example, impolite to assert *P*, she could change the topic, hedge, etc.

for rejection. The Knowledge Norm, which is the strongest for assertion, is also the weakest for rejection. A speaker who is inclined to make many assertions, while few can be rejected, is still better off using the same norm for assertion and rejection.

So unless a norm theorist wishes the norm for rejection to be of a different nature (non-epistemic) than the norm for assertion, I have trouble finding possible situations where it would be beneficial for a participant in a conversation to uphold different epistemic standards for assertion and rejection.

The idea behind highlighting the articulation of assertion and rejection as equally primitive, incompatible speech acts, gives a setting in which the normative conditions for assertion and rejection are interdefinable as dual conditions. When a speaker asserts P , they are not in a position to reject P . When they reject P , they are not in a position to assert P . This is possible only if the rules for assertion and rejection are dual to each other.

Norms of Rejection.

For $\neg\varphi$ a condition that opposes the corresponding condition for assertion,

One must: [reject] P only if $\neg\varphi P$.

In the same way as the norm for assertion, the dual norm for rejection does not compel the speaker to reject anything. It states the conditions under which a speaker may reject something. The Truth Norm does not compel the speaker to assert every true statement. The Truth Norm for rejection does not compel a speaker to reject every non-true statement.

With this blueprint to characterise rejection, we can now answer the question: which norm accounts best for rejection? In particular, how do the different norms for assertion deal with weak rejection?

2.4 Norms of Rejection

Following normative duality, I build a norm for rejection corresponding to each candidate norm of assertion. I assess how they account for weak rejections, and evaluate them.

2.4.1 The Truth Norm.

The Truth Norm account posits as a rule for assertion that S can assert P only if P is true (Weiner, 2005). Its proponents put forward as an argument for the Truth Norm of assertion that a speaker who asserts falsehoods appears blameable. It is pretty straightforward to state a converse norm for rejection:

The Truth Norm of Rejection.

One must: reject P only if P is not true.

This norm accounts for strong rejection. Strong rejection is the notion of rejection that Rumfitt (2000) and Smiley (1996), *inter alia* consider: where from the rejection of P , the assertion of $\neg P$ follows.

The introduction and elimination rules for strong rejection ($-$) and assertion of a negation ($+\neg$) are as follows:

$$\frac{-A}{+\neg A} \quad \frac{+\neg A}{-A}$$

These rules result in harmonious rules for negation, conversely to unilateral classical logic (Dummett, 1991; Rumfitt, 2000). In a bivalent picture of truth and classical logic, $\neg P$ is true if and only if P is not. So, if a speaker can reject P only when P is not true, then they reject P only when they would be licensed (and willing) to assert $\neg P$.

The issues for a Truth Norm of rejection start with *weak* rejections, where the speaker might not be willing to assert the negation of P .

- (9) A: Is it the case that Adèle or Bernard will win the election?
 B: No, Charles is also a candidate.

In (9), B is ready to reject *Adèle or Bernard will win the election*, but would not assert *Neither Adèle nor Bernard will win the election*. But in a system such as Rumfitt's (2000), that has a strong notion of rejection, classical truth, and classical negation, *Neither Adèle nor Bernard will win the election* follows from B's rejection.

$$\frac{-\frac{-(A \vee B)}{+\neg(A \vee B)} (+\neg \text{ intro})}{+\neg(A \vee B)} \text{ (De Morgan's laws)}$$

If the proponent of the Truth Norm is to take weak rejection seriously, they would have to say that it doesn't follow from P not being true that $\neg P$ is true (relax bivalence) or use a non-classical logic such that $\neg(A \vee B)$ is not equivalent to $\neg A \wedge \neg B$ (no neg-raising).

The second issue for the Truth Norm is that the Truth Norm for assertion allows to qualify as assertions accidentally true utterances (10).

- (10) Jana, without looking at a clock, confidently utters: *It is five*. The content of her utterance is true, but accidentally so.

Speakers who make accidental true utterances seem to be missing some epistemic component necessary to ground assertions. Similarly, the Truth norm for rejection encounters the question of rejecting statements who *happen* to be false.

2.4.2 The Belief Norm

The Belief Norm corresponds to a weak account of assertion (Bach and Harnish, 1979; Mandelkern and Dorst, 2021). The idea behind the Belief Norm is that we assert a lot of things, some of which aren't true, as soon as we have some credence in them. In a dual fashion, the norm for rejection should be pretty strong: where a speaker only deems unassertible what they don't believe.

The Belief Norm of Rejection.

One must: reject P only if one disbelieves P .

The notion of disbelief in this norm is not trivial. Conversely to the literature on epistemic logic, the literature on agnosticism does not equate disbelief in P to a belief that $\neg P$ (Baltag and Renne, 2016; Draper, 2017). Agnosticism treats disbelief as an attitudinal state of suspension of belief. So, it holds possible to disbelieve both a statement and its negation, leading to a state of *epochè*, or suspension of judgment. The Belief Norm of rejection uses a similar notion of disbelief, where disbelieving P corresponds to not believing P . This corresponds to a characterisation of rejection as a refusal to commit to a statement. A speaker might reject a statement while they do not believe the negation of that statement.

Weak rejection is *built in* the Belief Norm for rejection. The stronger aspect of rejection follows naturally. When a speaker asserts $\neg P$, they also reject P . Following a Belief Norm for assertion, the speaker asserts $\neg P$ only when they believe $\neg P$. Under some plausible restrictions for rationality, someone cannot believe both P and $\neg P$. So, if the speaker asserts $\neg P$, they are also warranted in rejecting P .

2.4.3 The Knowledge Norm

The Knowledge Norm of assertion aims to distinguish assertions from a whole lot of other *sayings* (Cappelen, 2011). A weak norm of assertion, such as the Belief Norm of assertions, treats as assertions utterances that the speaker only has some credence in (Mandelkern and Dorst, 2021). However, weak assertions appear infelicitous when the speaker does not make their guesswork apparent. Compare (11i) to (11ii):

- (11) John has a terrible work-life balance. His slack status is usually *offline* only when he is sleeping. But, once in a blue moon, he takes a week-end away from his phone. On Saturday at 10h, his status is *offline*.
- i. # John is sleeping.
 - ii. ? John is sleeping. After all, he took a city trip last weekend.

In fact, in contexts of guessing or conjecture, where speakers have some credence in the content of their utterances, *hedged assertions* such as *I think that P*, *I believe that P* are often more felicitous. Accounts of hedged assertions interpret hedges as a way to weaken speakers commitments (Benton and van Elswyk, 2020) or to express a type of knowledge that is *probabilistic* (Moss, 2018).

To separate assertion from guesses, DeRose (2002) and Williamson (1996) propose that the norm of assertion is knowledge instead:

The Knowledge Norm of Assertion.

One must: assert P only if one knows P .

In Williamson's account, knowledge is a primitive mental state, world-involved in that one can only know what is true, but unexplainable in terms of truth, belief, or justification (Williamson, 2002, p. 4). However, from knowledge does follow justification and certainty. Grounding the norm of assertion in this notion of knowledge addresses the fact that speakers who do not properly know the content of their assertions, but merely believe it with high level of credence, even if they turn out true, still appear blameworthy (*lottery cases*). A dual norm for rejection is:

The Knowledge Norm of Rejection.

One must: reject P only if one does not know P .

A Knowledge Norm for rejection provides a pretty low epistemic standard to reject a statement: a speaker can reject any statement that they do not know. That includes any statement they are not entirely certain of. This norm then covers both weak and strong rejection. In cases such as (12), B rejects the statement *There is a seminar tomorrow* on the basis of their uncertainty.

(12) A: There is a seminar tomorrow.

B: No, there might not be. I haven't heard from my supervisor.

In the case of strong rejection (13), we infer from the speaker asserting $\neg P$ that they know $\neg P$. If they know $\neg P$, given that according to Williamson, they can only know what is true, they do not know P . Because they do not know P , they also perform a rejection.

(13) No, sharks are not mammals.

The Knowledge Norm, for assertion and rejection, explains why we can challenge speakers making assertions. When speakers fail to justify their statements and prove knowledge, they can be blamed for their assertions, such that other participants can ask them to retract their utterances (MacFarlane, 2014; Marques, 2018).

2.4.4 The Belief⁺ Norms

Lackey (2007) and McKinnon (2013) claim that it is enough for a speaker to have good justification to make assertions. They propose norms for assertion that are stronger than a Belief Norm, weaker than a Knowledge Norm, focusing instead on plausibility, or justification. Lackey's Reasonable Belief norm addresses cases where a speaker does not know, or does not believe, the content of their claim, but holds it to be the most plausible. McKinnon's Justified Belief norm covers cases of *Gettiered assertions* (Turri, 2016), where the speaker asserts something they don't properly *know*, but are justified in believing.

The Reasonable to Believe Norm of Assertion.

One must: assert P only if it is reasonable to believe that P .

The Justified Belief Norm of Assertion.

One must: assert P only if one has justified belief that P .

An issue to raise with the Reasonable Belief Norm for assertion is that it is geared specifically towards contexts where a speaker sincerely asserts something they do not believe (14).

- (14) A highschool teacher believes that human beings are descended from Adam and Eve. However, she teaches her students that human beings evolved from an ancestor in common with apes. According to her evidence, the theory of evolution is reasonable to believe, even though she doesn't subscribe to it.

Under the Reasonable Belief Norm, the teacher in (14) is not performing an insincere assertion. She is providing her students with the claim she considers most reasonable to believe.

However, the context in (14) is precisely the type of context where it is expected that a speaker might utter, with quasi-assertoric force, statements that they do not subscribe to. As in the cases of guessing contexts for the Belief Norm (11), the teacher appears markedly more blameable if she asserts *human beings evolved from apes* in a context where participants expect her to speak in her own name, and subscribe to the content of her assertions.

The Justified Belief Norm is presented as a way to account for assertions made with a good justification, but without complete knowledge.

- (15) Marianne heard Paul, a week ago, claim he would not go to the party. Unbeknownst to Marianne, Paul changed his mind. However, Paul's COVID test came out positive, and he needs to isolate.

Marianne says: *Paul will not come to the party.*

In (15), Marianne asserts something that she sincerely believes, and has good reasons to believe. Therefore, she is being sincere, to the best of her knowledge. However, the reasons for which she believes *Paul will not come to the party* have been made invalid by Paul's change of mind. Marianne's assertion is a classic Gettier case (Gettier, 1963; Turri, 2016).

Intuitions differ on whether Marianne's assertion *Paul will not come to the party*, is blameworthy. While the claim of sincerity is a strong one, that underlies our intuitions about responsibility for the content of assertions, I think it loses its strength when we interrogate responsibility for one's assertions under the light of *retractions* (Caponetto, 2020; MacFarlane, 2014; Marques, 2018). A participant in the conversation, who is aware that Paul had changed his mind about coming to the party, could ask Marianne to retract her statement, as she didn't have complete evidence for it. And Marianne wouldn't be *at fault* for making an assertion on good enough evidence, but she would still be responsible to acknowledge that her evidence was incomplete.

Despite these issues, converse norms can be established for rejection:

The Reasonable to Believe Norm of Rejection.

One must: reject P only if it is not reasonable to believe P .

The Justified Belief Norm of Rejection.

One must: reject P only if one does not have justified belief that P .

Both Belief⁺ Norms of Rejection account for rejection, given a similar understanding of disbelief as in the case of the Belief Norm of Rejection (section 2.4.2). If, for example, not having a justified belief that P is enough to have a justified belief that $\neg P$, then it would follow from a rejection that the speaker is in the epistemic position to assert $\neg P$. Under an agnostic assumption, not having reasonable belief that P does not provide reasons to believe that $\neg P$. On the other hand, a reasonable belief that $\neg P$, supposing an agent is rational, provides reasons to disbelieve that P .

This first survey of the candidate norms for rejection, based on normative duality, shows that epistemic norms account for weak rejection under fairly uncontroversial assumptions. The assumptions that cover weak rejection under epistemic norms are that disbelieving P does not amount to believing $\neg P$, or that not knowing P is not equivalent to knowing $\neg P$.

On the other hand, a Truth Norm for rejection, if it is to cover weak rejection, poses significant challenges. If a speaker may reject P only when P is not true, then, with bivalence and normative duality, a speaker may reject P only when they are in a position to assert $\neg P$. Accounting for weak rejections, where speakers reject P when they are not in a position to assert $\neg P$, requires relaxing

bivalence, or renouncing classical logic for negation. This is not an implausible move. Intuitionist, paraconsistent and multivalent logics readily claim that, from the falsity of $\neg P$, the truth of P does not follow (Dummett, 1991; Malinowski, 1993; Priest, 2006). However, it is a significant move; and given the other concerns surrounding the Truth Norm of assertion, I elect to focus the rest of this paper on the different epistemic norms for rejection: the *Belief Norm*, the *Justified Belief* and *Reasonable Belief* Norms, and the *Knowledge Norm*.

2.5 Discussion

In this section, I discuss further the candidate norms for rejection, and give arguments that support a weak norm of rejection, the Knowledge Norm. A weak norm of rejection covers cases that are intuitively acceptable, but excluded by the Belief and Belief⁺ norms of rejection. A weak norm of rejection also carries non-ideal consequences: it licenses rejection in a lot of cases. However, I argue, this is not really a problem for the norm of rejection, and more of a problem of explaining why speakers don't do everything they can technically do.

2.5.1 Which norm for rejection?

To evaluate norms for rejection, I analyse whether they predict the rejections we find acceptable. The examples in this section concern speakers' intuitions about acceptability. These examples, while abundantly discussed within the debate surrounding norms of assertion, are still *armchair philosophy*. It would be most valuable to run a study to assess, empirically, what speakers find acceptable to reject, following Turri's seminal experimental work on assertion (2013). Specifically, the hypothesis is that rejecting statements one believes or holds reasonable to believe is acceptable and sincere, while rejecting statements one knows is insincere.

The first epistemic norm of rejection, the Belief norm, states that a speaker should reject a statement only if they disbelieve it. But speakers routinely reject statements that they do believe.

- (16) Suppose that B knows that John's favourite food is Indian, followed by Chinese, and then by Italian. B thinks that John has a 55 % chance of bringing Indian, 20 % chance of bringing Chinese, 15 % chance of bringing Italian. So B believes that John will bring Indian food. However, B does not know it.

A: John will bring Indian. After all, that's his favourite.

B: No, he might bring Chinese or Italian.

The dialogue in (16), building on Mandelkern and Dorst's example (2021), is acceptable. However, in this dialogue, B believes the statement that they are rejecting. Under a Belief Norm of rejection, this would not be acceptable; the speaker would appear to, somehow, break a norm.

The Reasonable Belief, or Justified Belief norms of rejection lower the epistemic standard for rejection: where a speaker may reject a statement only if they do not have a reasonable, or justified belief in it (but they could still believe it). The lottery case in (17) provides an example of rejection of a statement that the speaker thinks is reasonable to believe, and that she is reasonable in believing: that her ticket is not the winning ticket.

- (17) There are 1,000 tickets. There is only one winning ticket. Jane has one ticket. She reasonably believes that her ticket is not the winning ticket. However, she does not know it. So she might reasonably reject the sentence *My ticket is not the winning ticket*.

Jane: My ticket might still be the winning ticket.

The examples generalise outside of the edge cases: we often reject claims that we deem reasonable to believe, and justified, on the basis that we don't *know* them (Williamson, 2002).

By contrast, there is something very wrong about a speaker rejecting something they know.

- (18) B knows that Charles is dropping out of the election. The only remaining candidates are Adèle and Bernard.

A: Adèle or Bernard will win the election.

B: No! Charles is also a candidate.

B's rejection in (18) appears insincere.

Its acceptability improves somewhat in a context where it is understood that the speaker cannot disclose personal knowledge. For example if B in (18) is a campaign manager under a nondisclosure agreement. More generally, this interpretation applies to test cases, where one of the participants explicitly pretends that their knowledge aligns entirely with what has been said in the conversation.

These different examples support the Knowledge norm of rejection. The Knowledge Norm of rejection is the weakest epistemic norm of rejection. After all, if I don't believe P , or don't think that it is reasonable to believe P , then I also do not know P . On the other hand, it also licenses rejecting statements that the speaker does believe, or holds reasonable to believe, on the grounds that they do not have knowledge.

2.5.2 Consequences of the Knowledge norm of rejection

The Knowledge Norm for rejection has drawbacks. However, these consequences are, to a certain extent, shared by any normative approach of speech acts. Identifying a necessary condition for a speech act raises the issue that, in principle, speakers could perform the speech act in question on any content that meets the condition. This is a known issue of norm-based approaches to assertion. A necessary condition only provides circumstances in which one can, or cannot, perform the speech act. It does not tell you what is a *good* way to perform the speech act, except from a performance that respects the rules (Maitra, 2011).

The skeptic problem. In the case of the Knowledge Norm of assertion, the issue is not too pressing: a speaker could in principle assert anything they know, so what? Except from breaking a quantity maxim (Grice, 1975) and being an insufferable participant, they do not threaten the conversation. However, with the Knowledge Norm of rejection, a speaker could in principle reject anything they don't know. This seems much more of a problem.

Following Williamson, we suppose that knowledge is a primitive epistemic attitude which is not explainable in terms of justification and belief. Which means that, as with Achilles and the Tortoise (Carroll, 1895), we could well encounter a speaker who doesn't know (or pretends not to know) anything. They would be licensed to reject every and any statement, and never update their publicised beliefs within the conversation.

The skeptic that Hume decries is one such character (Hume, 1748). With his figure of the skeptic, Hume criticises specifically the late Antiquity Pyrrhonist school of skepticism, best known from Sextus Empiricus works (see Mates' 1996 edition). Skeptic scholars provided rhetorical and logical arguments to allow their students to doubt any statement. Their goal was to show that any knowledge may be doubted. By applying skeptic arguments to the assertion and rejection of a statement, the skeptic could disprove both the truth and falsity of a statement. They would then arrive at a state of *epochè*, suspension of belief, where they can refute any sort of epistemic or doxastic statement.

The skeptic, then, who does not know anything, would not break the norm of rejection by rejecting everything that they do not know. Under a Knowledge Norm of rejection, they can reject any statement.

A humane solution. The reply I have to such a worry might not be very satisfying. But ultimately, the problem with the skeptic is *not* really a problem of speech acts. This does not mean it is not worth addressing within a theory of constitutional norms. In fact, when Maitra (2011) criticises norm-based views of assertion, on the basis that they can't teach anyone what is a good assertion in a language game, she nails the issue on the head. The skeptic problem for the Knowledge Norm of rejection is grounded in epistemology: with impossibly

high standards for knowledge, the skeptic refuses to come to know something, no matter how many arguments their interlocutor provides. In the same way as norms of speech acts are not *sufficient* to know how to play the language games, a norm of rejection is not sufficient to bar malicious or uncooperative rejections. And I argue that it is not its job either.

The skeptic *can* reject everything they don't know, and never update their knowledge in order to keep everything in doubt. But they do not *do* so. They do not reject everything they do not know in conversation; they do not refuse all possible updates of their knowledge. How the skeptic still comes to know things, or act, is an explanation that epistemology, or philosophy of action has to provide. The point is that total *epochè* is unsustainable. For much the same reason that Hume highlights: were the skeptic to doubt everything, no action or conversation would be possible.

The great subverter of Pyrrhonism or the excessive principles of scepticism is action, and employment, and the occupations of common life. [...] a Pyrrhonian cannot expect, that his philosophy will have any constant influence on the mind: or if it had, that its influence would be beneficial to society. On the contrary, he must acknowledge, if he will acknowledge anything, that all human life must perish, were his principles universally and steadily to prevail. All discourse, all action would immediately cease; and men remain in a total lethargy, till the necessities of nature, unsatisfied, put an end to their miserable existence. (Hume, 1748, pp. 126–128)⁶

I don't think it is a matter of conversational norms to prevent a speaker to reject everything they do not know; it is more a matter of epistemology to show that refusing to come to know anything is unsustainable. It is also worth noting that *epochè* is not nearly as threatening to conversation as we could think. The skeptic doubt, or absence of knowledge, combined with a Knowledge Norm of rejection, would, at most, license the skeptic to perform *weak rejections*. By doing weak rejections, the skeptic merely *blocks* possible updates to the conversation; they are not removing possibilities from the conversation, but keeping them. So they do not bring the conversation to an absurd or contradictory state. Worst case scenario, the skeptic stalls a conversation by weakly rejecting every statement, and everyone walks out frustrated.

2.6 Conclusion

Norm-based approaches to speech acts seek to uniquely characterise speech acts with a constitutive condition. To 'treat assertion and rejection as distinct activ-

⁶Ironically, Hume here lands very close to the rich tradition of the skeptic school in finding criteria for action that are distinct from knowledge.

ities on all fours with one another' requires to define the constitutive norms of both assertion and rejection.

In this paper, I start from the notion that assertion and rejection are *incompatible* activities to understand their complementarity. When a speaker asserts a statement, they undertake certain commitments, so that they are not in the position to reject it. Similarly, when a speaker rejects a statement, it is understood that they cannot assert it. The incompatibility of assertion and rejection allows a dual definition of the rules of assertion and rejection (section 2.3). From this perspective, I establish candidate norms for rejection that mirror norms of assertion (section 2.4). A closer discussion of plausible norms for rejection, that takes into account speakers use of weak rejection, favours a Knowledge Norm for assertion and rejection (section 2.5).

This analysis of the conditions of rejection sheds some light on the debate surrounding the norms of assertion. Given that the conditions for rejection are the conditions under which a statement is unassertible, examining what we take to be acceptable rejections also brings some clarity to the underlying criteria for assertion. While speakers often assert things they do not know, they can be challenged upon their statements. Noting that speaker can reject statements that they hold reasonable to believe, but do not know, bring some more support to the Knowledge Norm of assertion. In everyday conversations, speakers flout the Knowledge Norm and take belief to be 'good enough'. However, it is natural for audiences to examine and challenge speakers epistemic grounds for assertions. And the only situation in which a speaker cannot sincerely reject a statement is when they know it.

Chapter 3

Yes, but No: a pragmatic picture of *but*'s contribution

Speakers use the adversative marker *but* in several ways, to indicate contrast. The type of contrast *but* contributes differs depending on context (section 3.1). Because of this diverse empirical profile, multiple accounts of the adversative marker attempt to give a unifying account of its contribution (section 3.2). To do so, these approaches take one use of *but* as primitive. Based on this one primitive use, they characterise the type of contrast that they take *but* to contribute to a sentence. Then, they rely on *pragmatic processing* to fit all uses of *but* under the type of contrast they have characterised.

The problem with this strategy is that extant accounts of the adversative marker use pragmatic processing without restrictions, to fit the diversity of uses of *but* into their one paradigmatic examples. In addition to being *ad hoc*, this appeal to unrestricted pragmatic processing *overgenerates*: by allowing for loosening the chosen contrast condition, it makes virtually *any but*-sentence acceptable (section 3.3).

In this chapter, I give a principled explanation of the pragmatic processing speakers use when assessing *but*-sentences (section 3.4). My approach is agnostic with respect to extant accounts of the adversative marker. It is formulated in terms of *indirect speech acts*, that correspond to attitudes of speakers. Its goal is to *complement* extant theories of the adversative marker by showing *how* speakers pragmatically recover relevant contrasts from the context to assess the felicity of *but*-sentences.

3.1 Empirical profile of the Adversative Marker *but*

Since *but* is an additive marker, it conjoins two propositions. As an adversative, it conveys a contrast between the conjuncts. But how does it contribute such a

contrast?

Early discussions situate the specific contrastive contribution of *but* in between pragmatics and semantics. Frege makes it to be a *colouring* (Frege, 1956, p. 296; Sander, 2019), part of the meaning of the word but not its truth conditions. Grice makes it to be a conventional implicature (Grice, 1961, p. 127; Grice, 1975; Potts, 2005), a lexicalized pragmatic inference.¹ *But*'s projection behaviour is in line with these early findings, that situates its contribution between pragmatics and semantics.

3.1.1 *But*'s projection behaviour

Characterising *but*'s projection behaviour provides insight on the relation of its contribution to the discourse context. It shows the conditions that make a *but*-sentence felicitous and the type of inference *but* triggers. To characterise *but*'s projection behaviour, I apply Tonhauser et al. three diagnostic tests for Strong Contextual Felicity, for Projection, and for Obligatory Local Effect. These are empirical tests devised to quickly categorise the projection behaviour of lexical items contents.

I test the assumption that *but* triggers an inference that the two conjuncts stand in some sort of contrast. In what follow, *but* is the *trigger item*: the lexical item that triggers an inference to a projective content. The *projective content* is that *there is some sort of contrast between the conjuncts*.

The results of Tonhauser et al. diagnostic tests characterise the projection behaviour of the contrast contributed by *but*: it obeys a *weak felicity constraint*. That means: the existence of some sort of contrast between the conjuncts must not be ruled out by the context for *but* to be felicitous.

Strong Contextual Felicity. The first diagnostic tests regards contextual felicity requirements. Some items associated with a projective content can only be used in a context that already entails the content. This is a contextual felicity constraint. To test for strong contextual felicity, we assess whether an utterance of the trigger item is felicitous in a context that does not entail the projective content.

For example, in (19b), the use of the trigger item *too* in *Our bus driver is eating empanadas too* implicates that *someone else is eating empanadas* (projective content).

- (19) a. Context: [Malena is eating her lunch, a hamburger, on the bus going to town. A woman who she doesn't know sits down next to her and says:]

¹For a deeper study of the distinction between *colourings* and *conventional implicatures*, see chapter 7.

- b. # Our bus driver is eating empanadas too. (Tonhauser et al., 2013, p. 78)

The context in (19a) is *neutral* with respect to the projective content *someone else is eating empanadas*. It does not entail or exclude the projective content. In this neutral context, the use of *too* is *infelicitous*. This helps determine that *too* imposes a strong contextual felicity: it can only be used in contexts that entail its projective content.

Conversely to *too*, *but*-sentences are felicitous in a context that is neutral with respect to the existence of a contrast between conjuncts (its projective content). The typical Gricean example (20) can be uttered in a context where no specific beliefs concerning a contrast between poverty and honesty are attributed to the speaker. There is no strong contextual felicity constraint.

- (20) She is poor, but she is honest

But could still impose a weak contextual felicity constraint – which is, be unacceptable in contexts that are negative with respect to the projective content. The test in (21) shows that the *but*-sentence in (21c) is infelicitous or borderline in a context that rules out the most obvious contrasts between conjuncts.

- (21) a. I think Alpacas are fluffy.
 b. To me, being an alpaca and being fluffy are both desirable qualities.
 c. ? Jim is an alpaca, but he is fluffy.

The problem to assess whether *but* imposes a felicity constraint is that speakers are very good at deriving possible contrasts to accommodate *but*'s contribution. While the context in (21) blocks the most obvious contrast, it is near impossible to build a context that blocks all possible contrasts between the conjuncts. So, it is hard to show with such a test that a *but*-sentence is unacceptable in contexts that rule out a contrast between its conjuncts. However, some *but*-sentences are infelicitous in virtue of their conjuncts, such as (22).

- (22) # She's tall, but she is very tall.

In (22), the conjuncts don't stand in contrast according to world knowledge – as being *very tall* is a sub-case of being *tall*. By default, contexts rule out a contrast between being very tall and tall. And, in most contexts², the utterance is unacceptable. If there were no felicity constraint, (22) would be acceptable in most contexts. Therefore, it seems that *but*'s content imposes a weak felicity constraint on the context. *But*-sentences are felicitous only in contexts that do not rule out the existence of a contrast between its conjuncts.

²More on improved acceptability judgments for contrasts that are infelicitous according to world knowledge in section 3.2.

The Family of Sentences diagnostic for Projection. Tonhauser et al.'s second test helps determine whether a content associated to a lexical item has a projection property. If a content is *projective*, it is inferred from the utterance of its trigger even in non-assertoric contexts. The family of sentences diagnostic assesses whether the content follows from the trigger items in variety of sentences without assertoric force.

Take *too* in (23). When a speaker uses *too* in (23), their audience infers that someone else is eating empanadas. In the family of sentences (23), the inference is preserved, as the utterances are unacceptable in context (23a) and acceptable in context (24a). This comparison shows that the contextual felicity constraint that *too* imposes (*the context must verify that someone is eating empanadas*) is preserved in non-assertoric sentences.

- (23) a. Context: [Malena is eating her lunch, A HAMBURGER, on the bus going to town. A woman who she doesn't know sits down next to her and says:]
 b. # Our bus driver is eating empanadas too.
 c. # It's possible that our bus driver is eating empanadas too.
 d. # If our bus driver is eating empanadas too, I am going to leave the bus.
- (24) a. Context: [Malena is eating her lunch, EMPANADAS, on the bus going to town. A woman who she doesn't know sits down next to her and says:]
 b. Our bus driver is eating empanadas too.
 c. It's possible that our bus driver is eating empanadas too.
 d. If our bus driver is eating empanadas too, I am going to leave the bus.

This amounts to the following property:

Projection: A content m of expression t is projective (i.e. has the property of projection) if and only if m is typically implied by utterances of atomic sentences S containing t and may also be implied by utterances of family-of-sentences variants of S . (Tonhauser et al., 2013, p. 82)

To diagnose projection, we verify whether family of sentences variants of an atomic *but*-sentence trigger the inference to a contrast between conjuncts.³ If *but*'s associated content imposes the same felicity constraints in non-assertoric sentences, it is *projective*.

³I apply the version of the test without a strong felicity constraint.

- (25) a. Context: [the speaker's beliefs with respect to a contrast between poverty and honesty are unknown.]
 b. If she is poor, *but* honest, she stayed hungry and did not steal the bread.
- (26) a. Alpacas are fluffy.
 b. I like alpacas and fluffy animals.
 c. ? If Jim is an alpaca, *but* he is fluffy, I want to pet him.

In (25), *but* projects out of conditional antecedents. When *but* occurs in a non-assertoric context – here, a conditional antecedent – the inference that there is a contrast between the conjuncts still holds. However, in a context where the speaker does not believe that there is a contrast between the conjuncts (26), a sentence with a *but* conjunction in the antecedent is infelicitous or borderline. The earlier contextual felicity results are preserved with non-assertoric sentences. Hence, *but*'s content, the inference of a contrast between the conjuncts, is projective.

Obligatory Local Effect. Finally, Tonhauser et al.'s test for obligatory local effect distinguishes projective content according to their embedding behaviour under propositional attitudes verbs. If a projective content has obligatory local effect, then, when it is embedded under a belief predicate, it is part of the content targeted by the belief predicate. If a projective content does not have obligatory local effect, when it is embedded under a belief predicate, it is not part of the content that the belief predicate targets.

The examples in (27) illustrate different behaviours with respect to local effect. Generally, the use of the verbal form *to stop V-ing* carries a *prestate implication: the subject used to V. Stopped smoking* implicates that the subject used to smoke (prestate implication). This is the projective content we test in (27a). Generally, the non-restrictive relative clause, *who is Sue's cousin*, implicates that Bill really is Sue's cousin. This is the projective content we test in (27b).

- (27) a. # Jane believes that Bill *has stopped smoking* and that he has never been a smoker.
 b. Jane believes that Bill, *who is Sue's cousin*, is Sue's brother (Tonhauser et al., 2013, p. 92)

In (27a), the utterance is unacceptable: it assigns contradictory beliefs to Jane. This tells us that the prestate implication of *stop* has a local effect: it assigns to Jane the belief that Bill used to smoke. By contrast, in (27b), the projective content of the relative clause does not have a local effect: its utterance does not assign contradictory beliefs to Jane.

The obligatory local effect condition is thus defined:

Obligatory local effect: A projective content m with trigger t has obligatory local effect if and only if, when t is syntactically embedded in the complement of a belief-predicate B , m necessarily is part of the content that is targeted by, and within the scope of, B . (Tonhauser et al., 2013, p. 93)

Without a strong felicity constraint, the test creates an utterance where the local context is negative with respect to the projective content, and the global context is neutral.

(28) ? Tara believes most clever people are handsome, and that John is clever, but handsome.

If *but*'s content had an obligatory local effect, then the sentence (28) would assign contradictory beliefs to Tara. *But* under a belief predicate would indicate that the belief-holder subscribes to the existence of a contrast between the conjuncts. However, the case in (28) is borderline. Therefore, *but*'s content does not have an obligatory local effect. Instead, it seems that the projective content of *but* needs to be not ruled out by the *global* context of utterance.

Alternative projective content. I also ran the diagnostics with respect to a stricter projective content: the Denial Condition that formal contrast approaches take to be *but*'s contribution (see section 3.2.1). Results are very similar as far as contextual felicity and projection are concerned. The Obligatory Local Effect is clearer than with a general notion of contrast. When *but*'s contribution is understood to be the denial condition, the local context is required to be *m-neutral*:

(29) # Jane believes that Nina is a bassist, and that Tara plays the bass, but Nina plays the guitar.

This further supports my claim that *but*'s content imposes a weak contextual felicity constraint, such that its projective content can be accommodated.

The projection behaviour of *but*. These diagnostics classify the contribution of *but*. *But*'s content does not have a strong felicity constraint (but possibly a weak felicity constraint, with accommodation), it is projective, and it does not have obligatory local effect. This is of interest when trying to characterise the contribution of *but*.

But's content imposes a weak felicity constraint with accommodation on the context. This explains how we assess *but*-sentences: we try to recover a salient contrast from context. If there is no such salient contrast, we rely on world knowledge to infer possible contrasts, relevant to the conversation and compatible with the context. Therefore, acceptability of *but*-sentences can improve given the

right context. That *but*'s content is projective is no surprise. It shows that the part of meaning that *but* contributes is not purely truth-conditional, but also difficult to cancel. Thus, it is not just a conversational implicature. Finally, the obligatory local effect diagnosis is undecided, but points out that the speaker has to assent to the existence of a contrast, even when uttering a belief report. This supports the Fregean and Gricean claim that *but* contributes a contrast at a pragmatic level and introduces its inference as a defeasible entailment in the conversation.

By characterising *but*'s projection behaviour, we understand better how it contributes its contrastive content to sentences. The contrast between conjuncts is a lexicalized inference, that triggers in non-assertoric contexts. However, for *but*-sentences to be felicitous, the contrast *but* contributes needs to not be ruled out by the context. In this sense it is at least somewhat context-dependent.

3.1.2 The multiple uses of *but*

But's variety of uses shows that the contrast it contributes varies depending on its context. *But* may highlight a contrast based on the conversational topic between conjuncts with a parallel structure as in (30). *But* may convey that the second conjunct challenges an expectation raised by the first as in (31) or that the conjuncts provide opposing arguments as in (32). *But* may also locally correct a false proposition as in (33).

Formal contrasts

(30) Natalie plays the drums, but Tara plays the bass.

Denial of expectations

(31) Tara plays the bass, but she cannot read music.

Argumentative contrasts

(32) This ring makes you invisible, but it drives you mad.

Local correction

(33) Natalie doesn't play the piano, but the drums.

But's variety of uses is such that Izutsu (2008) proposes an ambiguity explanation: that *but* actually corresponds to different, homophonic lexical items *but*₁, *but*₂, *but*₃, encoding different types of contrast. Izutsu takes as an argument that the different uses of *but* in English are lexicalized by different markers

in other languages: for example, Spanish, among many languages, has a specialized corrective marker *sino*. However, Jasinskaja and Zeevat (2008) note, the cross-linguistic evidence is not a decisive argument for an ambiguity thesis. They argue that, in languages that have a specialized corrective marker, the specialized marker (*sino* in Spanish) *prevents* the use of the underspecified general adversative (*pero* in Spanish) in corrective contexts. In English, that does not have a specialized corrective, the contrastive meaning of *but* is underspecified, but can be given a unified account.

Two families of approaches, that I recall in section 3.2, aim to account for the diverse profile of *but* in a unified manner: formal contrast approaches and inferential contrast approaches (Jasinskaja, 2012). They each take one of the previously highlighted uses of *but* as paradigmatic. Formal contrast approaches focus on the semantic contrastive structure in (30). Argumentative contrast approaches focus on the pragmatic argumentative structure underlying denials of expectations and argumentative uses of *but* (31, 32).

Both families of approaches appeal to pragmatic processing to account for the non-paradigmatic uses of *but*. As is often the case when pragmatic processing is invoked, there are little explanations as to its limitations. However, some *but*-sentences like (34) are clearly infelicitous except in specific contexts.

(34) # Vivian is tall, but she is taller than Paul.

This means that pragmatic processing is not unrestricted: at least some contrasts are off the table.

The general problem with pragmatic processing is to determine what kind of pragmatic inferences are available to assess the felicity of *but*-sentences. In section 3.3, I present new data to advocate for restrictions on pragmatic processing. Then, in section 3.4, I apply a feature of formal contrast approaches, the denial condition, to characterising pragmatic processing of *but*-sentences instead. In formal contrast approaches, the denial condition semantically encodes *but*'s felicity conditions (section 3.2.1):

Denial Condition 1. The use of *but* is felicitous only if combining the focus alternative introduced in the second conjunct with the background in the first conjunct may result in a false proposition given the conversational context.

In section 3.4, I reinterpret the denial condition as an indirect speech act of rejection. This provides needed restrictions on pragmatic processing while being compatible with both formal contrast and inferential approaches to the meaning of *but*. Explaining meaning contribution with indirect speech acts is in line with an inferential expressivist theory of meaning: according to it, the meaning of various lexical items can be analysed in terms of deductions from speakers' attitudes (Incurvati and Schlöder, 2021). This opens prospects to apply inferential expressivist approaches to meaning outside of their original realm of metaethical or logical discourse.

3.2 Extant accounts for the adversative marker

Because of its diverse use profile, the exact meaning contribution of *but* is difficult to pinpoint.

[...] with it one intimates that what follows is in contrast with what would be expected from what preceded it. (Frege, 1956, p. 296)

[...] what is implied is (very roughly) that there is some contrast
[...] (Grice, 1961, p. 127)

The consensus is that *but* is an additive-adversative marker. As an additive marker, it conjoins two propositions. As a contrastive marker, it contributes a contrast. However, the exact nature of the contrast that *but* contributes is the crux of the discussion between formal contrast approaches (Roberts, 1996; Umbach, 2005) and inferential approaches (Kripke, 2017; Winterstein, 2012). According to the former, the contrastive contribution of *but* can be subsumed by a formal, semantic rule that connects the adversative sentence to the conversational question under discussion (QUD) (Beaver, Roberts, et al., 2017). According to the latter, the contrast that *but* provides is inherently pragmatic, as it depends on the speaker's argumentative goals.

3.2.1 Formal contrast approaches

Formal contrast approaches focus on cases where *but* is used to express a parallel structure and formal contrast between the two conjuncts. They take topic-based contrasts (35) as the prototypical use of *but* (Saebø, 2004; Umbach, 2005).

(35) Natalie plays the drums, but Tara plays the bass.

In topic-based contrasts, the adversative marker *but* conveys a contrast that is *semantic*, given by the parallel structure of the conjuncts. Like *and*, *but* is an additive marker. The use of *but* adds new discourse elements from the second conjunct to sets of topics under discussion. In (35), the second conjunct brings the alternatives *Tara* and *bass* into sets of discourse elements $\{Natalie, Tara\}$ and $\{drums, bass\}$. The sets are given by the parallel structure of the conjuncts, which groups the elements given their role in the utterance (subject, object).

As an adversative, *but* contributes the additional inference that there is a contrast between conjuncts by way of a *denial condition* (Umbach, 2005).

Denial Condition 1. The use of *but* is felicitous only if combining the focus alternative introduced in the second conjunct with the background in the first conjunct may result in a false proposition given the conversational context.

Let's come back to the formal contrast example, and reconstruct the Question Under Discussion (QUD) its conjuncts partially answer.

- (36) a. Which instruments do the band members play?
 b. *Natalie* plays the DRUMS, but *Tara* plays the BASS.

The formal contrast analysis of the example appeals to the focus vs. background structure of the conjuncts. In (36b), the backgrounds give the alternative pair {*Natalie*; *Tara*}. These are the band members, we already know them. The foci are the pair {drums; bass}. These are the new contributions of each conjunct, that give partial answer to the question under discussion. Umbach's denial condition states that the use of *but* in (36) is felicitous only if combining the focus of the second conjunct, *bass*, with the background of the first conjunct, *Natalie*, can result in a false proposition within the conversational context. In other words: if the conversation, or world knowledge, is such that the proposition *Natalie plays the bass* is true, the use of *but* comes out as infelicitous.

Problems for formal contrast approaches. Formal contrast approaches aim to evacuate context-sensitivity out of the analysis of *but*. Thus, they do not interpret the contrast that *but* contributes as *given* by the context. Winterstein (2012) argues that, while a formal contrast approach relying on the denial condition correctly predicts that (37) is incorrect, it does not account for the improved acceptability of (38).

- (37) # Vivian is tall, but she is taller than Paul.
- (38) a. CONTEXT: Paul plays Prusias in *Nicomède*. To play Arsinoé, his evil wife, we need a tall woman, but no taller than the actor for Prusias.
 b. Vivian is tall, but she is taller than Paul.

The formal contrast approach predicts the infelicity of (37) in contexts where the question under discussion is *How tall is Vivian?* For (37) to be felicitous, the denial condition would require that *being taller than Paul* somehow negates *Vivian is tall*. But given world knowledge, there is no extant contrast between *being tall* and *being taller than Paul*.

The context in (38) brings about another possible contrast. But this possible contrast does not overrule the fact that *being taller than Paul* is no contradiction with Vivian being tall. To predict the improvement in acceptability in (38), formal contrast approaches have to take context and salience into account *via* pragmatic processing. According to Winterstein (2012), this is an argument for inferential approaches. If *but* contributes an argument-counterargument structure, the context in (38) supplies a new question under discussion: *Who should we hire?* In this case, the *but*-sentence in (38b) answers the question *Should we hire Vivian?* with the first conjunct arguing for the affirmative, the second for the negative. So the salience of context is built-in for approaches that take the contrast that *but* contributes to be inferential.

3.2.2 Inferential contrast approaches

Inferential contrast approaches focus on denial of expectation or argumentative uses of *but* (31,32) (Kripke, 2017; Winterstein, 2012). In that, they come back to a pragmatic Fregean root:

[...] with it one intimates that what follows is in contrast with what would be expected from what preceded it. (Frege, 1956, p. 296)

For argumentative approaches, *but*, in a sentence, indicates that the first conjunct argues *for* a salient pivotal inference, and the second conjunct argues *against* it. In (39), the pivotal inference of the *but*-sentence is *Putting on the ring*. The first conjunct argues for putting on the ring; the second conjunct argues against.

(39) The ring makes you invisible, but it drives you mad.

Problems for inferential contrast approaches. To derive topic-based contrast uses of *but* (40), argumentative approaches enrich the pivotal inference. In (40), it is difficult to reconstruct a salient pivotal inference such that the first conjunct argues for it, and the second conjunct argues against it.

(40) Natalie plays the drums, but Tara plays the bass.

Inferential approaches' purported solution is to pragmatically enrich the pivotal inference, and conjuncts, in formal contrast cases. For (40), the enrichment is that the pivotal inference would be: *Everybody in the domain plays the drums*. In turn, the first conjunct, *Natalie plays the drums* argues for it. Someone in the domain plays the drums. The second conjunct *Tara plays the bass*, triggers a *only* implicature. If Tara played both the bass and the drums, it would have been strictly more informative to utter *Tara also plays the bass* or some such (Grice, 1975). Given that the speaker did not say so, the inference is that Tara *only* plays the bass. The enriched second conjunct then argues against *Everybody in the domain plays the drums*.

However, this way of enriching formal contrast uses of *but* seems quite artificial, and brings new problems. With unrestricted pragmatic processing, the question becomes: why would a speaker exclude (41) out of the blue?

(41) Vivian is tall, but she is taller than Paul.

After all, it may well be that the first conjunct can be pragmatically enriched such that *tall* implicates *tall, but not very tall*. The appeal to pragmatic processing runs the risk of overgenerating.

Moreover, accounting for corrective uses of *but* proves difficult in inferential contrast approaches. In (42), the enrichment has to be done completely differently than in (40).

(42) Natalie doesn't play the piano, but the drums.

The inferential story for (42) would be that the pivotal inference is *Natalie plays an instrument*, where *Natalie doesn't play the piano* is an argument against her playing an instrument, and *Natalie plays the drums* is an argument for it. This is far-fetched.

3.2.3 QUD approach

To answer the challenges to inferential approaches while keeping their use of context, Jasinskaja (2012) argues for a third type of approach. She develops a question-based approach to the contribution of *but*. At heart, it is a formal approach focusing on the semantic structure of *but*. However, with the use of QUD, it integrates conversational context.

In QUD approaches, conjunctions give distinct answers to *wh*-questions: *Yvette is pretty, and she is clever* gives two different, compatible, answers to the question *What are Yvette's properties?*. *But* differs from *and* in that it conjoins a negative and a positive proposition in corrective uses. Jasinskaja formalizes this polarity as a requirement that a *but*-sentence answers a *wh-whether* question, where one variable is a normal *wh*-question, the other is a *yes-no* variable.

(43) Natalie doesn't play the piano, but the drums.

In this analysis, (43) answers the QUD *What instruments does / doesn't Natalie play?*. The first conjunct has a negative polarity: the instrument that Natalie does not play is the piano. The second conjunct has positive polarity: the instrument that Natalie does play are the drums.

Jasinskaja's account straightforwardly covers argumentative uses of *but*. The *wh-whether* question that (44) answers is *Should you (not) put on the ring?*.

(44) The ring makes you invisible, but it drives you mad.

Problems for the QUD account. In topic-based contrasts, the *yes-no* polarity is difficult to recover. Jasinskaja acknowledges that one could use pragmatic processing for such cases, in the same way as with inferential approaches. By pragmatically enriching the QUD and conjuncts, it is possible to recover a question such that the conjuncts provide a *yes-no* answer. However, this appeal to unrestricted pragmatic processing encounters the same problems as the use of pragmatic processing in inferential approaches. So, instead, Jasinskaja favours introducing hybrid cases where *but* answers a double *wh*-question.

(45) Natalie plays the drums, but Tara plays the bass

In this way, (45) answers the double *wh*-question *Who plays which instrument?*

Unfortunately, this solution collapses the specific contribution of *but* with that of a simple additive marker like *and*. Instead of contributing a yes-no polarity, *but* simply serves to coordinate two partial answers to a double *wh*-question. So the problem with introducing hybrid cases as a way to deal with formal contrasts is then to explain why *but*-sentences are not acceptable in all contexts where *and*-sentences are.

3.2.4 Intermediate conclusions

Because of the diverse empirical profile of *but*, all unifying approaches appeal to pragmatic processing or relaxing the denial condition to account for cases that are not their paradigmatic cases.

This appeal to pragmatic processing is sound. In general, when an utterance introduces a presupposition or an implicature that was not previously part of the shared information, speakers are willing to accommodate presuppositions that are not explicitly excluded by the context (Beaver, 2001; Beaver and Krahmer, 2001; Lewis, 1979). Similarly, speakers are willing to consider contrasts that are not excluded by the context to make *but*-sentences felicitous. Pragmatic processing illustrates this willingness.

However, not every context is up for grabs. Having completely unprincipled pragmatic processing, in addition to providing *ad hoc* explanations of acceptability judgments, also runs the risk of overgenerating.

- (46) a. CONTEXT: Paul plays Prusias in *Nicomède*. To play Arsinoé, his evil wife, we need a tall woman, but no taller than the actor for Prusias.
- b. Vivian is tall, but she is taller than Paul.

Without the context in (46a) that introduces a salient question under discussion (*Should we hire Vivian?*), the *but*-sentence in (46b) is infelicitous. This judgment for (46b) out of the blue reflects the fact that some contexts are more readily available to speakers than others. For example, it is easier for a speaker to accommodate a context where *playing the drums* and *playing the bass* form a contrast than a context where *being tall* and *being taller than Paul* form a contrast. So, an appeal to pragmatic processing should be principled enough to explain why not every context is available all the time to evaluate *but*-sentences. Some *but*-sentences are unacceptable out of the blue, and should remain so. In section 3.3, I provide more examples of *but*-sentences that require principled pragmatic processing, either to account for them or exclude them.

My goal, in the rest of this paper, is to refine the understanding of pragmatic processing. To do so, I encode the denial condition *pragmatically*, in terms of indirect speech acts. This is compatible with the three accounts of the adversative marker: formal contrast approaches *à la* Umbach (2005), inferential approaches

as advocated by Winterstein (2012), and QUD approaches as the one Jasinskaja (2012) proposes. A pragmatic encoding of the denial condition provides a way to appeal to context (for any of these approaches) in a principled manner, which gives the intended results for non-paradigmatic cases.

3.3 Challenging data

In this section, I present three features of *but* that challenge the appeal that extant accounts make to pragmatic processing. These features are not addressed in extant accounts of the adversative marker, and they motivate setting restrictions on pragmatic processing.

The first class of features are iterations. Most iterations of *but* are rightfully excluded by the strict versions of *but*'s accounts. However, when enriched with unrestricted pragmatic processing, existing accounts of *but* overgenerate: they accept too many cases, including iterations.

The second class of features are negationless corrections. Negationless corrections show that the contrast that the audience infers from the use of *but* is not only at the level of the sentence. Instead, it bears on the conversational context. This motivates further an account of *but* that integrates contextual cues to analyse the contrast that *but* contributes.

The third class of features concern the ability of *but* to conjoin different speech acts, or perform a topic change. The use of *but* to conjoin different speech acts motivates an analysis at the interface of pragmatics and semantics. The contrast that *but* contributes cannot be only between the semantic contributions of the conjuncts.

These three classes of features, together, motivate my proposal to posit restrictions on pragmatic inferences available to assess *but*-sentences. Moreover, negationless corrections and uses of *but* to conjoin speech acts provide additional evidence that the contrast *but* contributes also operates at the level of speech acts. I use this evidence in my proposal in section 3.4.

3.3.1 Iterations

Most iterations of *but* are deemed unacceptable by speakers.

(47) # Tara plays the bass, but Ramona plays the guitar, but Paul plays the drums.

(48) ? The ring makes you invisible, but it drives you mad, but it is very powerful.

Inferential approaches straightforwardly exclude iterations (Winterstein, 2012). Take *but*'s contribution to be that there is a pivotal inference *P*. The first con-

junct argues for P , the second argues against P . The third conjunct does not have a direction of argumentation.

In formal contrast analyses, the denial condition does not prevent a *but*-sentence from being a conjunct in a *but*-sentence. The denial condition predicts that the use of *but* in (47) to coordinate *Tara plays the bass but Ramona plays the guitar* and *Paul plays the drums* is felicitous only if combining the focused alternative (*Paul* or *drums*, depending on the prosody) with the background of the first conjunct may result in a false proposition. Let there be a situation where: *Paul plays the bass*, *Paul plays the guitar*, *Ramona plays the drums*, *Tara plays the drums* are all propositions ruled out by the context. Then, the denial condition's felicity conditions are fulfilled.

If, on the other hand, we exclude iterations by requiring that the alternatives are pairs and thus don't have a set as an element, the answer is *ad hoc*. In addition, this answer prevents the conjunction of *and*-sentences with *but*, which are usually felicitous (49).

(49) He is rich and happy, but a crook.

Iterations cause a similar problem in a QUD-account. Instead of relying on pragmatic processing, Jasinskaja introduces hybrid cases of use of *but* to account for formal contrast uses. In hybrid cases, one can use *but* without a yes-no polarity to answer a double *wh*-question. But with hybrid cases, the account does not exclude iterations. If the contribution of *but* is to provide pairs of lexical items to answer a double *wh*-question, then nothing excludes that one element of these pairs of items is itself a pair. In this way, the iteration in (47) would answer the double *wh*-question *Who plays which instrument?* with partial answers. The first conjunct, *Natalie plays the drums but Tara plays the bass* would be one such partial answer, and the second conjunct *Paul plays the drums* another partial answer.

The fact that iterations are generally deemed unacceptable seems a counter-argument to the appeal to hybrid cases. More generally, excluding iterations requires restricting the available pragmatic processing.

3.3.2 Negationless corrections

The second challenge to extant analyses of *but* is the case of negationless corrections. When *but* is used for local correction, it articulates a first conjunct where a negation pinpoints a specific problematic element, and a second conjunct that introduces an alternative. The local correction implies that, if we replace the problematic element with the alternative and remove the negation in the first conjunct, we get a true proposition (50).

(50) She didn't go to Paris, but to Milan.

However, some corrections are negationless (51).

- (51) a. I love these crimson curtains.
 b. The curtains are red, but scarlet.

For an argumentative approach that determines the pivotal inference only with respect to the scale of the sentence, it is difficult to determine a pivotal inference such that the first conjunct argues for it, but the second argues against. In (51), the pivotal inference is directly contributed by the context of utterance. The first conjunct argues for *the curtains are crimson*, as crimson is a subtype of red, and the second conjunct argues against *the curtains are crimson*, as scarlet is a distinct shade of red.

Formal contrast approaches in their strict reading predict that *but* introduces an inference that the curtains might not be red. This nonsensical result encourages to expand Umbach's denial condition to integrate more the context of utterance.

The example of negationless corrections shows that both approaches need to take into account the context of *but*-sentences. This is our first clue towards characterising pragmatic processing: the discourse context is what provides salient elements to enrich denial conditions, or pivotal inferences.

3.3.3 Combination with other speech acts

The last phenomenon I draw attention to is that *but* can be used to coordinate distinct embedded speech acts (Krifka, 2014). It can transition between an assertion and a question (52), or an assertion and an exclamation (53).

- (52) She is poor, but is she honest?

- (53) I have to work, but dammit!

This proves problematic for formal contrast theories that consider that the contribution of *but* bears on the semantic role of the conjuncts.

Even considering argumentative approaches, it is relevant to ask *how* the non-assertoric speech act in the second conjunct constitutes an argument with respect to a proposition under discussion.

Take home messages. Extant accounts of the uses of *but* appeal to pragmatic processing to explain non-paradigmatic uses of *but*. However, the examples in this section show that we need pragmatic processing to take into account contextual and pragmatic cues. Pragmatic processing has rules. It operates in a context of utterance that restricts what is up for grabs to make pragmatic inferences. The rest of this paper focuses on the relation between pragmatic processing and context of utterance.

3.4 Denial condition as a speech act

To characterise pragmatic processing, I take as a starting point that formal contrast and inferential contrast approaches agree that *but*'s contribution can be characterised as implying a *denial*. They disagree on how to determine what is denied. Hence, the characterisation of the *denial condition* in formal contrast approaches (section 3.2.1) is relevant to the denial in other uses of *but*. First, I show how to extend the denial condition to include salient propositions under discussion. Then, I propose to understand the denial condition as an implicit speech act of rejection. This picture explains how we enrich inferences to assess the felicity of *but*-sentences. It also restricts pragmatic processing.

3.4.1 Expanding the denial condition

In QUD discourse theories, two questions determine the information an utterance carries. First, an utterance brings information to the conversation by answering contextual questions. Those can be inferred from the topics of conversation, etc. Second, the specific question a sentence answers in function of the focus structure determines its information structure. Umbach calls this second question the *retrospective quæstio*.

The *retrospective quæstio* is the question that clarifies a sentence's background vs. focus structure. The information already implied by the quæstio is the background, while the focus is the new, salient information that the utterance contributes. Prosody is helpful in reconstructing this quæstio, as the prosodic focus usually falls on the new salient piece of information a sentence contributes.

- (54) a. Who had what to drink?
 b. NATALIE had coffee.
 i WHO had coffee?
 c. Jim had TEA.
 ii WHAT did Jim get?

(54a) is the question under discussion. It is a contextual question to which the utterances in (54b, 54c) give partial answers. By contrast, (54i) and (54ii) are the retrospective quæstiones that correspond to the prosody in (54b, 54c). They clarify the information structure of (54b, 54c). In (54b), NATALIE is the focus of the utterance, and *had coffee* the background. *Natalie* is the new information that the utterance contributes, under the assumption that somebody had coffee.

Enriching the Denial Condition. I propose enriching the denial condition by combining the contextual QUD with the background of the retrospective quæstio with a simple conjunction. Hence, the denial condition bears on two types

of questions. One is the implicit question the conversation is trying to solve (conversational QUD) and the other explicates the information structure of a *but*-sentence. In (55), the backgrounds of the discourse QUD (55a) and of the retrospective quaestio (55b.i) combine to form the background *x had y to drink & x had coffee*. We plug in the focused element *Natalie* in place of *x* and recover the information: *Natalie had something to drink & Natalie had coffee*.

- (55) a. Who had what to drink?
 b. NATALIE had coffee.
 i WHO had coffee?

For *but*-sentences, this gives rise to a modified denial condition which applies to the enriched background. Instead of requiring that the focus of the second conjunct, combined with the background of the first conjunct, may result in a false proposition in context, I spell out the role of context by integrating it in the background of the first conjunct. If, in (55), one were to continue (55b) with *but Jeremy had tea*, the idea would be that combining the focus of the second conjunct, *Jeremy*, with the conjunction given by the first conjunct, *had a drink & had coffee* results in a false proposition. Thus, I modify the original denial condition from formal contrast approaches to incorporate a rich background.

Denial Condition 2. The use of *but* to conjoin two propositions is felicitous only if combining the focus of the second conjunct with the *rich* background of the first conjunct results in a false proposition.

Note that this is just a way of explicitly adding in the denial condition the role of context (that it doesn't rule out that the proposition is false).

Applying the extended Denial Condition. We run the extended denial condition on some cases. For formal contrast, take example (56).

- (56) a. What did the children do today?
 b. The SMALL children stayed at HOME,
 i. WHICH CHILDREN stayed at home?
 c. but the BIGGER ones went to the zoo.
 ii. RQ: WHICH CHILDREN went to the zoo?

The background of (56b) is *we know that some children did something & some children stayed at home*. The focus of (56b) tells us that the children who did something (specifically staying home) are the *small children*.

The denial condition states that the use of *but* is felicitous only if combining the focus of (56c), *bigger children* with the background of (56b) may result in a

false proposition. *The bigger children did something (specifically, stayed home)* can be false given the context. It is not explicitly given that the bigger children both spent part of the afternoon at home and went to the zoo.

By now, the reader has noticed that, for most topic-based contrast cases, having an enriched denial condition is going to yield the exact same results as the first denial condition. In topic-based contrasts, the contextual QUD is generally a *who does what* situation, where the *but*-sentence gives two partial answers: *x does a, y does b*, with the added inference *y does not do a*. So felicity results for topic-based contrast uses of *but* are preserved.

The contribution of an extended denial condition is to explain how the acceptability of *but*-sentences (57b) improves when a context provides a different salient contrast (57a).

- (57) a. CONTEXT: Paul plays Prusias in *Nicomède*. To play Arsinoé, his evil wife, we need a tall woman, but no taller than the actor for Prusias. Which actresses can we hire?
- b. Vivian is TALL, but she is TALLER THAN PAUL.

To get the background, I break down the general QUD *Which actresses should we hire?* into a complex double question: *Which actresses are how tall compared to Paul, and thus can be hired?* This corresponds to a contextual background *If an actress has tallness x , we should hire her*.

The *but*-sentence in (57b) provides two candidates *foci* for *tallness x* . The first conjunct, *Vivian is tall*, plugs the focus *tall* into the contextual background. It results in a proposition that is true in the context: *If an actress is tall, we should hire her*.

According to the new denial condition, I combine the background with the focused alternative *taller than Paul* in the second conjunct. However, the context rules out *If an actress is taller than Paul, we should hire her*. This explains why the *but*-sentence in (57) turns out felicitous. When the context does not provide an alternative salient QUD, however, (57b) comes out as infelicitous. In absence of an additional QUD, the contrast it provides is between *being tall* and *being taller than Paul*, which are generally not in contrast.

3.4.2 Denial condition: an indirect rejection

In section 3.4.1, I introduced a broad denial condition. Because the broad denial condition appeals to the context, in addition to the utterance itself, it is at the interface between pragmatics and semantics. I analyse the broader denial condition as performing an indirect speech act of rejection: when using *but*, the speaker indirectly rejects some salient proposition in the context.

Indirect speech acts (Asher and Lascarides, 2001; Searle, 1975) occur when an utterance performs both an *overt* and a *covert* speech acts. Within a discourse, a

speaker may make multiple utterances, that result in performing different speech acts (Asher and Lascarides, 2003; Austin, 1975; Searle, 1969). The speaker in (58) performs a series of wishes.

- (58) I bless the bell of Koningsdam, as an extension of this vessel. I wish upon her fair winds and a flowing sea. I wish safety for her operation, wisdom for her officers, good health to her crew, and comfort for her guests. (Queen Máxima blessing and naming the M.S. Koningsdam during the May 23rd, 2016 dedication ceremony)

But some utterances, in virtue of conventions, may perform a secondary indirect speech act in addition to the overt speech act. When Queen Máxima blesses the ship's bell, she is also thereby naming the ship.

This small detour with indirect speech acts allows to refine the contribution of *but*:

Denial Condition 3. The use of *but* is felicitous only if combining the focus alternative introduced in the second conjunct with the *rich* background of the first conjunct results in a proposition *that the speaker rejects*.

Instead of taking the combination of second focus with first background in terms of *proposition that can be false* in context, I take it as *proposition the speaker indirectly rejects*. Why is it relevant? Because independently motivated features of the speech act of rejection show that a proposition need not be false for the speaker to reject it. Two features of rejection are particularly important here. First, rejection can be weak (see chapter 2). Secondly, rejection can bear on the performance of an utterance (Dickie, 2010; Price, 1983). These features help us cover the data from section 3.3.

3.4.3 Weak Rejection and *but*-sentences

Rejection, expressed by the force marker *No!* expresses the refusal to assert a proposition. In this, it corresponds to an attitude of the speaker, who disagrees with a proposition. Rejecting a proposition – disagreeing with it – does not correspond to asserting its negation (Incurvati and Schlöder, 2017).

- (59) a. A or B will win the election.
b. No! C is also a candidate.

In (59), the second speaker rejects the statement asserted in the first utterance. However, they do not (necessarily) subscribe to its falsity. The speaker in (59) would not be willing to assert the negation of the first statement: *Neither A nor B will win the election*. Speakers use weak rejections to express that they are unwilling to commit to a claim. It is the opposed move to strong assertions, by

which speakers commit to a statement (Brandt, 1983). As such, the essential effect of weak rejections is to block the acceptance of a statement in a conversation (Incurvati and Schlöder, 2017). A subcontrary speech act, weak assertion, blocks accepting the negation of a statement (Incurvati and Schlöder, 2019; see also chapter 4).

Rejection and inferential uses. Weak rejection explains argumentative uses of *but*. Even if we follow a formal contrast story, the combination of the focus of the second conjunct with the background of the first conjunct needs not result in a false proposition. It is enough that the speaker does not wish to assert the proposition. If we adhere instead to an argumentative story, the weak rejection targets the pivotal inference. Argumentative uses of *but* generally do not lead to a clear-cut decision. In (60), the use of *but* indicates that the speaker is not willing to assert *We should buy the ring*.

- (60) a. Should you put the ring on?
 b. The ring makes you invisible, but it drives you mad.

The speech act definition of the denial condition (3.4.2) accounts for argumentative intuitions as to where the contrast lies: the sentence is felicitous because the speaker weakly rejects *You should put the ring on*. That the rejection is weak is consistent with the fact that argumentative uses of *but* rarely give a definite answer to the question under discussion. Instead, they present arguments in both directions that further discourse may elaborate on.⁴

Rejection and iteration. Iterations are most commonly unacceptable. The problem was to explain why. Saying that *but* only coordinates simple sentences sounded *ad hoc*, and would not do: *but* interacts just fine with complex *and*-sentences. A speech act picture, where *but* alternates an indirect assertion and an indirect rejection, provides arguments to exclude *but* iterations. According to the new denial condition, *but* coordinates an assertion and a rejection answering a unique QUD. This does not iterate well. First, coordinating a *but*-sentence with a *but* would amount to perform twice an indirect rejection.

- (61) ? The ring makes you invisible, but it drives you mad, but it is very powerful.

The extended denial condition shows us how to interpret the contribution of the first *but*: *it drives you mad* constitutes an evidence against *You should put the ring on*. It is a reason that the speaker has to reject *You should put the ring on*.

⁴While Asher and Lascarides (2003) originally assume that only the second conjunct is available for elaboration in a contrast structure, Asher and Vieu (2005) refine this account. Asher (2008) revises the hypothesis that contrast always makes the first conjunct unavailable for anaphora. He highlights some of the constraints at stake in elaborating on contrast and parallel discourse structures.

Now, *it is very powerful*, when coordinated with the first two conjuncts, should be interpreted as a reason the speaker weakly rejects the contribution of the first two conjuncts. However, the contribution of the first two conjunct is itself an assertion / rejection structure. Such that the last conjunct should provide a reason both to refuse to put on the ring, and to accept to put on the ring. This complex attitude, for a speaker, cannot be expressed with only one proposition.

Secondly, treating *but* as performing an indirect rejection explains why the contrastive content of a *but*-sentence projects out of embedding (63) (Krifka, 2014). This is also in line with *but*'s projection behaviour (section 3.1.1).

(62) If my damn bike has a flat, I will have to take the tram.

(63) If she is poor but honest, she didn't steal the bread.

In (62) and (63), *damn* and *but* exhibit similar projection behaviour. In (62), even though the expressive *damn* is part of a conditional antecedent, the speaker is still committing to a certain dislike of her bike (Potts, 2005). Similarly, in (63), the speaker appears to commit to a contrast between the qualities *poor* and *honest*. Treating *but* as performing an indirect weak rejection predicts the effect of embedding it in a conditional antecedent: it amounts to embedding a weak rejection in a suppositional speech act. Even in a suppositional context, the speaker is assumed to subscribe to this weak rejection.

3.4.4 Pragmatic rejection and *but*-sentences

The speech act of rejection also bears on the *performance* of speech acts (Price, 1983). A *No!* may oppose an assertion that does not respect conversational norms – because it is impolite, irrelevant, etc. – without questioning its content. This pragmatic notion of rejection explains some uses of *but* for negationless correction and denial of expectation. In (64), the speaker rejects the pivotal inference (that we should hire Tara as a musician) as it is not *relevant* to the conversational goals.

(64) Tara plays the bass, but we need a harpist.

Similarly, when *but* conjoins an assertion and an exclamation, or an assertion and a question, the implicit rejection bears on whether the first assertion can be performed.

(65) I have to work, but dammit!

In (65), the speaker uses the second conjunct – the exclamation – to express her dismay with the first conjunct.

Treating the denial condition as an indirect speech act of rejection, independently motivated features of rejection accounts for acceptability judgments in

more cases. In argumentative uses, *but* conveys a weak rejection: it shows that the first conjunct fails to solve the QUD. The pragmatic notion of denial explains why *but* coordinates different speech acts, or changes topic of conversation.⁵

3.5 Conclusion

Because of its diverse empirical profile, the contribution of the adversative marker *but* is hard to pinpoint. Traditional analyses of this contribution rely on extra pragmatic processing to account for all its uses.

In this chapter, I developed a proposal to analyse the contribution of *but* in context, both semantic and pragmatic. This account, based on the denial condition from formal contrast approaches, explains the role of pragmatic processing in understanding *but*-sentences and shows how pragmatic processing relates to context. Section 3.4.1 expands the denial condition to bear on contextually salient propositions. I integrate the conversational question in the background of an utterance. This connects the focus and background structure of an utterance to the information structure of the conversation where it occurs. We understand how the contrast formalized by the denial condition picks propositions that are salient in the discourse context. This enrichment of the denial condition bridges the gap between formal contrast and inferential contrast approaches.

Both formal contrast approaches and inferential approaches agree that *but* contributes a yes-no polarity. In section 3.4.2, I interpret this polarity as a speech act alternation between an assertion and an indirect rejection. Independently motivated features of the speech act of rejection explain some performative behaviours of *but*. Argumentative uses of *but* rely on weak rejection: the second conjunct brings new information to reject some inference that the first conjunct defeasibly imply. Rejections can also bear on the performance of a speech act: this explains cases where the contrast between conjuncts appear to bear on pragmatic inferences, as when *but* coordinates distinct speech acts, or changes topic of conversation. This proposal explains the problematic cases exposed in section 3.3.

The goal of my proposal is not to compete with extant accounts of the adversative markers *but*. Instead, I aim to complement them with a principled approach to pragmatic processing that is compatible both with a formal contrast and an inferential contrast theory.

More broadly, this account of *but* corresponds to an inferential expressivist take on the contribution of a lexical item traditionally associated with Fregean colourings, or Gricean conventional implicatures. The profile of such items is still puzzling: they trigger inferences that are not purely semantic, but that are still non-defeasible. Taking these inferences to be conventionalised performances of indirect speech acts gives a new direction to their analysis.

⁵In the latter case, the indirect rejection questions the relevance of the first conjunct.

Chapter 4

Perhaps, might and weak assertion: evidence from speakers practices

4.1 Weak assertion

In the study of speech acts (Austin, 1975; Searle, 1975), assertion gets the lion's share (MacFarlane, 2011b; Stalnaker, 1978). It is the speech act with which the speaker can impart new information on her audience. Be it because she knows (Williamson, 1996), believes (Bach and Harnish, 1979), thinks reasonable to believe (Lackey, 2007), or holds true (Weiner, 2005) the content of her assertion, she gives to her statement an assertive force. One of the intended effects of an assertion is that her audience also comes to know, believe, hold reasonable to believe or hold true her statement (Bach and Harnish, 1979).

With great power comes great responsibility. When making an assertion, a speaker endorses the responsibility of justifying the content of her assertion, if she is so asked (Brandt, 1983). She can be held responsible for the content of her assertion. And other participants in the conversation may repeat the content of her assertion, taking her as an authority: she issues a warrant to repeat and use the content of her assertion (McKinnon, 2015). Because of the importance of these commitments, I argued in chapter 2 for a Knowledge Norm of assertion and rejection, that explains why speakers often refuse to assert claims they do not know.

However, speakers often wish to introduce conversational information that they do not know. In such cases, speakers sometimes break purported norms of assertion, which led many theorists to cast doubt upon norm-based accounts of assertion (Cappelen, 2011; MacFarlane, 2011b; Maitra, 2011). Or they perform *hedged assertions*: *I think that P*, *P, if I am not mistaken*, that weaken their commitments (Benton and van Elswyk, 2020) or express probabilistic knowledge (Moss, 2018).

But sometimes, speakers don't even wish to *add* information to the conversation, even with hedged or probabilistic assertions. Sometimes, speakers only wish

to leave possibilities open. *Weak assertions*, according to Incurvati and Schlöder (2019), such as *Perhaps P*, do just that. *Weak assertions* preemptively reject removing a possibility from the set of possibilities compatible with what the participants in the conversation agree upon. They only commit a speaker to a weak rejection of $\neg P$.

(66) Perhaps there is a seminar tomorrow.

When she utters (66), the speaker aims to keep open the possibility that there is a seminar tomorrow. In Common Ground terms (section 2.2.2), weak assertions do not propose an update so much as block an update. Specifically, the weak assertion in (66) does not aim to add to the Common Ground the information that *there is a seminar tomorrow*. Rather, if it is accepted, the result of the weak assertion is to *block* a potential update of the Common Ground that removes worlds where *there is no seminar tomorrow* (fig. 4.1). Thus, Incurvati and Schlöder (2019) take the strong assertion of *might P*: $(+\diamond P)$ and the weak assertion *Perhaps P* $(\oplus P)$ to be inter-deducible. In this picture, *might* is a sentence modifier (in the same way as *not* switches a truth value, *might* changes the epistemic modality of a proposition) corresponding to the force modifier *Perhaps*.

Now, it is all nice and well to postulate a speech act: *weak assertion* – corresponding to a speaker’s attitude: refusing to exclude a possibility. But is there linguistic evidence for it? According to Incurvati and Schlöder, yes! They analyse *perhaps* as a force marker by which the speaker indicates weak assertion; and *might* as a sentence modifier (section 4.2). If Incurvati and Schlöder’s analysis is correct, there are things we expect regarding the linguistic behaviour of *perhaps* and *might* (section 4.2.3). To amass more evidence in support of Incurvati and Schlöder’s analysis, I gathered examples across the *British National Corpus* and the *Corpus of Contemporary American English* (section 4.4).

4.2 Assertion and weak assertion, the theoretical picture

4.2.1 The theoretical argument for weak assertion

In first instance, the existence of weak assertion is *theoretically* justified. Bilateralism, in logic, stems from understanding rejection as a distinct activity from assertion of a negation (section 2.1.1). By doing so, bilateralists build logical systems where premises and conclusions of deductions may be asserted or rejected (Rumfitt, 2000; Smiley, 1996). This is done by understanding the result of assertions and rejections in terms of incompatible information, represented as possible worlds. A strong assertion of *P*, if accepted, *rules out* $\neg P$ -worlds. A strong rejection of *P*, if accepted, *rules out* *P*-worlds. To the Smileian picture,

where assertion and rejection are inter-deducible, Incurvati and Schlöder (2017) add *weak rejections* (section 2.1.1) With a weak rejection of P , speakers *rule out ruling out* $\neg P$ worlds: they block an assertion of P .

But then, if speakers have the possibility to block an assertion of P , they should also be able to block the assertion of $\neg P$. This leads Incurvati and Schlöder to postulate the existence of a speech act that mirrors weak rejection: *weak assertion* (2019). Where a weak rejection ($\ominus P$) blocks the strong assertion of P ($+P$), a weak assertion ($\oplus P$) blocks the strong rejection of P ($-P$). The multilateral rules for contradiction are a consequence of this picture:

$$\frac{+P \quad \ominus P}{\perp} \quad \frac{-P \quad \oplus P}{\perp}$$

Incurvati and Schlöder provide linguistic grounds for adding weak assertion to their theoretical framework. In ‘Die Verneinung’ (1918; *Negation* in (Beaney, 1997, pp. 346–361)), Frege describes asserting P as answering the thought-question *Is it the case that P ?* with *Yes*. Rejecting P is a *No* answer. For Incurvati and Schlöder, there is an equivalent for weak assertion: a *Perhaps* answer. They also argue that this *perhaps* answer has a sentential equivalent: the epistemic modal *might*. In Incurvati and Schlöder’s inferential expressivist framework, the force modifier *Perhaps* has a similar relation to the epistemic modal *might* as the force modifier *No* to the sentence modifier *not*. *Not* is the sentence modifier which, within an assertion of *not P* , indicates that the speaker also rejects P .

Incurvati and Schlöder connect the speech act of *weak assertion* to a speaker’s epistemic state *not ruling out* a possibility. Epistemic compatibility is what a speaker commits to when she weakly asserts P : keeping P open as an option. *Might* is the modal which, within an assertion of *might P* , indicates that the speaker *weakly asserts P* .

$$(+\neg \text{ elim}) \frac{+\neg P}{-P} \quad \frac{+\diamond P}{\oplus P} (+\diamond \text{ elim})$$

Accordingly, in (67), the speaker wouldn’t answer the question *is there a seminar tomorrow?* with either *yes* or *no*. She would say *perhaps*.

(67) Don’t make any plans yet! There might be a seminar tomorrow.

Incurvati and Schlöder provide compelling examples, and open an exciting direction of analysis for the epistemic possibility modal *might*. To strengthen their analysis, I provide more examples from corpora, and link them to projection behaviour that we can expect from indicators of weak assertion. This is a first step in the direction of quantitative testing to assess whether the weak assertion analysis of *perhaps* and *might* is backed by conversational practices.

4.2.2 Weak assertion and Common Ground

A Common Ground picture of conversation focuses on the information compatible with what the participants have agreed upon (Stalnaker, 1978, 1999, 2002). In this picture, the effect of an assertion, if accepted, is an update that removes all worlds incompatible with the content of the assertion from the set of worlds compatible with what the participants have jointly agreed upon (see fig. 2.1 in section 2.2.3).

Assertions are powerful: an assertion of P , if accepted, makes non- P worlds incompatible with what the participants agree upon. But a speaker may well consider that P is possible, while not wanting non- P worlds to be ruled out. If she were asked *Is it the case that P ?*, she wouldn't answer with a resounding *Yes* or *No*. Instead, she would want to answer *Is it the case that P ? Perhaps*. We can model this essential effect of the speech act of *weak assertion*. When a speaker utters *Perhaps P* , she expresses her wish to not remove the possibility that P from what the conversation leaves open (fig. 4.1).




Figure 4.1: Updating the Common Ground with a weak assertion *Perhaps P*

To understand what it means for a speaker to weakly assert P , we can turn to a commitment analysis of the respective effects of strong and weak assertion (Brandom, 1983). When the speaker asserts P , she is doing two things:

- (i) She issues a warrant to his audience, so that they can assert themselves that P , and
- (ii) She undertakes a commitment to P :
 - (i) She commits to behave as if she held P true, and
 - (ii) She takes on the responsibility of providing an adequate justificatory procedure for P if she is challenged.

When a speaker *weakly* asserts P , she expresses a refusal to commit to $\neg P$. She may also issue a challenge to anyone committed to $\neg P$ to justify their commitment to $\neg P$.

4.2.3 The weak assertion account of *perhaps* and *might*

***Perhaps* as a force modifier.** According to Incurvati and Schlöder (2019), weak assertion is expressed in natural language with *perhaps* and *might*. The role of *perhaps* in (68) is to weaken the speaker’s assertion of *it is raining*.

(68) Perhaps it is raining.

Incurvati and Schlöder present evidence that *perhaps* is a force modifier that corresponds to weak assertion:

Our evidence for *perhaps* being a force-modifier is the following: (i) *perhaps* exhibits the embedding behaviour that one would expect of a particle operating exclusively at the speech act level; (ii) the role of *perhaps* cannot be reduced to that of commenting on one’s performance of a speech act; (iii) in polar questions, *perhaps* appears not to modify the core proposition; (iv) in commands, *perhaps* appears to modify force.

Defending an analysis of *perhaps* as a force modifier requires characterising its embedding patterns. I do so, first by theoretically characterising force-modifiers projection patterns in line with other lexical items that operate at the level of speech acts: speaker-oriented adverbs and utterance modifiers (section 4.3.1). Then, in section 4.4.1 and section 4.4.2, I gather examples that I analyse closely to determine its embedding behaviour.

***Might* as an epistemic modal.** The role of *might* is to express epistemic possibility: that a speaker’s mental state does not rule out a possibility. Incurvati and Schlöder’s treatment of the modal *might* differs from the treatment *might* receives in Update Semantics (Groenendijk, Stokhof, and Veltman, 1996; Veltman, 1990, 1996). In Update Semantics, the context is taken to be a collection of *information states* that support propositions. Information states are defined over the powerset of atomic sentences. Sentences are associated to update clauses. Among them, Veltman treats *might* updates as *non-classical tests*:

With A a set of atomic sentences, for every sentence φ of L_1^A and information state σ ,

$$\begin{aligned} \text{might} : \quad \sigma[\text{might } \varphi] &= \sigma \text{ if } \sigma[\varphi] \neq \emptyset \\ &= \emptyset \text{ if } \sigma[\varphi] = \emptyset \end{aligned}$$

Might behaves like a test that verifies whether an update of σ with φ is possible. But it does not perform an update.

By contrast, Incurvati and Schlöder’s treatment of *might* is that it is an operator that is introduced when we move from weak to strong assertion. In this sense, *might* is an update:

$$(+\diamond \text{ intro}) \frac{\oplus P}{+\diamond P} \quad \frac{+\diamond P}{\oplus P} +\diamond \text{ elim}$$

However, they also take *perhaps it is raining* and *perhaps it might be raining* to be inferentially equivalent.

$$\oplus\diamond \text{ intro}) \frac{\oplus P}{+\diamond P} \quad \frac{\oplus\diamond P}{\oplus P} \oplus\diamond \text{ elim}$$

4.3 What a weak assertion picture supposes

To give more evidence in favour of their weak assertion analysis, I assess the embedding behaviour of *perhaps* and *might*. If *perhaps* and *might* are markers that correspond to a speech act, we can postulate that their linguistic use will follow patterns that have been associated with similar markers. So I first lay out expectations regarding their embedding behaviour (sections 4.3.1 and 4.3.2). Then, in section 4.4, I gather corpus evidence of their pattern of use and compare it to these expectations. While this project is not a full-fledged quantitative study, it provides additional insight on the linguistic behaviour of the two lexical items.

4.3.1 Expectations regarding *perhaps* embedding behaviour

First, if *perhaps* is a force modifier that operates at the speech act level, there are linguistic patterns of use that we can identify. Force modifiers change the effects of an utterance, either in terms of the information update it carries or in terms of how it affects the deontic structure of a conversation. Sticking *perhaps* in front of a strong assertion makes it *not a full-fledged assertion anymore*. Similarly, *perhaps* applies to commands (69) and makes them weaker.

(69) Perhaps close the window?

In that sense, force modifiers differ from other ways speakers use to signal their attitudes, like *speaker-oriented adverbs* or *utterance modifiers*.

Utterance modifiers, speaker-oriented adverbs and force modifiers. *Utterance modifiers* like *frankly*, *confidentially*, *presumably* affect the commitments that the speaker is willing to endorse after her utterance. They modulate the deontic effect of an utterance. For example, *Off the record* is an utterance modifier that speakers often use to preempt being used as a public warrant for the content of their utterance. When they say *Off the record, P*, speakers indicate that they do not wish the content of their utterance to become part of a public conversation. Speakers use utterance modifiers for all sort of reasons. For example, they might use hedges (*I think that, I believe that*) (Benton and van Elswyk, 2020) to lessen their commitment to an utterance by expressing uncertainty. Or they might use

modifiers like *probably* to show that the epistemic grounds for their utterance are probabilistic rather than certain (Moss, 2018).

Speaker-oriented adverbs, like *fortunately*, *surprisingly* reflect attitudes of the speaker with respect to the content of their utterance. They serve to introduce, by way of conventional implicatures, the mental state of the speaker while not modifying the *at-issue content* of the utterance (Potts, 2005).

Force modifiers like negation, *perhaps*, *probably* affect the update that an utterance produces. They change the illocutionary effect of an utterance (Austin, 1975). When the speaker in (68) uses *perhaps*, she indicates that she does not wish to perform an assertion in the conversation with *it is raining*. She merely wishes to leave open the possibility that *it is raining*. So she performs a weak assertion: she expresses her refusal to remove *P*-worlds from the set of worlds compatible with what the speakers agree upon in conversation.

(68) Perhaps it is raining.

While these lexical items differ, they have in common that they operate at the level of the speech act that an utterance performs, not at the level of the content. Speaker-oriented adverbs introduce a judgment of the speaker with respect to the content of the speech act she performs. Utterance modifiers modulate the performance of the speech act: where the speaker adds specifications concerning the deontic effects of her speech act. Force modifiers change the speech act being performed: negation operates a switch from assertions to rejections, commands to interdictions, etc. As all of these items operate at the level of the speech act being performed, they share embedding patterns.

Projection behaviour. The key commonality between *speaker-oriented adverbs*, *utterance modifiers* and *force modifiers* is their projection behaviour.

Speaker oriented adverbs project out of attitude reports (Potts, 2005). With *unfortunately* in (70), the speaker expresses her attitude of disappointment with respect to the content of her assertion.

(70) Unfortunately, Beck survived the descent. (Potts, 2005, p. 14)

This type of expression tends to escape local contexts.

(71) # He believes that Beck unfortunately survived the descent, and I am stoked that Beck is alive

In terms of *projection behaviour*, speaker-oriented adverbs require strong contextual felicity (Tonhauser et al., 2013) (see also section 3.1.1). In simpler terms: speaker-oriented adverbs require that the speaker actually endorses their contribution. When speaker-oriented adverbs occur in local contexts (such as conditional antecedents, belief contexts), they *escape* the local context to characterise the attitude of the speaker with respect to the embedded content.

(72) If Beck unfortunately survives the descent, they will fall into a bear trap.

(73) He believes that Beck unfortunately survived the descent.

To clarify: in (72) and (73), *unfortunately* does not qualify an attitude of the speaker in the conditional antecedent (resp. an attitude of the belief holder). Instead, it qualifies the attitude of the speaker with respect to the hypothetical antecedent (resp. the reported belief).

Similarly to speaker-oriented adverbs, *utterance modifiers* such as *frankly*, *confidentially*, *off the record*, do not embed under attitude reports. They project out of them.

(74) # He believes off the record that Beck survived the descent.

Utterance modifiers escape local contexts to modify the whole utterance. The same is true concerning conditional antecedents.

If *perhaps* is a force modifier that applies to the performance of a speech act within an utterance, similar behaviour is expected. If *perhaps* did embed easily within attitude reports and conditional antecedents, it would mean that it is not the speech act it modifies, but the proposition expressed by the utterance. For these reasons, part of the study I conducted on corpora of English was dedicated to gathering examples of *perhaps* occurring in conditional antecedents (section 4.4.1) or attitude reports (section 4.4.2) Then, I took a closer look at these examples to see if they showed actual embedding behaviour, or if *perhaps* escaped the local context to modify the speech act of the speaker.

4.3.2 Predictions regarding *perhaps* and *might*

The second aspect of Incurvati and Schlöder's predictions concern the modal *might*. They interpret *might* as a sentence modifier corresponding to *perhaps* that expresses epistemic possibility. If *might* corresponds to the same attitude of 'leaving open a possibility' from the speaker as *perhaps*, then we expect one of two things.

- i. Either *perhaps* and *might* should co-occur, in the same way as *no* and *not* co-occur. *Perhaps* indicates the force of the sentence, *might* is the associated verb.
- ii. Or, more plausibly, *might* and *perhaps*, in the same way as *Probably* and *it is probable*, wouldn't occur together much, as it would cause a redundancy effect.

We already know that (i) is not the case. There are many examples where speakers use *perhaps* with an indicative verb, like (75). There are many examples where speakers use *might* without the force modifier *perhaps*, like (76).

(75) Perhaps he went to Daniel’s party.

(76) She might be scared of spiders.

So, since *might* and *perhaps* do not co-occur in the same way as *no* and *not*, my goal is to gather fine-grained examples of the cases where they do co-occur. These examples can help determine whether *perhaps*-utterances and *perhaps might*-utterances are, in fact, inferentially equivalent. If, instead, *perhaps* and *might* produce a strong redundancy effect in the resulting examples, it would weakly support the idea that *might* is a specific type of update. Bluntly: if *perhaps might* examples sound like the speaker is doing two moves, weak assertion and epistemic possibility, this would counter the claim that *perhaps* and *perhaps might* are inferentially equivalent. Or, at the very least, this would call for an explanation of the way that *perhaps* and *might* can sometimes invoke different contexts to frame their update (Incurvati and Schlöder, 2019). In section 4.4.5, I searched for *perhaps might* utterances, and tried to determine whether they produced a redundancy effect.

The present study aims to find examples for these preexisting hypotheses concerning the embedding behaviour of *perhaps* and *might*. I do so by analysing their occurrences in specific contexts: conditional antecedents, together, and under attitude reports. My goal is to support the budding analysis of weak assertion. With a wide array of examples that qualify the pattern of use of its associated expressions, we can understand better the way speakers use weak assertion.

4.4 Gathering examples

I conducted a study to assess the behaviour of *might* and *perhaps* with respect to what their analysis as weak assertion indicators predict. This study was done by querying on the British National Corpus 1994 (BNC) and the Corpus of Contemporary American English 2021 (COCA) for occurrences of *perhaps* and *might* in three types of contexts: (i) in conditional antecedents (ii) under attitude reports (iii) in co-occurrence.

To assess the projection behaviour of *perhaps* and *might*, I observed whether we find examples of their occurrence under conditional antecedents. The reasoning is that, if *might* is the sentence modifier corresponding to the force modifier *perhaps*, in the same way as *not* is the sentence modifier corresponding to the force marker *No!*, then *might* should embed under conditional antecedents where *perhaps* cannot.

In (77), the sentence modifier *not* is appropriate under the scope of the conditional antecedent where the force marker *No* is not.

(77) a. # If ‘Is it raining? No!’, then we don’t need an umbrella.

- b. If it is not raining, then we don't need an umbrella.

Treating *perhaps* as a force modifier that operates at speech act level supposes that it isn't appropriate in the context of conditional antecedents (section 4.3.1). Conversely, if *might* is a sentence modifier, it should be appropriate in conditional antecedents.

The second test to determine whether *perhaps* operates at the speech act level was to assess its behaviour under attitude reports. Because discussions of *perhaps* and *might* situate their contribution at the level of *epistemic certainty* (Knobe and Yalcin, 2014), I restricted myself to verbs that speakers use to report epistemic attitudes: *think*, *believe* and *know*. The idea was to avoid attitude verbs that indicate desire-like attitudes, such as *wish*, *hope* or *fear*.

Thirdly, I wanted to assess the co-occurrence of *perhaps* and *might*. Do they often co-occur, or do they produce a redundancy effect by expressing the same attitude? First, I compared the number of cases where *perhaps* and *might* occurred together to the total number of their occurrence. This comparison provides a quick and dirty way of assessing the frequency of their co-occurrence. I then examined closely token examples of cases where *might* occurred under the scope of *perhaps*. The goal in analysing specific examples was to determine whether *perhaps* and *might* together gave the impression that the speaker was performing the same speech act twice (redundancy) or whether the speaker seemed to use *perhaps* and *might* to express two distinct attitudes.

Corpora. The BNC¹ and COCA², among the BYU corpora, are two contemporary corpora for British and American English.³ They are freely accessible online, and provide an integrated querying system, enriched with part-of-speech tags, that I used. They are remarkable for both their quantity of entries and their qualitative variety. The BNC contains 100 million words, the COCA 1 billion words. The BNC ranges from the 1980s to 1993, the COCA from 1990 to 2019. Both are balanced corpora: they cover a wide range of text categories, with similar representation of each text category. Examples come from spoken discourse, fiction, popular magazines, newspapers, academic texts, and, for the COCA, TV and movie subtitles and web pages.

As a baseline, all the results concerning the behaviour of *perhaps* and *might* within relevant contexts should be compared to the total number of occurrences of *perhaps* and *might* in both corpora. The total number of occurrences of *perhaps* in the BNC is 33306. In the COCA, *perhaps* appears in 191436 entries. *Might*,

¹<https://www.english-corpora.org/bnc/>

²<https://www.english-corpora.org/coca/>

³To sidestep methodological concerns regarding the construction of the COCA compared to the BNC (Egbert, Biber, and Gray, 2022; Egbert, Larsson, and Biber, 2020), I chose to study both corpora.

as a verb, occurs 58344 times in the BNC, and 521997 times in the COCA.⁴ By contrast, the occurrences I gathered of either *perhaps* or *might* in conditional antecedents, with each other, or under attitude verbs (*think*, *believe*, *know*) are comparably rare. While the goal of this chapter was more to find examples of *perhaps* and *might* behaviour than to develop a full-fledged corpus study, I summarise some numerical results in table 4.1.

		BNC	COCA
total frequencies	might	58344	521997
	perhaps	33306	191436
conditional antecedents	might	4	678
	perhaps	11	23 (out of 200)
redundancy	Perhaps S might	130	543
attitude reports	think S might	49	578
	think that perhaps	30	181
	think that S perhaps	0	3
	believe that perhaps	2	40
	believe that S perhaps	0	2

Table 4.1: Numerical results

4.4.1 Conditional antecedents

If* and *might To observe the embedding behaviour of *might* under conditional antecedents, the first step was to search for constructions of the form *if + pronoun + might* and *if + noun + might*, using the corpora query systems and PoS tagging. On the COCA, this yielded the following results: *if + noun + might*: 310 occurrences. *if + pronoun + might*: 3,188 occurrences. Afterwards, I filtered these results by close reading to remove false positives where *if* was part of comparative or concessive clauses, such as *as if*, *even if*, indirect questions (introduced, for example by *ask if*, *wonder if*, but also other verbs) or speech reports (*tell if*, *say if*). The goal was to keep only cases where *if* signals an actual conditional antecedent. The COCA contained 678 instances of embedding *might* under a conditional antecedent. 24 had the form *if + noun + might*, 654 the form *if + pronoun + might*. The construction *if + noun + might* yielded 23 results, of which 4 actual conditional antecedents in the BNC. The list of instances is available, in table form, at github.com/lwennb/weak_assertion_examples.

⁴I used PoS tagging to exclude cases where *might* is used as a noun, synonymous with *strength* or *force*.

- (78) there was a need for a public authority to play a role in securing access to and observance of the law, where the issues involved justified it, if individuals might otherwise be deterred from bringing cases to the courts because of their complexity, the financial costs involved or for any other reason. (BNC:HP3, *Industrial Law Journal*. Oxford: OUP, 1993)

Within the COCA results, 96.5% of *if [subject] might* constructions were of the form *If + Pronoun + might*. Among these, 74% were of the form *If I might* (484 entries).

- (79) They really haven't thought this whole thing through, if I might be permitted to state the obvious. (COCA:2012:BLOG splcenter.org)

As exemplified in (80) *If I might* is a paradigmatic formulation to politely request permission for an action:

- (80) Sheila Altman had green eyes and a soft voice. She said "If I might" and "Would you mind," and never forgot to say "Please" [...] (COCA: FIC: *The Atlantic Monthly*, 2004 (May). Vol. 293, Iss. 4; pg. 149, 13 pgs, *foaling season*, Aryn Kyle)

Speakers, like the one in (79) use the formula *if I might* to preface a speech act. By doing so, they express their wish to perform a speech act or action conditional on acceptance from the audience. In this sense, these are occurrences of *might* within a conditional antecedent. However, the consequent of such conditional is not a truth-apt proposition, but the performance of an act. This reading as a request for permission is consistent with the over-representation of *if I might* entries among *if + pronoun + might* in the COCA. I hesitated whether to consider such examples, including those of parenthetical *if I might*, relevant. I decided to keep them: the use of *might* as a deontic operator to request permission gives evidence that speakers sometimes hedge their utterances by asking for permission to perform a speech act. Further analysis of the resulting commitment structure would be needed; however, at first glance, it seems that an interpretation as *conditional update* is plausible.

- (81) kill him with my third round. GANT: Captain, if we might have a word... Algren stops, turns, raises his REVOLVER and FIRES (COCA:2003:FIC Mov:LastSamurai)
- (82) That is patent balderdash! - Is it? Gentlemen, if we might keep things cordial, we may make some progress. (COCA:1999:MOV Topsy-Turvy)

If and perhaps. I conducted a similar search to observe the embedding behaviour of *perhaps* under conditional antecedents. a research for *if + perhaps* yielded 402 entries in the COCA, 58 entries in the BNC. I selected the BNC entries, and a randomised 200 lines sample from the COCA examples. Then, I annotated them manually from close reading to keep conditional antecedents. The annotated sample is available in table form at github.com/lwennb/weak_assertion_examples. Among this sample of 258 entries, there were lots of false positive results, with 116 constructions of the form *wonder if*, and 14 of the form *ask if*. Thus, I disregarded 33 such BNC entries.

In addition, the construction *if perhaps* is used to modulate an adjective in 37 cases, as in (83). I also eliminated these false positive examples.

- (83) For many centuries people assumed that animals had a similar, if perhaps more limited, view of the world (BNC:FEV)

That left some actual conditional antecedents containing *perhaps*: 11 in the BNC, 23 out of a 200 lines randomised sample of the COCA.

- (84) If you want to change your sex life, if perhaps your needs are different to your partner's, then read on. (BNC:B3G)

The question is to assess whether, in conditional antecedents, *perhaps* takes a narrow or wide scope.

Among these conditional antecedents, some were transcribed with commas (*If, perhaps, ...*) indicating that *perhaps* was used as a parenthetical. In these cases, the contribution of *perhaps* scopes out of conditional antecedents.

- (85) If, perhaps, Katherine was a little more credible (and Gerardine Hinds's rather uncertain playing doesn't help), there might be more to grapple with. (BNC:A1D)

That *perhaps* projects out of the conditional antecedent in (85) is confirmed further by the use of *might* in the consequent and the *little more* hedge. The whole conditional is weakly asserted: the speaker leaves open the possibility in the conversation that *if Katherine was a little more credible, there might be more to grapple with*.

On the other hand, there are no commas around *perhaps* in (86) to signal it as a parenthetical.

- (86) If perhaps the Government were to fund victim support properly, Erm, Mr (—) who was burgled and has never got over it might perhaps have had a visit and some counselling from a victim support worker, and that would be a very good thing. (BNC:JS9)

To determine whether *perhaps* takes a narrow scope or wide scope, then, I observe which scope pattern fits best. (87a) is a wide scope reading while (87b) is a narrow scope reading.

- (87) a. Perhaps (if the Government were to fund victim support properly, Mr (—) might perhaps have had a visit)
- b. If (the Government were perhaps to fund victim support properly)(Mr (—) might perhaps have had a visit)

In (87), *perhaps* is ambiguous between a wide and narrow reading.

To force a narrow reading of *perhaps* under the conditional antecedent, I also searched for *If [subject] perhaps* constructions. For the COCA, there were a total of 22 such constructions across the whole corpus; and 2 occurrences in the BNC. Excluding false positives, 11 occurrences in the COCA, and 2 occurrences in the BNC. Given these results, it seems fair to say that cases of embedding *perhaps* under conditional antecedents with a narrow reading are extremely rare, and in the examples, the reading is ambiguous between a narrow and wide scope at best.

- (88) If Earvin perhaps just came out and said, "Don't do sex at all," that may be a little difficult for kids to accept. (COCA:SPOK: ABC Brinkley, 1991)
- a. Perhaps (If Earvin just came out and said: "Don't do sex at all", it would be difficult for kids to accept).
- b. # If (perhaps Earvin just came out and said: "Don't do sex at all"), (it may be difficult for kids to accept.)

Paraphrasing the example (88) to obtain a wide and narrow reading shows that the narrow reading does not resist further examination: it is incorrect. So even in examples *if [subject] perhaps*, that I gathered in an attempt to force narrow reading of *perhaps* in conditional antecedents, the contribution of *perhaps* projects *out* of conditional antecedents and takes wide scope over the sentence.

4.4.2 Attitude reports

The second aspect to examine was the predicted difference in behaviour of *perhaps* and *might* in attitude reports. For these, I searched for the respective occurrences under *think*, *believe* and *know*. Since in the results for *think that perhaps*, *perhaps* is ambiguous between a narrow scope and wide scope, I also compared the occurrences of *think that [subject] perhaps*, to force a narrow reading.

A search for *Think that + [subject] + might* shows that *might* embeds well under attitude reports. In the COCA, there are 578 such constructions; in the BNC, 49. By contrast, *think that + Perhaps* occurs 181 times in the COCA; 30 times in the BNC. Strikingly, among the 181 COCA examples of *perhaps* under *think that*, 84 are spoken and 80 are from movie or Web sources, where speakers often mimic spoken patterns. While the construction does occur, it does so in cases where speakers do not prepare their speech.

When I tried to force a narrow scope reading of *perhaps*, within constructions like *think that + [subject] + perhaps*, there were only 3 hits on the COCA, 0 on the BNC. A similar result holds for *believe that perhaps*, with 40 occurrences in the COCA, 2 in the BNC; while *believe that + [subject] + perhaps* occurs 2 times in the COCA, nowhere in the BNC. And the results for *know that perhaps* are identical.

4.4.3 The case of indirect discourse

Within the examples for *perhaps* under *think*, there are cases where *think that* introduces indirect discourse with elements of free indirect discourse.⁵

- (89) While she lay in bed, Molly began to think that perhaps his first ever girlfriend was actually a boy, which would have been fine, of course. She began to think that perhaps his first ever girlfriend was actually imaginary, which would have been... less fine. (COCA:2018:FIC Crazyhorse)
- (90) But my husband! What will he think?' 'He will think that perhaps you are dead,' said the old woman. (COCA:2015:FIC Modern Age)

In indirect discourse, or indirect thought, indexicals and utterance modifiers are typically transposed in a process of reporting (Wales, 2011, p. 224sq). By contrast, free indirect discourse dispenses with the speech act verb and some transpositions (Wales, 2011, p. 175). The distinction between indirect discourse and free indirect discourse is sometimes muddy. In a rigid version, indirect discourse should transpose most of the speaker's speech markers in reported speech; to the extent that some scholars define reported speech that contains any sign of direct speech features as free indirect speech (Leech and Short, 2007). In less rigid version, indirect discourse does tolerate some mimetism of direct speech.

It might depend on one's definition of indirect discourse whether to classify examples (89, 90) as indirect discourse or free indirect discourse. However, in both examples, the speaker uses *think* to introduce a proposition corresponding to the attitude holder's mental discourse. So in this sense, these examples don't show an embedding of *perhaps* under an attitude report, but the use of the attitude report verb to introduce an indirect discourse, reproducing what the holder of the attitude would say about their mental state. In (89), the verb *think* is used to introduce a stream of consciousness within which the subject performs weak assertions. A similar structure governs (90). In both these examples, *think* introduces an indirect discourse rather than a subordinate proposition.

⁵In this discussion, I use the terminology *indirect discourse* and *free indirect discourse* to refer to the stylistic phenomenon by which a speaker reports utterances or thoughts. (*Free Indirect Discourse*, in style analysis, is also known as *indirect speech* or *indirect style*. However, since in other parts of my thesis I discuss *indirect speech acts*, the terminology *indirect speech* risks creating confusion.

4.4.4 Think that perhaps: a superhedge

Another type of cases are cases where *think* and *perhaps* are used together as a sort of *superhedge*. This is particularly salient in (91) and (92), where the speaker is trying to lessen the impact of a negative statement.

- (91) As a singer, look, I think that perhaps she's making a whole lot of sense, again, artistically. But she says she's trying to bring her fans and Chris's fans together. (COCA:2012:SPOK CNN Showbiz)
- (92) I also think that perhaps the author struggled with the purpose of this book at points. Is it for academicians? For therapists? For people wanting to just be better? (COCA:2012:WEB Goodreads)

The constructions with *think that [subject] perhaps* do seem to introduce a weak assertion under an epistemic attitude. Example (94) is particularly remarkable. In it, the speaker is reporting her epistemic state, where a possibility is left open: *Geilie had been right*. This contrasts with the case in (95), where *perhaps* takes wide scope. It refers to the beliefs of the speaker who leaves open the possibility that dreams have a practical function, rather than the beliefs of researchers. The group of researchers would not *weakly assert* that its possible that dreams have a practical function: a portion of it would *strongly assert* it, while others would assent to this update of the Common Ground.

- (93) Have you come to think that people perhaps expect too much of it and its ability to effect change? (COCA:2016:SPOK: PBS NEWSHOUR)
- (94) I began to think that Geilie perhaps had been right in considering this a fairly lenient sentence, given the overall state of current Scottish jurisprudence, though this didn't alter by one whit my opinion as to the barbarity of it. (COCA:1991:FIC *Outlander*, Diana Gabaldon)
- (95) Most researchers now think that dreams perhaps have a practical function than - that's not symbolic. (COCA:2012:SPOK NPR: *Talk of the Nation*)

The embedding behaviour of *perhaps* under attitude reports (or, more exactly, absence thereof) corresponds to the pattern predicted by Incurvati and Schlöder's analysis of *perhaps* as a force modifier. If *perhaps* modifies an utterance at the speech act level, its use in free indirect discourse is expected. Conversely to indirect discourse, that transposes a lot of idiomatic expressions and indexicals, free indirect discourse often keeps the speakers' speech markers. Otherwise, its use under attitude reports, comparatively rare, often leads to wide scope readings.

4.4.5 Redundancy

The last step was to assess whether *perhaps* and *might* produce a redundancy effect. A search for *Perhaps + [subject] + might* resulted in 130 hits in the BNC, and 543 in the COCA. Among the COCA results, 479 were of the form *perhaps + pronoun + might*. Compared to the total frequencies of *perhaps* and *might* in the corpora, these are pretty low numbers (table 4.1). However, these numbers show that *might* also embeds under *perhaps*.

The examples I gathered of *might* occurring under *perhaps* include permission requests and suggestions.

(96) Perhaps we might arrange a dinner of our own? Tomorrow? (COCA:FIC ‘Fallacious’, Sean Vivier, in *Analog Science Fiction & Fact*, Vol. 136 (7/8), 2016, pp. 125-129)

(97) ‘Now, perhaps I might have a look at the injury?’ (COCA:FIC *Memoirs of a Geisha*, Arthur Golden, 1997)

In the cases of suggestions or permissions requests, speakers seem to use *perhaps* and *might* as a way to hedge the utterance. Of course, this way of hedging could have a politeness explanation. The speaker adds modals and weak expressions so that she doesn’t overtly make suggestions or request permissions. (96) would then be equivalent, with a flourish, to saying *Let us arrange a dinner of our own* and (97) an equivalent to *Can I take a look at the injury?*

But the use of *perhaps* and *might* could also be given an Incurvatian explanation: the use of *perhaps*, as a force modifier, serves to indicate a *weak* speech act. When a speaker requests permission, or suggests a course of action, she opens the possibility of her action. While cases of command or advice-giving are not truth-apt, they can be represented as updates in Common Ground models that integrate deontic effects of speech acts (section 2.2.2). Their weak counterparts, in turn, can be modelled as keeping a possibility open: ruling out a counter-update.

Compare with cases where a speaker gives a command, even if it is formulated as a request.

(98) Please close the blinds.

When a speaker gives a command, if she is successful, she reduces the context set to only those worlds where the content of her command is on her addressee’s to-do list (Portner, 2004). This result of removing possible worlds from the context set is identical if she formulates a request with *might*, for politeness.

(99) Please, might you close the blinds?

A request is a speech act that requires uptake from the addressee in the form of a *yes* or *no* answer – a *closed call* in Caponetto’s terminology (2017). While its

illocutionary effect is not directly to narrow down the context set, it updates the Common Ground by way of dividing it in alternative sets: these worlds where the addressee answers yes vs. these where they answer no. In the same way, weakened requests or suggestions modified by *perhaps* seem to aim to keeping a possibility open. With (96) *Perhaps we might arrange a dinner of our own?*, the speaker does not overtly put an expectation that the addressee will answer by *yes* or *no*. She brings out to salience that the possibility of arranging dinner is open.

Another way the construction *perhaps [subject] might* is used is to introduce hypothetical scenarii:

- (100) Perhaps I might have been a lot less defensive if I had known then what I know now. (COCA:MAG ‘The laughter prescription’, Norman Cousins, *Saturday Evening Post*, 262 (6), pp.32-110, 1990)

In these cases, similarly to the use of *might* in conditional antecedents, the combination of *perhaps* and *might* seems to indicate a possibility that is more remote epistemically from the speaker. In fact, the possibility that she considers is already closed. The speaker is simply musing on an unactualised possibility.

These contrast with cases where the combination of *perhaps* and *might* indicate weak assertion:

- (101) Perhaps I might take Lady, my collie, for a nice, leisurely walk through town and window-shop on our way to the park. (COCA:FIC *Every hidden fear: a Skeet Bannion mystery*, Linda Rodriguez, 2014)

In these cases, *perhaps* and *might* indicate a possibility that is not yet closed, that the speaker does not wish to reject. These cases where *perhaps* and *might* are both used to perform a weak assertion do not seem to force a reading of *might* as a specific update. In fact, in the extended context of the speaker’s stream of consciousness in (101), *perhaps I might* seems equivalent to the speaker having said *I thought I might*. If we remove the past tense transposition of indirect discourse, a direct discourse equivalent would also be: *I thought: ‘Perhaps I will take Lady for a walk’*.

- (101) I’d thought I might read a book for pleasure or sit on the front porch and knit in the unseasonably warm weather. One of the advantages of giving up a hard-hitting career with the Kansas City Police Department and moving to a small college town as head of campus police was the slower pace of life here right outside the city. Perhaps I might take Lady, my collie, for a nice, leisurely walk through town and window-shop on our way to the park.

While the results for the co-occurrence of *perhaps* and *might* are not yet decisive for the inferential equivalence of *perhaps + indicative* and *perhaps + might*, they seem to license cautious optimism. Using *perhaps* and *might* together is uncommon, as signalled by the low number of occurrences. In the cases where

they do occur together, they seem to perform the same *weak* update as *perhaps* or *might* alone.

In requests or other such formulations, *perhaps + might* does behave differently than *perhaps* or *might* alone. It introduces a sort of *superhedge*, making the request weaker than just *perhaps* or *might*. However, this superhedge effect could also be explained as politeness, given by the pragmatic effect of combining two weak indicators.

4.5 Conclusion

In this chapter, I investigated a complementary speech act to weak rejection: *weak assertion*. Incurvati and Schlöder (2019) provide theoretical reasons to integrate weak assertion in their multilateral system. In their 2017 paper, they coin the notion of *weak rejection* to address the fact that the speech act of *rejecting P* at play in bilateral inferentialist works (Rumfitt, 2000; Smiley, 1996) does not always provide grounds to assert $\neg P$ in natural language. *Weak rejection* is the speech act whose essential effect is represented as *blocking* an update of the Common Ground with the rejected statement. When a speaker *weakly rejects P*, she aims to leave open a possibility that $\neg P$. If some speech acts can be weak, so that they *block* updates instead of proposing them, it makes theoretical sense to postulate a mirror speech act: *weak assertion*, where a speaker *blocks* removing a possibility from the Common Ground. When a speaker *weakly asserts P*, she aims to leave open a possibility that *P* in the Common Ground.

Incurvati and Schlöder also provide linguistic evidence for the speech act of weak assertion. The utterance modifier *perhaps* can be used by speakers to answer polar questions. Frege (1918b) and Dummett (1981) analyse an *assertion* of *P* as the answer *Yes* to the question *Is it the case that P?*. In the same fashion, Incurvati and Schlöder analyse a *weak assertion* of *P* as the answer *Perhaps* to the question *Is it the case that P?*. Thus, they treat the utterance modifier *perhaps* as a *force modifier*; and take *might* to be the sentential operator corresponding to weak assertion.

The aim of this chapter was to test the claim that *perhaps* embedding patterns correspond to what is expected of a *force modifier*. To do so, I conducted a study to assess the occurrence of *perhaps* and *might* in contexts relevant to assessing their projection behaviour: their occurrence in conditional antecedents, their co-occurrence, and their occurrence under attitude reports. This, of course, leaves room to further study of the embedding pattern of *perhaps* under quantifiers, and whether *perhaps* also weakens commands.⁶

The use of *perhaps* and *might* by speakers supports their treatment as indicators of weak assertion. Occurrences of *perhaps* and *might* under conditional

⁶The large number of uses of both *perhaps* and *might* to hedge suggestions, commands, or requests in (4.4.5) is an encouraging sign.

antecedents and attitude reports are comparatively rare, when compared to a total frequency of use. I was conservative in disregarding examples from *if + perhaps* and *if + might* results. I kept a lot of examples that might be better analysed as concessives or conditional speech acts. The cases of actual conditional antecedents in a propositional sense are even rarer. However, these first findings regarding occurrences of *perhaps* and *might* under conditional antecedents show that there are few cases where *perhaps* gets a narrow reading under conditional antecedents.

Similarly, while it is used in indirect discourse, *perhaps* almost never embeds with narrow reading under attitude reports. This supports the hypothesis that *perhaps* performs a discourse move of the speaker. In contexts like attitude reports, the occurrences of *perhaps* modify an indirect discourse attributed to the attitude holder. That is: in examples like (89), the attitude verb *think that* is used to introduce an indirect discourse that mimics the attitude holder mental state.

- (89) Molly began to think that perhaps his first ever girlfriend was actually a boy, which would have been fine, of course. (COCA:2018:FIC Crazyhorse)

In examples like (91), *think that* does introduce a proposition qualifying the speaker's epistemic state. But it is unclear whether the speaker uses *perhaps* in this example as a speech act modifier or as a sort of *superhedge* to reduce the impact of her statement.

- (91) As a singer, look, I think that perhaps she's making a whole lot of sense, again, artistically. But she says she's trying to bring her fans and Chris's fans together. (COCA:2012:SPOK CNN Showbiz)

One thing is certain though: with other utterance modifiers (*frankly, confidentially*), *perhaps* does not like plugs. It occurs rarely under attitude verbs. Under conditional antecedents, *perhaps* usually takes wide scope, or is ambiguous between a narrow and wide scope.

But these examples go further than merely supporting Incurvati and Schlöder's claims that *perhaps* pattern of use is consistent with it being a force modifier. Throughout the different embedding patterns I tested, there were indeed cases where speakers use *perhaps* and *might* to indicate a weak assertion, grounded on their epistemic state allowing for a possibility. But there were also a number of cases came out where *perhaps* and *might* are used by the speakers as a way to *hedge* their utterances. There are cases where a speaker hedges their assertion of a negative judgment, but also cases of requests weakened towards mere suggestions.

Does that mean that *perhaps* in those examples loses its force-modifying effect? Not really. In a commitment analysis of speech acts, it makes sense that speakers would weaken their speech acts as a way of hedging their utterances. Suppose a strong speech act, like an assertion or a command. When a speaker

performs it, they propose an update of the Common Ground, and endorse responsibilities (to justify their assertion, to have the authority to give a command). However, when they perform a weak speech act, like a weak assertion or a suggestion, they manifest that they leave open a possibility. When a speaker makes it known that they wish to leave open a possibility in the conversation, they also make this possibility *salient*. Combined with pragmatic reasoning and a cooperative hypothesis that speakers make their contribution *relevant* to a conversation (Grice, 1975), these weak speech acts have the effect of expressing to the audience that the speaker is sympathetic to a possible update of the Common Ground, but doesn't wish to endorse all the commitments associated with it. This is how to make sense of examples like (96). The speaker overtly makes a weak suggestion – that aims to keep the possibility of arranging a dinner in the mutual to-do list –, but indirectly proposes an update.

- (96) Perhaps we might arrange a dinner of our own? Tomorrow? (COCA:FIC 'Fallacious', Sean Vivier, in *Analog Science Fiction & Fact*, Vol. 136 (7/8), 2016, pp. 125-129)

Studying the embedding behaviour of *perhaps* and *might* gives further evidence for the speech act of *weak assertion*, already theoretically motivated. Moreover, it opens promising avenues to discover other weak speech acts, and understand their relation to conversational pragmatics. I gave the example of analysing how speakers can hedge requests as suggestions, while still making it obvious, by way of implicatures, that they wish their suggestion to be followed. This is one out of many examples of ways speakers use weak speech acts in conversation. Studying the relation between stronger and weaker speech acts, as well as how speakers weaken and hedge speech acts, would give us much needed insight in a cartography of speech acts as they relate to speakers' commitments.

Chapter 5

Retraction and Common Ground[†]

5.1 The Retraction Dilemma

Often, speakers wish to cancel previous discourse moves they have made: assertions, questions, promises, imperatives. To do so, they use *retractions*: the act one performs when saying “I retract that” or “I take that back”.

- (102) a. Shannon: Albert should be here in 15 min.
b. Shannon: I take that back – he texted me that he got held up at work.

Retractions are the only speech act that can cancel the illocutionary effect of a previously made, successful utterance. Our goal in this chapter is to explain how retractions work, and affect the conversational picture, by cancelling the illocutionary effects of utterances. In particular, we want to understand better the status of speech acts that have occurred between a target utterance and its retraction.

Retractions pose an interesting dilemma. On the one hand, they are needed so that speakers behave responsibly. When a speaker realises that a speech act they performed was infelicitous, for example that their assertion is not true, or that they won't be able to keep their promise, the cooperative thing to do is to retract (MacFarlane, 2014). A speaker who does not retract a claim that they know is false, or a promise that they won't be able to keep, exhibits blatantly uncooperative behaviour. While examples in politics abound, let us build a more innocent one.

- (103) Shannon never retracted her assertion in (102a). An hour after her assertion *Albert should be here in 15 min*, Albert has yet to arrive, and you ask her about it. She breezily answers:

[†]This chapter is based on joint work with Incurvati, Sbardolini and Schlöder. For an overview of the authors' respective contributions, see section 1.3.

– Oh yeah, he texted me that he got held up at work.

In (103), Shannon did not stand by her assertion anymore. However, she never shared this relevant information. Because of that, she risks losing status as a trustworthy participant in the conversation. P. Brown and Levinson define the notion of face:¹ *the public self-image that every member [of society] wants to claim for himself* (P. Brown and Levinson, 1999, p. 311). According to them, speakers who exhibit *blatant noncooperation in an activity* threaten the addressee's face (p. 314) by showing disregard for the addressee's wants. They go against a certain social construction:

In general, people cooperate (and assume each other's cooperation) in maintaining face in interaction, such cooperation being based on the mutual vulnerability of face. That is, normally everyone's face depends on everyone else's being maintained, and since people can be expected to defend their face if threatened, and in defending their own to threaten others' faces, it is in general in every participant's best interest to maintain each other's face [...] (P. Brown and Levinson, 1999, p. 311sq)

Retractions serve to maintain trust, in that they enable speakers to cancel utterances that they are not willing to commit to anymore in a conversation.

On the other hand, retractions have a significant cost. First, the speaker has to admit she is wrong. In P. Brown and Levinson, she has to accept a certain loss of face. Secondly, retractions have a cost for the conversation. This is expressed in the ambiguous treatment of retractions in argumentation theory.

Whether forced or not, retraction of an original thesis always has the serious repercussion that (part of) the dialogue is lost. (Krabbe, 2001, p148)

Krabbe (2001) comes to retractions from the perspective of argumentation theory. Therefore, he analyses the effect of retractions within an agonistic conception of argumentation, where one participant *wins* if they convince the other to retract their opposing claim. His analysis of retractions is that, in an argumentative conversations, speakers make claims that allow them to build deductions. Sometimes, participants challenge each other's claims. And when a participant challenges a

¹P. Brown and Levinson's notion of *face* is a non-phenomenological, social, take on Levinas' notion of *visage* (Levinas, 1984). According to Levinas, another person's *visage* is a pure manifestation of their personhood and alterity, which compels the subject to ethical behaviour.

“Le visage s'impose à moi sans que je puisse cesser d'être responsable de sa misère.
La conscience perd sa première place”

In P. Brown and Levinson (1999) work on politeness, the role of *face* as a construction of self that people present in social interaction compels the subject to social behaviour.

claim, their goal is to get the author of the challenged claim to retract it. In this sense, when a retraction occurs, *a part of the dialogue is lost* in two senses. First, the retractor has conceded a point to the other participants. She has to back down from her claim *A*. Secondly, if the speaker has used her claim *A* in a deductive reasoning from *A* to *B*, and then retracts *A*, her argumentation from *A* to *B* is also lost. As such, retractions do not only threaten the image of the retractor, but can also threaten the cohesiveness of a conversation and disrupt it.

The analysis of retractions in speech acts theory focuses on the effect of retractions on conversations. While it might be ultimately fruitful to study the social effects of retractions on how retractors are perceived, our proposal is to understand how retractions work and affect the course of a conversation. Analysing the conditions and effect of retractions in a conversation may also bring light to other speech acts. At least, this is the argument in the contextualist debate surrounding truth of assertions (MacFarlane, 2014; Marques, 2018): if we know under which conditions a speaker *has to* retract an utterance, we also know about the necessary conditions they needed to meet to perform this speech act in the first place.

So let's dive into the effects of retractions. In (102), when Shannon retracts, she attempts to remove the information she had introduced in the conversation. In particular, she wants her audience to stop believing (on account of her claim) that *Albert will be there in 15 minutes*. She also wants to let go of some of the commitments that her previous assertion brought about. Say, she implicitly committed to set the table in the next 15 minutes. For Shannon, the ideal scenario, when she retracts her assertion, would be that the conversation continues as if she had never made this assertion in the first place. But is the information Shannon retracted really *lost*?

- (104) a. Shannon: Albert should be here in 15 min.
 b. Shannon: I take that back – he texted me that he got held up at work.
 c. # Shannon: I never said Albert was going to be here soon!

If Shannon's retraction completely cancelled her contribution, and completely removed (104a) from the dialogue, then she could felicitously follow up with (104c). But (104c) is infelicitous: Shannon cannot cancel the fact that she made an assertion. Caponetto (2020) explains that retractions only cancel the illocutionary effects of utterances (Austin, 1975), i.e. the change in information and deontic norms that they brought about. The locutionary effect, that a speaker produced an utterance, linger. A retraction cannot make it that an utterance did not happen.

Two aspects, however, remain largely unexamined in speech acts accounts of retraction. Discussions of retractions in the truth-contextualist debate (MacFarlane, 2014; Marques, 2018) and to a lesser extent in the work of (Caponetto, 2020) share the assumption that retractions immediately and successfully affect

the conversation if the speaker is the author of the claim she retracts. Moreover, these accounts remain silent concerning potential further utterances in the conversation that happen between an utterance and its retraction.

Our goal in this paper is to build a unified account of how exactly retractions affect a conversation, and to address these two blind spots. We argue that retractions can be analysed, in a stalnakerian fashion, as *update proposals*. If accepted, retractions cancel a previous utterance's contribution. Nonetheless, they do not cancel the conversation that happened between this utterance and the retraction. We explain how this characteristic effect takes place. In section 5.2, we briefly cover extant accounts of retractions. Then, in section 5.3, we argue for understanding retractions as update proposals. In section 5.4, we explain the update that a retraction operates on the Common Ground, defined as in section 2.2.2, if accepted. In section 5.5, we account for the status of speech acts that occur between a target utterance and its retraction.

5.2 What Retractions undo

5.2.1 Retractions and assertions

Discussions of the speech act of retraction emerge as an attempt to shed light on the norms of assertion. Or, more precisely: if we assume that speakers aim at least somewhat towards truth when they make assertions (Weiner, 2005), we have to take into account the contexts in which they make assertions to evaluate the truth of their statements. Context is relevant to resolving indexicals (Kaplan, 1978, 1979), taste predicates (Lasersohn, 2005), or even, according to metaethical expressivists, moral predicates (Blackburn, 1984; Gibbard, 2003). *Contextualism* is the view according to which one evaluates the truth of assertions with respect to their context of utterance.

MacFarlane (2011a, 2014) argues instead for another type of contextualism: assessment-relativism (earlier, non-indexical contextualism (MacFarlane, 2009)). *Assessment-relativism* is the view according to which one evaluates the truth of assertions with respect to a context of assessment. The reference of some indexicals to elements of the context of utterance may be preserved; however, assertions are evaluated with respect to the context where the assessment is made.

One of MacFarlane's arguments relies on retractions. He claims that if a speaker made the assertion (105) at time t , and this assertion is false when evaluated from context of assessment t' , she is required to retract her assertion.

(105) Carroll (Tuesday): My car is done for.

Suppose that Carroll makes the assertion (105) on Tuesday. On Tuesday, their car is in a ditch, three flat tires, not starting whatsoever. On Tuesday, Carroll, to the best of their knowledge, is asserting a true statement. However, on Friday,

a mechanic has examined the car, huffed and puffed, and declared that it can be repaired. On Friday, if Carroll were to utter *My car is done for*, they would be asserting something false. Then, according to MacFarlane, Carroll *has to* retract their assertion of (105) on Friday.

- (106) a. Carroll (Tuesday): My car is done for.
 b. Carroll (Friday): Oh, actually, the mechanic said it was salvageable.

In this example, Friday is the *context of assessment*, distinct from the *context of utterance* (Tuesday). MacFarlane (2014) concludes: since Carroll has to retract their assertion if it is false on Friday, the truth of assertions is relative to the context of assessment, not the context of utterance.

Conversely, Marques (2014, 2018) maintains that retractions are not obligatory when an assertion is false at a context of assessment. And thus, contextualism still has some good days ahead. She takes specifically the example of taste assertions such as (107) (Marques, 2014).

- (107) a. Malena (5 years old): Bubblegum ice-cream is the tastiest!
 b. Malena (75 years old): Pistacchio ice-cream is the tastiest!

Both the statements that Malena asserts (107a,b) are true in their context of utterance. In the context of assessment where Malena is 75 years old, the statement she asserted in (107a) is false. An assessment-relative account predicts that Malena, at 75 years old, would be required to *retract* the assertion she made at 5 years old. According to Marques (2014), in most conversations, Malena would not be required to retract her previous assertion after her change of perspective. Marques (2018) develops her claim further. By reviewing different arguments in the literature (Knobe and Yalcin, 2014; Ross and Schroeder, 2013), she highlights a second way in which retraction and falsity come apart:

- (i) It is permissible to retract a past assertion that is true, and
 (ii) It may be appropriate to retract a true assertion for reasons that do not directly concern truth. (Marques, 2018, p. 3355)

Therefore, she concludes, retractions aren't mandatory when an assertion is false in a context of assessment; and falsity at a context of assessment isn't necessary to retract an assertion. This counters the retraction argument for assessment-relative truth.

5.2.2 Retractions as a speech act

Our interest in retraction is partly motivated by these debates, but our goal is to study retraction as a speech act for its own sake. Like Caponetto (2020), we treat

retractions as a self-standing speech act in their own right, not restricted to enforcing norms of assertions. Retractions target many different types of utterances, including non-assertoric speech acts (108, 109, 110).

- (108) a. Context: On a professional video call.
 b. Oh! Is that a cat?
 c. Please ignore that.
- (109) a. Context: In a taxi.
 b. To Dam Square, please.
 c. Oh, actually, can you drop me off at Leidseplein?
- (110) a. Context: a parent to their daughter throwing a tantrum.
 b. Yes, you will get a pet octopus if you behave.
 c. Well, never mind.

The retractions in examples (108, 109, 110) retract, respectively, a question, an order and a promise. A speaker can even target for retractions contributions that are mostly performative. Speakers can retract their thanks and welcome, albeit in a more comical manner as these speech acts have little in the way of illocutionary content.

- (111) The situation looks grim, but then the sky starts to flash. Sawyer turns to the sky and shouts "Thank you God!", only to find that they've flashed in the middle of a tropical storm, still in their canoe. Cue Sawyer shouting at God "I TAKE THAT BACK!" (COCA: Web 2012: *You Just HAD To Say It - Television Tropes & Idioms*; <http://tvtropes.org/pmwiki/pmwiki.php/Main/YouJustHadToSayIt>)

Caponetto (2020) characterises retractions as a *second-order* speech acts. *First-order* speech acts, like assertions, questions, commands, occur when a speaker targets a content with their speech act. In chapter 2, I argued for treating rejections as a first-order, primitive speech act that needs not target another speech act. By contrast, *second-order* speech acts like refusals, retractions, occur when a speaker targets another speech act (Caponetto, 2017). *Second-order* speech acts presuppose that a speech act that they can target was performed. For example, Caponetto (2017) characterises refusals as negative answers to *open calls*. When a speaker makes a genuine request, she writes, they invite an illocutionary response, like *acceptance* or *refusal*. The deontic structure of the conversation is such that the hearer is entitled to respond to the open call in virtue of being the addressee. And *acceptance* or *refusal* are illocutionary act that take as their target the speech act they respond to: the request itself, and

not only its content. Similarly, *retractions* take as their target a speech act *qua* speech act.

Thus, Caponetto (2020) characterises retractions as the speech acts by which one *undoes* a previous utterance. Retractions belong in the family of second-order speech acts for undoing, with annulments and amendments.

5.2.3 Typology of undoing

Annulments are used to show that a previous speech act was vitiated: it was not correctly performed.

The Annulment Strategy applies to fatally infelicitous acts that were mistakenly taken as felicitous and whose deontic effects were deemed to be binding for a while due to the ignorance of the involved parties. (Caponetto, 2020, p. 2404)

Caponetto takes the example of annulling a marriage ceremony officiated by an unqualified minister. In the sense of Austin (1975), some speech acts may fail to fulfil their intended effect because of an *abuse*. While everyone thought that the speech act was being performed correctly, it actually contained a fatal flaw that the participants failed to notice, rendering its normative effect as good as void. Cases of abuse generally involve a speaker not having the relevant authority to perform a speech act *and* the audience thinking they do have the authority.

- (112) a. Context: Peter reaches for a book, thinking it's Susan's, but it actually belongs to Antonio.
- b. Peter: Hey, Susan, can I borrow this one?
- c. Susan: Yeah, sure.

In (112), Susan issues a permission that she has no authority to give – as permissions to use a property can only be issued by the owner (here, Antonio) or someone authorised by the owner to issue permissions. The speech act of giving permission that Susan performs is *void*. When a speech act like (112) is fatally infelicitous, it can be annulled:

- (113) a. Context: Peter reaches for a book, thinking it's Susan's, but it actually belongs to Antonio.
- b. Peter: Hey, Susan, can I borrow this one?
- c. Susan: Yeah, sure.
- d. Antonio: That's actually my book, and I am still reading it, so certainly not.

Amendments are used to change the normative burden of a speech act. When speech acts are performed, speakers and addressee incur a certain amount of commitments and associated normative expectations. A command that the speaker has the authority to give compels the addressee to a certain course of action. A command is a closed call, to which the addressee responds by an action. On the other hand, an advice puts a different normative weight on the addressee. An advice is an open call, that the addressee may respond to with an illocutionary act, like acceptance or refusal. A speaker that is in position to do so may adjust the illocutionary force of an utterance. She could weaken or strengthen it, or even change completely the type of speech act.

- (114) a. You should find a place to hide and wait it out.
 b. But I can help!
 c. Find shelter and stay out of trouble, and that is an order.
- (115) a. There is no way I will make it to the top of this mountain alive.
 b. Or is there?

In (114), the speaker strengthens the illocutionary force of their advice (114a) to a command (114b). In (115), the speaker amends their rejection to a question.

Retractions target speech acts that were felicitously performed. Caponetto defines general felicity conditions for retractions:

General felicity conditions of retraction: a speech act A performed by a speaker S at a time t may be retracted at a later time t_n only if

- i. A was felicitously performed at t ;
- ii. a. A was performed by the retractor, or
 b. the retractor is a third party whose position grants her the right to take back A . (Caponetto, 2020, p. 2409sq)

5.2.4 The unique effect of retractions

Caponetto (2020) then describes the general effect of retractions. Retractions target the *illocutionary* effects of the target utterance. They aim to undo the information change and deontic changes that the utterance had brought about. As seen in section 5.1, in particular example (104), retractions do not make it the case that the utterance they target was never performed. They are not a magic eraser.

- (116) a. Context: a parent to their daughter throwing a tantrum.
 b. You will get a pet octopus if you behave.

c. Well, never mind.

In (116), the speaker made a promise (116b) in a context t – say, before the daughter threw a tantrum. Because they made a promise, which was accepted by participants in the conversation, the speaker updated the information structure and the deontic structure of the conversation. In stalnakerian terms, this can be described thus: they added to the Common Ground the information *if she behaves, then daughter gets a pet octopus* (section 2.2.2). The speaker also endorsed a commitment to respect their (conditional) promise that their daughter would get a pet octopus. In (116c), the speaker performs a retraction in context t' – the middle of the tantrum. By doing so, they wish to cancel their commitment to getting their daughter a pet octopus. If the cancellation succeeds, the audience (the daughter) will not hold the speaker to their commitment anymore. However, retractions have no power over the perlocutionary effects of utterances: what an utterance caused. In (116), the speaker cannot cancel the excitement of their daughter, her belief that she would get a pet octopus, and so on.

Retractions differ from annulments in that they target a speech act that was felicitously performed. As seen in section 5.2.1, they can occur when the speaker no longer stands by an assertion that they felicitously performed. But they can also occur in contexts where a speech act would still be felicitous. A speaker can retract a permission that they still have the authority to give, or a command, even an assertion that they are still committed to but don't deem relevant.

- (117) a. Sure, you can borrow this jacket.
 b. Hey, actually I'd prefer that you stop wearing my jacket.
- (118) a. Please go over this contract.
 b. Oh, the deal fell through, no need to review the contract anymore.
- (119) a. I'm quite hungry.
 b. Aah, scratch that, doesn't matter.

Retractions differ from amendments in that they actually cancel a speech act, while amendments modify its normative structure. Caponetto (2020, p. 2411) gives, as the example of an amendment, a speech where a speaker amends a request to a plea:

- (120) a. I have an impossible request of you. Is there any way that you could return to the old woman and bring the girls back here to us?
 b. Madame Archer, I'm begging you, (*Hotel Rwanda*, 2004)

Amendments *modify* the force of a speaker's previous utterance. By doing so, they also perform a second speech act that replaces the previous utterance. In

this case, the normative effect of the request is modified to the normative effect of a plea.

By contrast, retractions may occur *bare* (Vermaire, 2020): they do not necessarily replace the effect of the speech act they cancel with another speech act. In examples (117, 118, 119), the speakers do not replace the speech act they previously made. They retract its effects, but they do not add new information, or new normative commitments to the conversation. The speaker in (117) cancels the effect of their previous permission; but they do not edict a new interdiction. Rather, they bring back the preexisting deontic structure – where the addressee was not permitted to use their jacket. The speaker in (118) does not give a new command, or change their command to a request. Rather, they release the addressee from the obligations their previous command gave. The speaker in (119) does not make a new assertion, or comment on their previous assertion.

5.3 Retraction as a proposal

Our goal is to characterise the essential effect of retractions on conversations. Caponetto takes it to be that retractions cancel the illocutionary effects of a previous utterance. While we somewhat agree with this perspective, we want to clarify two points. First, not all retractions are successful. We challenge the view that retractions immediately undo the effect of the target utterance in a conversation (Caponetto, 2020; MacFarlane, 2014; Marques, 2018). Instead, we argue that retractions should be understood as a proposal to update the conversation. Second, we somewhat agree with Krabbe (2001) by claiming that a speaker retracting an utterance can affect the status of a following argumentation. We give a more precise account of how retractions target the illocutionary effect of an utterance and modify the conversation. This way, we clarify the status of utterances that occur between a target utterance and its retraction.

To address these two points, we adopt a (thoroughly idealised) Common Ground model of conversation. This model, introduced in section 2.2.2, includes a conversational record and the various commitments of speakers in the conversation. These three assumptions determine a public structure of conversation that is available to all participants. So the notion of Common Ground that we use, while being somewhat Stalnakerian in spirit, explicitly contains a deontic picture of the conversation and a conversational record. We then model the effect of retractions within this picture.

5.3.1 Retractions failures

Extant accounts of retractions characterise the essential effect of retractions as the undoing of previous commitments. When the speaker retracts an assertion, she

undoes her previous conversational update.² The first issue we take with extant accounts of retractions is that they take retractions to immediately affect the conversation. A common thread between MacFarlane (2014), Marques (2018) and Caponetto (2020) is that, while they consider that retractions can be *required* of a speaker, they do not mention how retractions can *fail*. For Caponetto, provided that the retractor has the authority to retract the target utterance, the effect of a retraction is immediate. We argue that such is not the case. Instead, retractions need to be accepted by participants in the conversation.

Our first argument is that speakers do not get to retract utterances they made in a conversation without the other participants being present.

- (121) Suppose that Heather, in conversation with Jack, made a promise that she would repair his car. Later on, Heather retracts her promise in a conversation with Jonah.

a. I know I said I would repair Jack's car, but I really do not want to.

As long as she has not retracted her promise *so that Jack is aware of this retraction*, Jack will still consider that she committed to repairing his car.

Retractions that occur in private do not successfully affect a conversation.

Secondly, and more crucially, retractions obey the usual norms of conversation; in particular relevance (Grice, 1975). Recall Marques' example:

- (107) a. Malena (5 years old): Bubblegum ice-cream is my favourite!
 b. # Malena (75 years old): I was wrong to think bubblegum ice-cream was tasty.
 c. Malena (75 years old): Pistacchio ice-cream is my favourite!

This is the core example that Marques uses to reject MacFarlane's claim that any assertion false at the context of assessment ought to be retracted. She sums it up in 2018:

If one's dispositions towards cartoons, or food, change across time, one's later responses and reactions are not more veridical than one's earlier responses and reactions. *A fortiori*, one has no obligation to retract one's earlier assertion, which was correct when made and complied with the constitutive Truth Rule of assertion, at a later time. (Marques, 2018, p. 3346)

We want to add to this analysis. A retraction in the case of (107) strikes us as bizarre. This effect is not only due to the nature of taste predicates. Yes, it is to be expected that Malena's taste may have varied somewhere in the past 70

²Section 5.4 will explain more precisely how a previous update is cancelled.

years. But if Malena were to retract, unprompted, the statement she made at 5 years old, other participants in the conversation would be very puzzled, and reply along the lines of: *what are you talking about?, this was so long ago*, etc. The participants' puzzlement is not only due to the fleeting nature of taste. Malena's retraction would appear *irrelevant* in the conversation.

This observation brings us to the core of our defence of retractions as proposals: like other speech acts, retractions can be rejected by other participants in the conversation (Vermaire, 2020). If a retraction is an attempt of the speaker to be *off the hook* from previous commitments, there are cases where other participants are able to *refuse* and force the speaker to keep their commitments. Hence, retractions do not always perform their intended effect on the Common Ground.

5.3.2 Rejected retractions

Let us see how participants reject retractions. To properly make a promise, a speaker needs to intend to fulfil the course of action they are promising. Similarly, when they make an assertion, a speaker can be held responsible for it: they can be expected to have a justificatory procedure for what they asserted, and other participants may refer to them when they repeat their assertion. Now, retractions are the means by which one undoes such deontic commitments. But the other participants might not agree with a speaker thus forgoing their responsibilities:

- (122) a. Context: A parent to their daughter throwing a tantrum.
 b. Yes, you will get a pet octopus if you behave.
 c. Well, nevermind.
 d. Child: But you promised!

What happens when the child rejects their parent's retraction in (122)? The parent has made a promise. Their retraction aims to cancel the deontic effect of their promise on the Common Ground. If their retraction is successful, their daughter will not expect them to buy a pet octopus anymore, will not believe that they broke a promise if they fail to do so, etc. However, in (122d), the child rejects their retraction. The daughter is not willing to let go of the promise that the parent made, and will not let go of their deontic commitment.

If retractions could not be rejected, the utterance of (122d) would be infelicitous, as the illocutionary effects of the promise would already have been cancelled. Such is not the case.

In (122), because of the power imbalance of the participants in the conversation, the child's rejection can be ignored by the parent. So the retraction operates purely at the level of the conversation. The daughter does not have the power to actually compel the parent to follow through on their promise. However, the daughter signals that, within the conversation, she still holds her parent to be

committed to their promise. She also signals that the parent not following through on the promised course of action is considered a breach of commitment.

By contrast, there are many conversations where the participants have more or less equal power and where they can, by rejecting retractions, compel a the speaker to follow through on their previously announced course of action.

- (123) a. I take back what I said yesterday, it was harsh.
b. A bit too late for that.

- (124) a. I wish I could fly. Oh, please ignore that.
b. No, no, please tell me more.

The result of rejecting a retraction in (123, 124) is to hold the speaker to the commitment she wanted to cancel. In (123), her audience holds her responsible for the harm her words caused. In (124), her audience requires her to justify and elaborate on an assertion. The actual result of the rejection depends on the social dynamics that govern the Common Ground management (section 5.3.3). In a setting where the rejector has power over the conversation management – they have a say in what goes in the information pool, either because their agreement is required or because they are an authority over the conversation –, they are able to keep the illocutionary effects of the retractor’s previous utterance in the conversation. They are also able to compel the retractor to follow on the course of action that they committed to. The rejector in (123) holds the retractor to the effects of their previous utterance, including the undesirable harshness. The rejector in (124) blocks the retraction from taking effect, and encourage the (would be) retractor to elaborate on their previous assertion.

That retractions can be rejected shows that they do not immediately affect the conversation. Instead, whether a retraction is successful depends on the other participants in the conversation accepting it. We argue that retractions are an update *proposal*, in a stalnakerian fashion. This approach has three effects.

First, if a retraction is modelled as a proposal to update the Common Ground, we account for rejection of retractions in a standard way. Rejection is the speech act by which one *blocks* or *prevents* an update of the Common Ground. In chapter 2, I focused on the conditions for rejections of statements, and established that speakers may reject statements that they do not have grounds to assert. However, when understanding rejection as blocking an update to the Common Ground, speakers can also reject questions, orders, promises (see sections 2.2.3 and 2.2.4).

Secondly, if a retraction is a proposal, then the second order speech act of retraction behaves similarly to other speech acts, like assertions, that are also proposals. This situates retraction in an egalitarian picture of speech acts. Instead of having to postulate that some speech acts may directly affect the conversation

while others require acceptance from the other participants, we can understand most speech acts as proposals. However, the way power dynamics affect the conversation can modify what is required for assent. This is a topic I develop more in detail in chapter 6, but that I also investigate in section 5.3.3.

Thirdly, if a retraction is a proposal, then retraction has a double effect in conversation. We model it the following way:

- i. When a speaker performs a retraction, it is first added to the conversational record in the Common Ground that she proposed an update.
- ii. Then, if the proposal is accepted, the retraction modifies the Common Ground according to its unique effect. We model the unique effect of accepted retractions in section 5.4.

As with the second point, this lines up retraction with other speech acts like assertion, and allows us to treat retraction as a speech act in its own right.

5.3.3 Retractions and Power Dynamics

Caponetto (personal communication) raises a counterexample to the treatment of retractions as proposals. She focuses on retractions of directives, as in (125):

- (125) Take a general who ordered her soldiers to attack at dawn. During the night, she changes her mind, realising that an attack might have unforeseen, disastrous consequences. Therefore, she retracts her order.

Do the soldiers need to accept this retraction for it to take effect?

Caponetto argues that such a retraction need not be accepted to take effect, and directly undoes the commitment that the general had put on her soldiers. Therefore, to free her soldiers from this commitment, the general does not need their approval.

It is our opinion that this counterexample points to interesting problems about how power relations can affect conversation management. I dabble in them in chapter 6. Here are some notions specific to retractions, and Caponetto's counterexample.

Idealised Common Ground models suppose that participants have equal say on what updates are performed, what commitments are put on everyone, etc. In everyday dialogue, the situation is more complicated. What we can assume is that, in our models, the Common Ground represents the set of possibilities that are not excluded by what the participants agree upon. In turn, this agreement requires that the participants have somehow (implicitly) agreed upon some rules that determine how the set of possibilities is managed, and who has the last say in their agreements. In this sense, participants engage in a joint project to manage the Common Ground.

But not all conversations are such that participants have equal power over the conversation. In a court setting, the judge has the last word on what goes (or not) in the (literal) public record of the conversation. A judge can declare that an objection is sustained, compelling one of the parties to retract their statement. We analyse as retractions these cases. That the objection is sustained and the speaker has to retract means that the target statement should be disregarded by the witness and jury and everyone. However, the fact that the statement was made in the first place still makes it to the transcript. In this sense, a judge sustaining objections compels a retraction, not an annulment. The judge can also reject objections. Effectively, she manages the conversational information, by deciding what will and will not be considered.

The role of power dynamics in conversation management requires nuancing the (admittedly very weak) stalnakerian notion of acceptance. For Stalnaker, any proposal that is not rejected is accepted, and modelled as an update of the Common Ground. In some contexts, acceptance is indeed a default attitude that can be attributed to the participants. If the audience does not voice a rejection, the update proceeds. However, in many contexts, the opposite is true. When a student defends a claim to her teacher during an examination, she will assume that her claim is insufficiently justified and thus rejected until it is explicitly accepted. When a parent tells their child to clean their room, they are likely to assume that their child remaining silent is not a good thing. That assent cannot be presupposed is consistent with the way speakers make use of multiple paralinguistic cues to indicate acceptance (or rejection) in conversational settings. I develop this analysis of silence in non-cooperative contexts in chapter 6. For now, it suffices to say that an absence of rejection does not always guarantee assent; but in certain contexts, it does.

This hypothesis underlays the treatment we propose of Caponetto's example. Caponetto assumes that the soldiers do not need to explicitly accept the retraction in virtue of the features of retraction. Because the general is retracting a commitment she put on the soldiers, this retraction takes effect immediately. The general does not need the soldiers' assent to free them of their commitments. Instead, we offer a treatment of the counterexample where the general's retraction is still a proposal. In virtue of a context where the general has authority over which updates enter the Common Ground, she does not need the soldiers' explicit assent: she gets to decide what goes on the soldiers to-do list. But the context can be changed such that the soldiers' presupposed acceptance of the general's retraction is not automatic. For example, if another general already started preparations for the joint attack, they may be able to reject the general's retraction: *My cavalry has already positioned themselves for the ambush, if you don't attack at dawn, they will be slaughtered!*

So we answer Caponetto's counterexample: yes, the soldiers need to accept the general's retraction for it to take effect. In some contexts (the military command line is one of them), the soldiers' acceptance can be presupposed. Since the

general is ultimately the one who decides what the soldiers are committed to, she can assume that her retraction goes through. However, this is a feature of the context, and the way conversation is managed in it; this is not a feature of retraction.

5.4 Retraction and Common Ground

Retractions aim to remove an information from the conversation, by *undoing* the effect of the speech act that brought this information in the conversation. We represent the effect of retractions in a Common Ground model. Take a simple example:

- (126) a. Shannon: It is raining. You should grab an umbrella.
 b. Albert, trusting Shannon that it is raining, proceeds as she advised.
 c. A few minutes later, Shannon looks out of the window.
 d. Shannon: Sorry, I was wrong. It is not raining. My bad.

By asserting *It is raining* (R), Shannon makes a proposal to update the Common Ground with R . The fact that she made a proposal is consigned in the conversational record. If Shannon's assertion is accepted by the participants, it results in updating the Common Ground by removing all non- R worlds (fig. 5.1). Shannon also endorses certain commitments: that, if challenged, she would provide a justificatory procedure for her statement, and that other participants in the conversation may take the content of her assertion as a premise in deductions. In this case, Albert might use her assertion as a justification for taking an umbrella. These commitments become part of the information that participants share.




Figure 5.1: Updating the Common Ground with an assertion that R

When Shannon *retracts* her claim that it is raining, in (126d), she proposes to *undo* the illocutionary effect of her utterance (Caponetto, 2020). She performs a retraction by uttering *Sorry, I was wrong*. To successfully retract her assertion

that *it is raining* is a necessary condition for the continuation of her discourse: when she says *It isn't raining*, she proposes a further update of the Common Ground, to remove all *R* worlds (fig. 5.2). If her retraction had not recovered some *R*-world to exclude, this further update is impossible, as it would result in an inconsistent Common Ground.

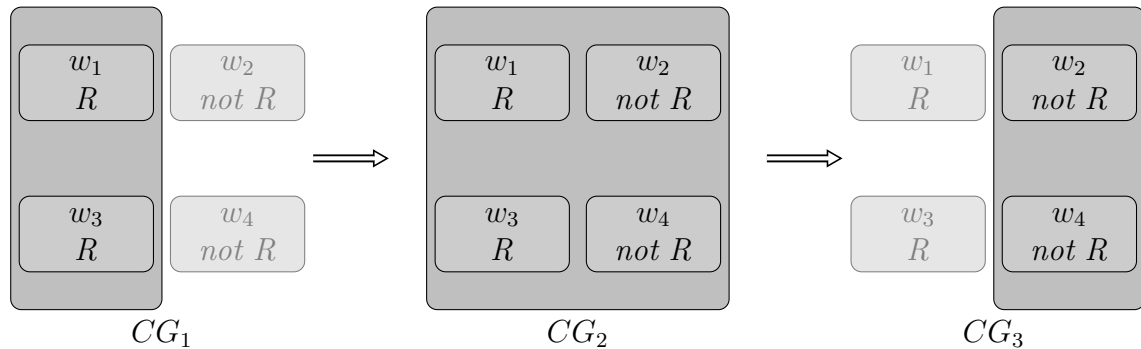


Figure 5.2: Retracting an assertion that *R*, and asserting that *not R*

In technical terms, a retraction is a proposal to *recover* some previously excluded worlds, and re-establish their compatibility with the participants beliefs. How does that work? A naive, only temporal, picture of the retraction in (126) would be that Shannon proposes to recover the Common Ground as it was before her utterance. Things are not that simple.

- (127)
- a. Shannon: It is raining. You should grab an umbrella.
 - b. Albert, trusting Shannon that it is raining, grabs an umbrella.
 - c. Albert: Would you look at that! Our umbrella is blue!
 - d. Shannon (after looking out of the window): Sorry, I was wrong. It is not raining. My bad.

In (127), after Shannon's assertion that *it is raining* (*R*), Albert asserts *Our umbrella is blue* thereby updating the Common Ground with *B* (fig. 5.3). If the effect of Shannon's retraction, if accepted, were just to recover the Common Ground as it was before her utterance, it would result in the first information state (CG_1). But in CG_1 , we do not have the subsequent information that Albert introduced: *B*. CG_1 contains both worlds where the umbrella is blue and worlds where the umbrella is not blue. So we need to refine the notion of *cancelling* the effects of a speech acts further than previous authors.

In (127), Shannon's retraction affects CG_3 , after Albert assertion has removed *not-B*-worlds (fig. 5.4). If the effect of her retraction were just to recover the Common Ground as it was before her assertion: CG_1 , it would also undo the effect of Albert's assertion. But intuitively, Shannon's retraction of the claim that *It is raining* should not undo the effect of Albert's assertion *Our umbrella is*

blue. The intended result of Shannon's retraction, if accepted by the participants, ought to be to reintroduce in the Common Ground *not R*-worlds that are also *B*-worlds: CG_4 in fig. 5.4.




Figure 5.3: Toy example: assertion that R , assertion that B




Figure 5.4: Toy example: retracting an assertion that R after an assertion that B

To model the more fine-grained effect of retractions, we use the fact that the Common Ground also contains a conversational record of the speech acts that acted upon it. Stalnaker (1978) divides the effect of an assertion in two. First, when a speaker makes an assertion, the participants in the conversation add to the conversational record the information that the speaker performed an assertion (i.e. made a proposal to update the Common Ground with a propositional content). So the fact that a proposal was made gets in the conversational record whether the assertion is accepted or rejected. This is the first effect of an assertion. Then, if the participants accept the assertion, they update the Common Ground by removing worlds incompatible with its content.

The first effect of Shannon's assertion: *it is raining* in (127) was to add to the Common Ground the information that she performed an assertion that it is raining, that she is committed to the proposition *it is raining*, etc. The second effect, after Albert accepted Shannon's assertion, was to remove worlds incompatible with the proposition *it is raining* (R) from the Common Ground. We represent this in fig. 5.5 by associating a conversational record (CR) to each step of the conversation. Assuming that the conversational record starts empty (no previous utterance has been made), it then records Shannon's assertion of R with both a proposal ($Prop : Assert_S(R)$) and the update it causes ($update : (R)$).

This modelling allows us to emphasise the notion of *proposal*. Moreover, we can account for a more fine-grained analysis of the effects of retractions.

As we claimed, a retraction amounts in a Common Ground model to a proposal to go back to a previous utterance and remove its essential effect from the Common Ground. Shannon, when she utters *Sorry, I was wrong* in (127), proposes to retract her assertion that *it is raining*. If Shannon's proposal is accepted, then the participants backtrack on the conversational record to find the target utterance. They jointly cancel the effects of the target utterance: that is, the participants recover the state of the Common Ground before to the utterance. However, that the target utterance was made is not erased from the conversational record. As such, Shannon is freed from her commitments, but it is not as if she never undertook them.

After recovering the Common Ground that predates the update of the target utterance, the participants perform again the updates that came after Shannon's assertion, *if they are compatible*.³ In effect, they go through the conversational record again and recheck the different conversational steps. This is represented in (fig. 5.5). If Shannon's retraction is accepted, the participants first recover the Common Ground as it was before Shannon's update (CG_1). Then, they re-perform the compatible subsequent updates. In our example, they re-do the update with the content of Albert's assertion *Our umbrella is blue* (B). That gives us the intended result of a retraction (CG_4).

³We explore in section 5.5 what happens if a subsequent update is incompatible with the retraction, i.e. depends on what has been retracted.

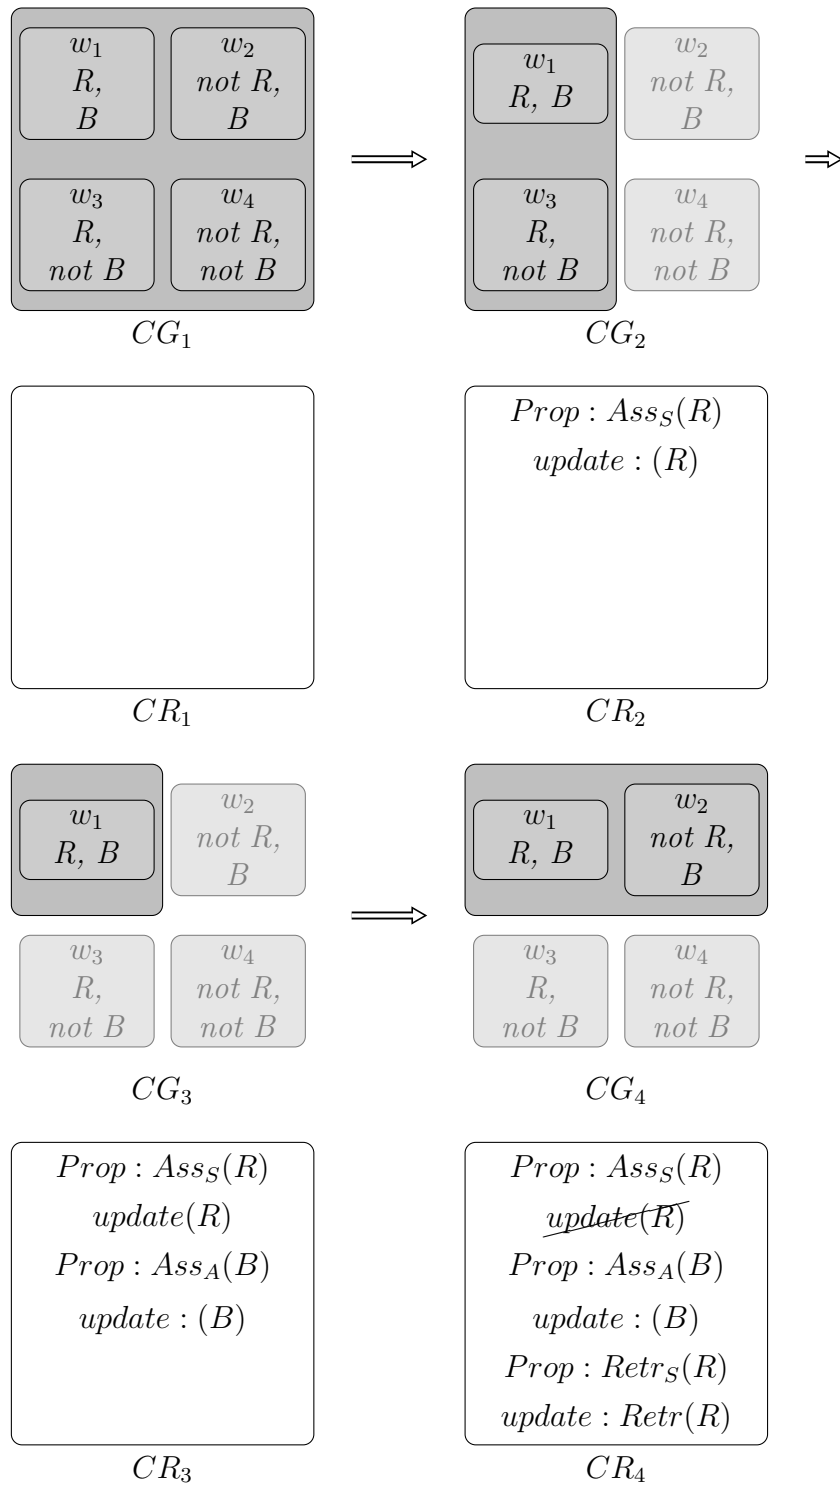


Figure 5.5: Toy example: the effect of retractions on Common Ground and Conversational Record

This picture may seem unnecessarily complicated and costly. But the reason for running through subsequent updates, and not only cancelling the effects of the target update is that speakers need to check for following updates depending on the target utterance. For example: suppose that after building a complete reasoning and plan on the assumption that it is raining, Shannon retracts her claim that it is raining. Then, the speakers might want to keep some parts of the conversation (the choice to wear warm clothing), as they are also deduced from extra premises; while other parts of the conversation (the choice to take an umbrella) depend entirely only on the retracted premise, and so should also be cancelled.

The fact that retractions require significant backtracking explains, for us, the failure of retractions that occur too long after the utterance, as in Marques' (2014) ice-cream example.

- (107) a. Malena (5 years old): Bubblegum ice-cream is my favourite!
 b. # Malena (75 years old): I was wrong to think bubblegum ice-cream was tasty.
 c. Malena (75 years old): Pistacchio ice-cream is my favourite!

Moreover, it explains why retractions are a discourse move that is *costly* for a speaker. Not only do they have to admit that one of their interventions was misdirected, but they also have to resolve the effects of this intervention together with the audience. And, the more hinges on the target utterance, the more likely the audience is to resist the retraction.

5.5 Problematic Retractions

Given this picture of how retractions work, we can now investigate some specific cases of problematic retractions and explain them.

5.5.1 Retractions and subsequent utterances

At the beginning of this paper, we set out to investigate the status of speech acts between a target utterance and its retraction. We have modelled what happens, when a retraction occurs, to utterances that were made between a target utterance and its retraction. The subsequent utterances are checked for compatibility following the cancellation of the target utterance. If they are compatible, the participants re-update the Common Ground with them. If they are incompatible, then the participants hit a snag and have to resolve the issue. So, for example, if a further update cannot hold without the retracted speech act, it needs to be undone *before* the retraction can succeed.

- (128) a. I am getting you a pet octopus.

- b. Oh sweet! we'll call her Luna.
- c. Actually, I looked into it and octopi are really hard to care for. We won't be getting one as a pet.

In (128), the update performed by naming the future pet Luna *depends* on the promise *I am getting you a pet octopus*. When the speaker retracts their promise, the naming update is voided. Thus, the participants also have to undo the naming update.

While in (128), participants resolve the retraction seamlessly, it is important to remember that retractions undo *only* the commitments introduced by their target utterance and its closure under logical consequence, not the ones that rely on it. Take the example of retracting the premise of a deductive reasoning. In (129), the retractor Sirin undoes the effect of her utterance. In particular, she undoes the warrant that her utterance may be used as a premise in deductive reasoning. Therefore, her audience is also led to undo the (flawed) deduction that hinges on the retracted premise.

- (129)
- a. Sirin: If it is sunny, you need to wear sunscreen.
 - b. Refke: Well, I live in the Netherlands, so I don't really need sunscreen.
 - c. Sirin: Oh, actually, dermatologists recommend wearing sunscreen even on cloudy days.
 - d. Refke: Urgh. I guess I do need sunscreen after all. Do you have any sunscreen to recommend?

Retractions undo retractor commitments only, not the commitments that others may undertake on the basis of the retractor's utterance. In (129), Refke has to retract her assertion (that she doesn't need sunscreen) separately from Sirin's retracting the justification of her assertion. This also applies to presuppositions:

Giorgio tells Julian about his pet: *I have a cat, Norma, at home*. Julian is then licensed to refer to Giorgio's cat with a definite description when speaking to Leila: *Giorgio's cat is named Norma*. But if Giorgio retracts the assertion that he has a cat (*No, I meant I am cat-sitting Norma, she's not my cat.*), Leila may point out that Julian's assertion had a presupposition failure: *Giorgio does not have a cat*.

Even if Giorgio has retracted his assertion, and thus the justification for Julian's commitment to the proposition *Giorgio has a cat*, Julian is still committed to the proposition. Julian has to undo his own commitments separately.

But there are also cases where the commitments linked to a target utterance cannot be retracted, because the illocutionary effects and commitments of the target utterance have been entirely resolved. Take the example of a marriage proposal. If the retraction happens a few days after the proposal, it can be

accepted into the conversation, undoing the commitment of the proposer to marry the proposee. On the other hand, after the wedding ceremony, the proposal cannot be retracted anymore, as the wedding ceremony is a subsequent speech act that resolves the effect of the proposal entirely. In the same way as it wouldn't make sense to retract a promise *after* having fulfilled it (and therefore resolved the associated commitments), it doesn't make sense to retract a proposal after the wedding that solves it.

5.5.2 When retractions go wrong

We also suggest an analysis for cases where a speaker wishes to retract, but the audience rejects it, and neither has enough power to make their decision final.

Vermaire (2020) builds just such an example of retraction to challenge the assurance view of testimonial justification. He defines the Assurance view he addresses as:

the claim that, in order for someone's belief to be testimonially justified, some testifier must be "on the hook for it" under the illocutionary norms that govern speech. (Vermaire, 2020, p. 3960)

The example he builds goes this way:

Die-Off Zane, an environmental scientist, investigates a recent trout die-off in Kingfisher Lake, and at the end of his investigation he communicates his conclusion to the local newspaper: "The primary cause of the die-off was industrial pollution from the Kemco paint factory." A few years later, Zane is under consideration for a high-paying position as Kemco's chief environmental officer. Eager for the job, he puts out another short statement to the paper, which is doing some retrospective reporting on the event: "I retract my earlier statement about the die-off in Kingfisher Lake." (Vermaire, 2020, p3967)

However, Athavi has formed the belief that Kemco caused the die-off on the basis of Zane's assertion. Zane's assertion is the testimonial justification for Athavi's belief. Zane's retraction, however, leaves him *off the hook* for his assertion. Therefore, an Assurance view would be compelled to say that Athavi's belief is not testimonially justified.

We propose an alternative analysis, grounded on our claim that retractions are proposals which may be accepted or rejected. Athavi knows that Zane's retraction was not motivated by truth, but by other reasons. Therefore, she *rejects* Zane's retraction, on the basis of his non-epistemic grounds to reject his assertion. In this case, while Zane sees himself as not committed to his statement anymore, Athavi keeps using his statement as a justification for her belief (and further assertions of it).

This can be explained as two participants disagreeing on whether to accept a retraction. We can model those cases, where neither participants has the ability of forcing the acceptance (resp. rejection) of the retraction in the Common Ground, as cases where the Common Ground breaks. This is not so dramatic as it sounds. Both participants form their personal copy of the Common Ground. In Zane's copy of the Common Ground, Zane's retraction was successful. He is freed of his commitments to the assertion that Kemco caused the Die-Off. In Athavi's copy of the Common Ground, it is still the case that Zane has issued a warrant to repeat his assertion that Kemco caused the Die-Off and use it in inferences. The retraction has not freed Zane from his commitments in Athavi's copy of the Common Ground. In other words: even if Zane retracts his assertion, Athavi can still hold him *on the hook* for his assertion. This addresses the most dramatic cases of disagreement concerning a retraction. Because a retraction is a proposal, participants may accept or reject it in the conversation. However, in much the same way as with other proposals, like assertions, sometimes participants disagree on whether to accept or reject a proposal; and this conflict can lead them to have conflicting beliefs about the conversation. These conflicting beliefs can be represented as conflicting copies of the Common Ground, similarly to what Camp does in her 2018 article.

5.6 Conclusion

We give a picture of the way retractions affect a conversation. In a model of conversations as updating a set of mutually accepted propositions, retractions were an outlier when assuming that they immediately undid the effect of a target utterance on the Common Ground. On the contrary, we argue that the effect of retraction is to propose an update to the conversation (section 5.3). This is shown by the fact that retractions can be rejected, and need to be accepted to affect a conversation. Moreover, this makes retractions analogous to other speech acts: assertions, rejections, questions, requests. . . do not directly update the Common Ground, but propose an update that has to be accepted by the participants to take effect. Thus, a retraction has two effects: it adds to the conversational record that a retraction was proposed; if accepted, it cancels the illocutionary effect of a target utterance.

Section 5.4 explains, by way of modelling retractions update behaviour, how cancelling the illocutionary effect of an utterance works. We show that felicitous retractions, while they cancel the target utterance's update on the Common Ground, do not affect the subsequent updates. These stay in place as long as they are compatible with the retraction. Section 5.5 explains how a further update can be incompatible with the retractions.

Chapter 6

Silence, Dissent, and Common Ground[†]

6.1 Introduction

You are talking with a friend. Suddenly, your friend remains silent – not only do they not speak back, but they do not engage in other para-linguistic replies. They don't nod, or frown. You take it for granted that your friend heard you, and could speak their mind if they so wished. How can you interpret their silence?

According to a certain analysis of conversation, *Quid tacet consentire videtur, ubi loqui debuit ac potuit*, 'They who are silent, when they ought to have spoken and were able to, are taken to agree'. In Common Ground pictures of conversation, in particular, it is enough that participants voice no objection for an utterance to take its full effect. When a participant remains silent, not as a result of a pause or of physical limitations, but of their own volition, they do not perform a speech act that expresses refusal. This absence of rejection grants weak assent in a cooperative conversation (Goldberg, 2020a). In this sense, 'silence gives assent' for Common Ground accounts (section 6.2).

However, in less cooperative contexts, silence does not guarantee acceptance (Lackey, 2018). In fact, *eloquent silences* (Tanesini, 2018) may be powerful expressions of dissent (section 6.3). This is particularly true in contexts where, while one participant is free to speak, they ostensibly refuse to engage with the conversation.

But if silence does not give assent, can we interpret silence in a structured manner, that is still useful to model conversation? I argue that we can.

The analysis of silence relies on a less idealised picture of conversation. Conversational contexts may be more or less adversarial. Participants might be more or less willing to agree with a speaker, positively answer their requests, etc. Given this variety of contexts, I propose a fine-grained analysis of silence. Crucially, a speaker can assess the cooperativeness of a conversation, based on priors and

[†]This chapter is based on published work (Bussière-Caraes, 2021).

pragmatic cues. Then, they ascribe a default attitude to the audience, which supplies an interpretation for silence (section 6.4). Silence works as the expression of a default attitude of the audience. With this picture of silence, one can observe how speech acts behave in non-cooperative contexts where silence doesn't mean acceptance (section 6.5). This gives full weight to an analysis of speech acts as update proposals: the update may require more than an absence of objection to go through.

While my analysis of silence is fecund for a proposal view of speech acts, it opens another line of questioning. If silence does not yield assent, why is silencing problematic, and how does it impact a conversation? Being serious about non-ideal conversations requires addressing this political point. I do so in (section 6.6).

6.2 Silence and Common Ground

In the classic Stalnakerian picture of Common Ground models (section 2.2.2), an update, proposed by an utterance, is accepted *unless it is rejected*:

To make an assertion is to reduce the context set in a particular way, *provided that there are no objections from the other participants in the conversation* . . . This effect is avoided only if the assertion is rejected. (Stalnaker, 1999, p. 86, emphasis mine)

This *default assent* hypothesis does not give much importance to the notion of proposal, except maybe as a way to differentiate assertions from presuppositions:

. . . normally, after the utterance of “I have to pick up my sister at the airport”, it becomes shared information that the speaker has a sister *even if the statement itself is rejected*. (Stalnaker, 2002, p. 714)

While the information that an assertion contains is not added to the Common Ground if the assertion is rejected, presupposed information is accommodated even if the utterance is rejected. Suppose that the speaker's assertion *I have to pick up my sister at the airport* occurs in a context where it wasn't already known that the speaker has a sister. If the assertion is rejected, it does not become shared information that the speaker has to pick up their sister. However, because the speaker made an utterance that contained the presupposition that they have a sister, the information that the speaker has a sister is added to the conversation.¹ As such, the update behaviour of assertions differs from presupposition accommodation.

Since Stalnaker characterises the update behaviour of speech acts in case they are accepted, it is enough for him to consider a weak acceptance, closer to *illocutionary uptake* (Austin, 1975; Searle, 1969), which assumes that acceptance

¹Under the assumption that the speaker is being sincere, and knows more about their familial situation than the rest of the participants.

is granted unless the participants state otherwise. The assumption is that, if a speech act is rejected, its normal effect does not go through.

Stalnaker's construction of silence as assent is explained by an idealised picture of conversation. Participants are competent speakers, free to speak, and cooperative, in a Gricean manner (Grice, 1975). They share a common goal to harmonise and maximise the shared information. In particular, they strive for sincere, relevant and concise utterances. Furthermore, power dynamics do not muddy the picture, and the audience is not biased against accepting the discourse moves of the speaker. In this scenario, no participant is prevented from expressing their commitments. If a speaker makes an utterance, the cooperative thing to do for other participants is to clearly indicate their assent – or dissent. So they have a responsibility to reject update proposals they do not agree with (they do not wish to commit to) (Pettit, 2002; Tanesini, 2018).

Because of this responsibility to reject proposals they disagree with, participants who remain silent appear to assent. This is in cases where silence is an intentional absence of communication – purposefully excluding cases of nonverbal communication or non-communicative silence (Saville-Troike, 1995). With nonverbal communication, such as writing or gestures, participants still deploy devices of communication and thus separately indicate assent or dissent. With non-communicative silences, like pauses to take a breath or process speakers contributions, participants do not convey an intention. By contrast, intentional silences show that a participant made a conscious decision not to speak.

Pettit (2002) and Goldberg (2016, 2018, 2020a) argue that in the absence of defeaters (power dynamics, silencing, inability of the audience to speak), silence conveys assent. Silence in the absence of defeaters can roughly be considered intentional silence. Not only participants are choosing silence, but this choice is free. When intentional silence follows a speaker's discourse move, the audience lets this move go through. They do not oppose the update to the Common Ground that the speaker proposes.

Silence is interpreted as assent in cooperative conversation, because assent emerges as the least marked response in cooperative contexts:

Assertions project confirmation, and therefore the move of confirming an assertion is the least marked next discourse move. (Farkas and Bruce, 2010, p. 93)

Intentional silence signifies that the participant opts out of performing an utterance. Thus, the least marked discourse move can be attributed to them. In the example, confirming an assertion. Similarly, in a context where a speaker has authority over the addressee, silence is interpreted as assent when it follows requests or orders.

Goldberg (2020a) likens the presumption of assent in idealised cooperative settings to a *tacit acceptance procedure*. In tacit acceptance procedures, the conversation is explicitly structured so that silence expresses assent.

- (130) A committee sets a report to be tacitly accepted unless an objection is raised by May 3rd.
- (131) The officiant instructs the audience to ‘speak now or forever hold your peace’. If they remain silent, the marriage takes effect.
- (132) The teacher tells her students to interrupt her whenever they have a question or objection.

According to Goldberg, the assumption of cooperativeness licenses attributing to the audience a default attitude of assent. In cooperative conversations where the audience is not prevented from expressing their dissent, the least marked conversational move is the move of assent. So, as with tacit acceptance procedures, it is the audience’s responsibility to voice their dissent. In such a setting, intentional silences convey an absence of rejection. With the extra assumption present in Common Ground pictures that assent is weak, and thus absence of rejection signals agreement, silence stands for assent.

6.3 Silent Dissent

What Goldberg, and many others, note is that the cooperative picture is rather idealistic; in real life, the situation is more complicated, and silence does not automatically mean assenting to an update (Lackey, 2018). Tanesini (2018) collates examples of *eloquent silences*: silences whereby the audience expresses something distinct from assent. She follows Saville-Troike (1995), as do I, in only considering intentional silences as possible eloquent silences. In cases of nonverbal communication, or unintentional silence, the silent participant does not make clear an intention to convey something by their silence. Thus, in a Gricean sense, they do not communicate anything by their silence. By contrast, eloquent silences have a communicative role.

Examples of eloquent silences abound: the politically loaded silence of white audiences when confronted to discrimination stories from persons of colour (DiAngelo, 2012); the pointed silence that follows an offensive remark at a dinner party, followed by a topic change; defiant silence faced with a question, or an order, ... Far from being cooperative silences, these eloquent silences signal that the audience refuses to cooperate in the conversation (Tanesini, 2018).

DiAngelo (2012) analyses the case of *white silence*. White silence is enmeshed within power dynamics. Its effects, during antiracist discussions, differ:

Silence has different effects depending on what move it follows. For example, if white silence follows a story shared by a person of color about the impact of racism on their lives, that silence serves to invalidate the story. People of color who take the social risk of revealing

the impact of racism only to be met by white silence are left with their vulnerability unreciprocated. [...] Conversely, when white silence follows a particularly problematic move made by a white participant, that silence supports the move by offering no interruption; in essence, white silence operates as a normative mechanism for these tactics. When white silence follows a white, antiracist stand (such as challenging one's fellow whites to racialize their perspectives), it serves to isolate the person who took that stand. This isolation is a powerful social penalty and an enticement to return to the comfort of white solidarity. In this context, white silence denies the support that is critical to other whites working to develop antiracist practice. (DiAngelo, 2012, p. 5)

The common thread between white silences is that they convey a refusal to engage. This refusal to engage is made salient by the fact that the white audience is expected to participate in the conversation. When a person in power refuses to acknowledge discourse moves from marginalised individuals, or to call out problematic discourse from another person in power, they implicitly agree to maintain the *statu quo* of power dynamics.

The converse also occurs. In imbalanced power dynamics, silence from an inferior party can convey disagreement in a form that is less susceptible to punishment.

- (133) Your boss makes a sexist joke. You are not in a position to call them out on their behaviour. You remain silent, exchanging loaded glances with sympathetic coworkers

White silence, or the silence in (133), differ from other types of eloquent silences that Tanesini (2018) analyses. She takes as an example the silence of a protester who is interrogated by the police about the names of their comrades, or the hostile silence of an objector who refuses to follow an order. The latter silences follow closed calls that require a response from a specific addressee, while the former examples follow open calls that address an audience (Caponetto, 2017). When silence follows a closed call, it immediately expresses dissent or noncooperation. The addressee refuses to answer the call, and thereby brings attention to a situation where they might not be safe outright rejecting the call. Conversely, white silence or the silence in face of a sexist joke (133) follow open calls. Silences that follow open calls show that the audience or part thereof refuses to engage with the speaker. When interpreting open calls as bids for connection, where the speaker is looking for someone to engage with what they say, the silence of the audience reflects a refusal to connect with the speaker.

What unifies cases of eloquent silences, whether they follow open or closed calls, is that they subvert the expectation that the silent participant continue the conversation. When the participant elects to remain silent, they signal that

something is awry with the conversational exchange (Tanesini, 2018). Eloquent silences have a specific illocutionary effect: they indicate that the participant refuses to engage in a conversation. Occurring after open calls, such as assertions, they can evoke disinterest, refusal to share the speaker's commitments, refusal to engage in certain topics. After closed calls, such as questions, or requests, they convey a refusal to answer the speaker, to fulfil their demands.

The person who is deliberately keeping silent, instead, indicates that something is amiss with the conversation which, therefore, cannot continue as normal. It may be worth noting in this regard how awkward silence often is in conversation. Silence is uncomfortable because it often marks the fact that things are not going well with the conversational exchange. (Tanesini, 2018)

Eloquent silences are distinct from the kind of silence at stake in the idealised Stalnakerian picture. They are also extremely common, which makes clear that non-cooperative contexts, where the audience is not assumed to agree with the speaker until they explicitly claim their agreement, are far from being exceptional. In fact, cooperative contexts are more of an exception than a norm (Lackey, 2018). But this does not mean that silence is always an expression of dissent, or uninterpretable. What this paper aims to show is that different readings on silence, that distinguish an absence of objection from assent, can be incorporated in a Stalnakerian framework.

6.4 Silence and default attitudes

My stance builds on Lackey's (2018) observation that cooperative contexts, where silence means assent, are few and far between. I develop this idea that silence is not, by itself, an expression of assent or dissent. Instead, a conversational context supplies a *default attitude* of the audience with respect to the speaker's utterances. When the audience remains silent, they are assigned this default attitude. In an idealised context, Stalnaker style, the audience is assigned a default attitude of assent when they remain silent. And, in the non-cooperative contexts exemplified by Tanesini (2018), the silence of the audience indicates a default attitude of defiance, or dissent, with respect to the speaker's updates.

6.4.1 Dialogue cooperativeness

To categorise the adversarial, or cooperative, levels of contexts, I help myself to the four scenarios of dialogue that Dutilh Novaes identifies for Prover-Skeptic argumentation games:

- (1) Prover and Skeptic have a common goal, that of establishing the validity or invalidity of proofs, and no (conflicting) individual

goals (they either win or lose together). They each perform a different task, but in view of a common interest (or converging individual interests). This is a purely cooperative, division-of-labor game, where neither player can ‘win’ alone; both players will benefit from achieving the overall goal of correctly identifying (in)validity.

- (2) Prover wants her proof to go through no matter what (as this counts as a win for her), regardless of whether it is a valid proof or not. Skeptic, by contrast, wants valid proofs to go through and invalid ones to be refuted, and is neutral with respect to ‘pay-offs’ of the game for him (no win or loss). Here, Prover can win or lose the game, and Skeptic can neither win nor lose (the outcome is neutral for him).
- (3) Skeptic wants to block (refute) the proof no matter what (as this counts as a win for him), regardless of whether it is a valid proof or not. Prover, by contrast, wants valid proofs to go through and invalid ones to be refuted, and is neutral with respect to ‘pay-offs’ of the game for her (no win or loss). Here, Skeptic can win or lose the game, and Prover can neither win nor lose (the outcome is neutral for her).
- (4) At a lower level, the game is a classical adversarial, zero-sum game: Prover wins if the proof goes through, Skeptic wins if the proof is refuted or otherwise blocked. But, at a higher level, they are in fact cooperating to establish whether a proof is valid or not.

(Dutilh Novaes, 2020, p. 55sq)

Dutilh Novaes applies these different scenarii to particular dialogues, whose practice is exemplified in mathematical proofs. However, the taxonomy also extends to casual conversations. I use it to gauge conversations level of cooperativeness or adversariality:

1. **Fully cooperative dialogue:** Speaker and Audience have a common conversational goal. For example, to establish the truth or falsity of a claim; to settle on a movie to watch; to maximise the number of sailing-related puns in their dialogue. They have no conflicting individual goals.
2. **Semi-cooperative dialogues:**² Speaker and Audience conversational goals partially align. Speaker wishes to convince an agnostic Audience to have

²This is in correspondence with Dutilh Novaes’ (2) and (3) scenarii of Prover-Skeptic dialogues, where Prover and Skeptic individual goals conflict in some respect. These scenarii are semi-cooperative in that one of the participants is neutral with respect to the conversational outcome.

a church wedding; Speaker presents their paper to an audience of peers; Speaker questions an Audience decision.

3. **Adversarial dialogues:** Speaker and Audience conversational goals do not align. Speaker debates an hostile Audience; a black Speaker wishes to challenge discrimination in front of a white Audience; Speaker orders their teenage child to clean their room.

Whether the dialogue is cooperative or adversarial impacts how Speaker perceives the silence of their audience. In a cooperative setting, Speaker is likely to interpret their audience's silence as assent. Their conversational move went through. They suggested an action movie, their Audience stayed silent and followed them to the DVD shelf. On the other hand, adversarial settings introduce dissent as a default attitude. When the teenage child stays silent after being asked to clean their room, they express their rebuttal of the imperative. Finally, semi-cooperative settings introduce an interesting read on silence. Speaker answers an Audience's question about their paper. Their Audience remains silent, not expressing assent or dissent. Speaker deduces that they have to provide more information to ground their claim.

6.4.2 Identifying dialogue situations

How the speaker perceives the dialogue situation, whether it is more or less adversarial, causes them to attribute different default attitudes to their audience. In a fully cooperative situation, the tacit acceptance procedure of idealised Stalnakerian models applies. By contrast, in adversarial situations, silence is perceived as hostile, a refusal to let a speech act go through. The perceived cooperativeness is influenced by the explicit discourse moves that the participants make in the conversation, and by various pragmatic factors. Discourse moves can convey cooperation or adversariality in a more or less overt manner: enthusiastic assent, explicit rejections, etc. If an audience rejects every move that the speaker attempts, the speaker will identify an adversarial situation. But speakers and audience also use pragmatic cues to assess cooperation, and thus attribute default attitudes. Here are some of these pragmatic cues.

Priors. A speaker engages in conversation with prior beliefs about the other participants and whether the conversation is cooperative. When they tell their teenage child to clean their room, they expect some resistance. When they discuss which movie to watch with their friend, they assume that their conversational goals align: find a movie that they will both enjoy.

These priors are defeasible. At any point, the speaker may reassess the situation, and discover that a setting is more, or less cooperative than they previously thought. They can also investigate whether their audience actually holds the default attitude they attribute them.

- (134) After asking their child to clean their room, and meeting only silence, the speaker checks in. ‘Did you hear me? Does that mean you will clean up?’

These priors are influenced by what the speaker knows of the audience’s beliefs and attitude, but also by the sociological factors at stake, such as power dynamics, epistemic authority, politeness, etc. Among other things, a speaker may deduce from their audience being part of a similar community that they may share conversational goals. Persons of colour and LGBTQ+ folks, among others, often report the benefits of having spaces specific to them. These spaces enable them to approach topics of discrimination with like-minded individuals who are more likely to engage cooperatively in such conversations.

Power dynamics and adversariality. In addition to belonging to the same community, other sociological factors influence the perceived cooperativeness of a situation. When a speaker is an authority in a conversation (be it because they have more sociopolitical power, or more knowledge), they are likely to take their audience’s silence as assent. Authority enables a speaker to make closed calls and expect positive replies. When a boss requests their employee to fulfil a contractual obligation, they expect assent to their request, as they have the authority to formulate it.

The case of open calls is more complex, but epistemic authority also plays a role. In (135), the speaker is the epistemic authority in the conversation. She can be reasonably sure that her open calls – her assertions that constitute the lesson, her questions to her students – will be heeded. The audience will update their copy of the Common Ground with the information her utterances convey. Conversely, if the audience is an epistemic authority as in (136), then the speaker can interpret silence as unconvinced, or hostile.

- (135) A teacher is giving a history lesson in front of her class. She is confident that she knows more on the lesson topic than her students.
- (136) A student has claimed that dolphins are mammals. Her professor asks her to justify her claim. The student goes through her evidence.

In (135), as long as there are no objections to clear up, the teacher can safely assume that her students’ silence stands for an assent to update the Common Ground with the contents of her lesson. In a context such as (136), the student assumes that, when her teacher keeps silent, she needs to provide more evidence. The update is not going through unless her audience explicitly allows it.

Power dynamics that do not correspond to epistemic authority also influence the perceived cooperativeness. When a participant has political power over other participants, the latter may not feel comfortable discussing their specific experience of oppression. Because of the power dynamics, the oppressed participants may remain silent on some topics, as they do not feel safe voicing their dissent. This is part of the phenomenon of self-silencing that Tanesini (2016) analyses:

Whenever illocutionary silencing is deployed to prevent dissent, the silenced individuals will soon learn that it is less risky to share the views of those who are capable of silencing them. (Tanesini, 2016, p. 90)

Goldberg (2020a) takes exchanges that are skewed by heavy power dynamics to fall under a non-conversation category:

NON-CONVERSATION: The particular speech exchange is not a conversation – it is not a cooperative exchange – in the first place. (Goldberg, 2020a, p. 175)

I am, with Lackey (2018), more reluctant to exclude such exchanges from the category of conversation (see section 6.6). But for now, it suffices to say that, depending on conversational goals, power imbalances affect how the speaker perceives the cooperativeness of a conversation. They might be confident that their audience will assent to their claims (because this audience has no other choice). Or they might expect their audience to oppose their challenges to the political *statu quo* (as in the cases of white silence).

Para-linguistic cues. A third type of contextual factors are para-linguistic cues. It is very rare that an audience is entirely ‘silent’, in the sense delineated by Tanesini. Instead, audiences nod, look sceptical, smile, frown, hum with interest, etc. These para-linguistic cues come into play for the speaker to distinguish between an audience that is silent because they assent to the update proposals, and an audience keeping silence out of disinterest, or boredom.

- (137) A teacher is giving a history lesson in front of her class. Upon seeing perplexed frowns, she realises that she needs to provide more details on the changes in U.S. foreign policy following the Cuba missile crisis.
- (138) I advise my friend to read *The Last Girl Scout*. He seems unsure, so I make a case for the book, until he hums with interest.

The para-linguistic cues allow the speaker to interpret in a more fine-grained manner the attitude their audience holds towards their discourse. Para-linguistic cues can change their perception of the cooperativeness degree of a conversation.

6.4.3 From cooperativeness to interpretation of silence

Following these different cues, and others, a speaker determines whether a dialogue is cooperative or adversarial. With this assessment, they attribute a default attitude to their audience. If the conversational situation is cooperative, they assume that their audience shares their conversational goals, and thus that the less

marked attitude is assent. In cooperative settings, silence is interpreted as an assent. If the conversational situation is adversarial, they assume that their discourse move is not accepted until an explicit confirmation is reached. In adversarial settings, silence is interpreted as dissent. If the situation is semi-cooperative, the speaker may be unsure whether their discourse move is accepted or rejected, unless there is an explicit confirmation or rejection from the audience. The semi-cooperative setting makes it likely that a silence is an invitation to ground their assertion with evidence, explain their question, show that their request makes sense. The speaker attributes these default attitudes to an audience until proven otherwise; either by an explicit discourse move (confirmation, rejection), or by pragmatic cues that indicate changes in the cooperativeness.

6.5 Silence and updates

The perceived cooperativeness of a conversation provides three possible interpretations of silence:

1. Fully cooperative conversation: silence means assent
2. Semi-cooperative conversation: silence means ‘tell me more!’
3. Adversarial conversation: silence means dissent

This allows to define distinct behaviours of conversational updates depending on conversational cooperativeness. Speech acts can be taken, for simplicity sake, to be proposals to update a (more or less complex) conversational Common Ground in certain ways (section 2.2.2). These proposals may be accepted or rejected by the participants in a conversation. Or the audience may suspend their judgment, leaving the Common Ground undecided on the speech act, and keeping both worlds compatible and incompatible with the update in the context set.

These different interpretations of silence require us to separate two essential effects of speech acts. First, the effect of adding to the conversational record that a proposal to update the Common Ground in a certain way was made. Second, the effect of updating the Common Ground *only if* the proposal is accepted. The requirements for acceptance can be determined depending on the cooperativeness of the dialogue. In some cases, acceptance only requires an absence of rejection. In others, acceptance needs to be made explicit, or even enthusiastic.

Consider two examples, one of open call – an assertion in (139) – and one of closed call – a request in (140)³.

³While the overt speech act in (140) is a question, I follow Asher and Lascarides’ (2001) account of indirect speech acts in interpreting its conversational effect as a request. After all, the speaker who would treat it only as a question would appear either incompetent or uncooperative (Grice, 1961, 1975).

(139) Paul is not coming to the party.

(140) Might you pick some tomato juice at the store?

The first essential effect of (139,140), no matter the cooperativeness of the conversation, is that they make public a proposal of the speaker. In (139), Speaker proposes adding to the information they share with their audience that Paul is not coming to the party (Stalnaker, 1978). The proposal is an *open call* that extends to anyone in the audience, and where anyone may pick up the conversation with a relevant contribution. In (140), Speaker proposes adding to the addressee to-do list to pick up tomato juice (Portner, 2004). The proposal is a *closed call* that only someone with the authority to do so (most often the addressee) may assent to or dissent from.

In cooperative contexts, it is enough that the participants voice no dissent for the proposed update to go through. That the conversation is cooperative, and the participants voice no dissent, means that the utterance fulfils its second illocutionary effect: updating the conversational Common Ground with the content of the update.

By contrast, semi-cooperative contexts require an expression of assent or dissent to determine the status of the update. The example of closed calls shows just how eliciting an expression of assent or dissent can be considered an illocutionary effect of a speech act in its own right. Closed calls stand in contrast to open calls as they more often require explicit assent or dissent, even in apparently cooperative conversations. Suppose Speaker utters (140) while making preparations for a party with Addressee, and Addressee committed to go pick snacks at the store. Except in cases of subordination, Speaker is likely to seek confirmation from Addressee that the task was added to their to-do list. A silence from Addressee would be, at the very least, uncomfortable.

Semi-cooperative contexts are the norm rather than the exception when Speaker seeks to elicit an active response from the addressee – as opposed to simple uptake. In the party planning context, Speaker and Addressee share a common goal: organising a successful party. The request from Speaker to Addressee (140) highlights that Speaker and Addressee's goals also differ in some respect. Speaker wishes to add *picking up tomato juice* to Addressee's tasks, while, presumably, Addressee seeks to keep their to-do list as short as possible. Crucially, if Addressee is not subordinated to Speaker, they are the sole authority on the update of their own to-do list.

So the update of Addressee's to-do list is a perlocutionary effect of the request (in much the same way as Addressee coming to believe the content of an assertion is a perlocutionary effect of the assertion). On the other hand, that something was required of Addressee and they are expected to answer this request is added to the Common Ground when Speaker makes the request. In semi-cooperative contexts, the second illocutionary effect of an utterance is to require uptake by the participants.

Finally, in adversarial contexts, Speaker assumes that their assertion, or request, is rejected unless the Audience explicitly assents. The first illocutionary effect of her utterance, to propose an update, still holds. But unless they receive assent from the audience, no other effect goes through on the conversational Common Ground.

I sum up essential illocutionary effects of utterances depending on contexts:

1. Fully cooperative conversation:

- (a) The utterance proposes an update to the Common Ground.
- (b) the update goes through *unless rejected by the participants*.

2. Semi-cooperative conversation:

- (a) The utterance proposes an update to the Common Ground.
- (b) The utterance introduces an expectation that the Audience expresses assent or dissent.

3. Adversarial conversation:

- (a) The utterance proposes an update to the Common Ground.
- (b) The update is rejected *unless explicitly accepted by the participants*.

6.6 Silencing without assent

Please note that the topic of illocutionary silencing is enmeshed with the notion of power imbalance in conversations, and the abuses it may foster. While I try to discuss it in an abstract manner, this section mentions breach of consent, sexual coercion and dynamics of abuse. If these are particularly difficult topics for the reader, I urge them to skip this section.

If silence does not project assent, then another issue arises: how to understand the problematic aspects of silencing. According to Goldberg (2020b), a picture in which silence gives assent explains the harm of silencing. When a participant remains silent, they are taken to assent, in a form of epistemic injustice. But in a picture where silence might well mean dissent, or uncertainty, I cannot help myself to the same explanation. Let us see how to make sense of the harms of silencing.

Illocutionary silencing is defined as a conversational injustice by which a participant's utterances are dismissed (Langton, 1993). It is a form of *uptake failure* that prevents a participant's speech act from taking effect (Caponetto, 2017). In this sense, silencing figuratively reduces a participant to silence: their utterances are treated *as if* they had remained silent. And, in a climate where they

are silenced, a participant might become *literally* silent, as they learn that their utterances are not taken seriously (Tanesini, 2016).

A classical explanation of the harms of silencing is that it simulates a participant's assent (Goldberg, 2020b). Feminist philosophy takes the example of how pornography, defined as a sexual representation of the subordination, dehumanisation and degradation of women (Langton, 1993) affects women's ability to refuse sex (Bird, 2002; Caponetto, 2017; McGowan et al., 2011; Mikkola, 2011). More broadly, silencing leads to disregard participants' dissent in conversational contexts. But if silence does not give assent, as I claim, how to explain silencing?

I keep the analysis of silencing as figuratively reducing a participant to silence. When a participant is being silenced, their utterances do not affect the conversation, as if the participant had not said anything. But silencing does not simulate assent, as, in my analysis of silence, a participant being silent may indicate other attitudes. Instead, silencing a participant casts them in a passive role, where they cannot decide where the conversation goes. In that sense, they are not an enthusiastic participant in a conversation – which maintains the analogy with sexual consent.

Since conversational contexts might be more or less adversarial, the absence of effect of the silenced participant on the conversation might not be interpreted as assent. And indeed: in this fiction of dialogue, the speaker decides of the cooperativeness of the setting – and the silenced participant cannot change this assessment by use of cues (section 6.4). The silenced participant is attributed a default attitude depending on which level of adversariality the speaker assumes. As the participant is silenced, they are not able to stray from this default attitude that the speaker assigned. This leaves the possibility that a speaker might silence a participant *while attributing them an attitude of dissent*. Example (141) covers a case where the speaker disregards signs that their audience has already accepted the update of the Common Ground they proposed. Example (142) depicts a situation where the speaker casts their Audience in a role of dissent. These situations are sadly common in contexts of emotional abuse.

- (141) Speaker defends their opinion. While their audience has, several times, indicated that they did agree and were convinced, and wished to change topics, Speaker continues to pile up their arguments.
- (142) An abuser thinks that their victim opposes them. While the victim has apologised, several times, and indicated their agreement with them, the speaker takes these behaviours as signs of defiance and insubordination.

In both examples, the analogy between illocutionary silencing and silence holds. It is as if the Audience did not say anything. However, the result is not that they are attributed assent. In fact, they are maintained in a default attitude the speaker attributed them. The harm of silencing remains, but it does not stem

from a fake assent. When they are silenced, the participant is not granted *agency* on the conversation – both their dissent, but also assent can be ignored. The speaker takes all control of the conversation, and casts the participant in a role (which may well be dissent) from which they can't deviate.

6.7 Conclusion

In this chapter, I take seriously the idea that silence does not always imply assent. To do so, I analyse silence as the expression of a default attitude. This default attitude is attributed to conversational participants on the basis of the cooperativeness level of the conversation. I delineate some ways in which participants determine the cooperativeness of a conversation.

Taking silence to express a default attitude, not necessarily assent, goes against the idealised Stalnakerian picture. It requires a more fine-grained analysis of the behaviour of speech acts in more or less cooperative contexts. But it also shows how we can model, and even make use of non-cooperative contexts when analysing the effects of speech acts on a less idealised picture of conversation. In this sense, taking silence seriously emphasises the talk of proposals when considering the effect of speech acts on Common Ground. While the full update effect of a speech act can only occur if the participants assent to the update, the first and main effect of most speech acts is to issue a proposal to update. No matter the reaction of the participants, the proposal goes on the conversational record. And considering that silence might not mean immediate assent also highlights the part that pragmatic cues play into determining uptake from the participants in a conversation.

A picture where silence does not mean assent also complicates the analysis of silencing. How to explain the harms of reducing a participant to silence, if this does not force their assent? I suggest thinking of silencing as a deprivation of agency, whereby the participant is not able to stray from a default attitude that the speaker assigned to them. This picture explains cases where the silenced participant is cast in an role of opposition, even while trying to express their agreement. My analysis of silence gives full space to a notion of enthusiastic participation in conversation, where a participant is able to express the full range of their agreement or disagreement. When silencing occurs, the participant is not always forced to assent, but they are certainly prevented from impacting the course of a conversation.

Chapter 7

Communicating with colourings[†]

7.1 Introduction

(143) My damn bike has a flat!

(144) One of my bicycle tires has emptied of air.

Utterances in (143) and (144) express similar thoughts, and are true under the same conditions. They differ in the choice of words that carry the meaning. Frege calls these differences in tone or *colouring* (Frege, 1879, 1892). *Colourings* are expressions, or grammatical constructions, that do not modify the truth of a sentence when substituted to an alternative expression, but may convey an additional content. The use of the formal address in languages that have one conveys that the relation between speaker and addressee is somewhat distant. The use of *but* instead of *and* to form a conjunction conveys that the conjuncts stand in some sort of contrast (chapter 3). The use of *damn* in (143) conveys the speaker's annoyance at her bike.

Colourings resemble Grice's conventional implicatures: inferences that are entailed by the conventional meaning of lexical items, but are compositionally independent of the explicit contribution of a sentence (Grice, 1975; Potts, 2005). However, Sander (2019) argues that Gricean implicatures do not subsume the category of colourings. According to him, some colourings are non-communicative.

My goal in this chapter is to challenge the notion of communication at stake in this argument. When Sander says that a speaker can *convey* some content by the use of a colouring without *communicating* this content, because she does not *intend* to communicate such content to her audience, he is relying on a picture of communication where speaker's intentions are the landmark of communication. However, I argue for an intersubjective picture of communication, where *what the audience infers* that the speaker intends to communicate is also relevant to what

[†]This chapter is based on published work (Bussière-Caraes, 2022).

is communicated. In such a picture, the use of colourings impacts what a speaker communicates when she makes an utterance.

After exposing the concepts of colourings and conventional implicatures at stake, I show that the claim that some colourings do not communicate anything requires a narrow notion of communication, grounded on speaker's intentions (section 7.2). If a speaker only communicates what she intends to communicate, then the content of some colourings is only conveyed. This allows to distinguish colourings from conventional implicatures, that communicate their content. However, a more intersubjective notion of communicative intentions is possible, and compatible with Grice's theory of conversation (section 7.3). Under this notion, a speaker who uses certain expressions is attributed a communicative intention by her audience based on what speakers generally intend when using these expressions. Since so-called non-communicative colourings trigger specific inferences, a speaker who uses them communicates these inferences. Therefore, I vindicate the communicative role of all colourings with content (section 7.4).

7.2 Colourings and Conventional Implicatures

According to Frege, if two sentences have the same *Sinn* and (arguably) the same *Bedeutung*, they express the same thought and are true under exactly the same conditions. However, they can differ in what Frege calls *tone* or *colouring* (Frege, 1879, 1892). Differences in colouring may affect connotations of a sentence, but do not touch its truth-conditions.

7.2.1 Colourings

The characterisation that Frege gives of colourings is mostly negative. Typically, colourings are triggered when a lexical item or expression contributes the same thought as an alternative expression but may differ in the connotations (hints) it carries. Frege gives a variety of examples: colourings include the difference between active and passive voice, formal and informal address (in languages that have such a distinction), evaluatives (*fortunately*, *damn*), and even the difference between *but* and *and*.

The way that *but* differs from *and* in that we use it to intimate [an-deuten] that what follows it contrasts with what was to be expected from what preceded it. Such conversational suggestions make no difference to the thought. A sentence can be transformed by changing the verb from active to passive and at the same time making the accusative into the subject. In the same way we may change the dative into the nominative and at the same time replace *give* by *receive*. Naturally such transformations are not indifferent in every respect

but they do not touch the thought, they do not touch what is true or false ... (Frege, 1918a, p. 331)

The sentences in (145) are a pair of sentences that do not differ in thought, but in colouring.

- (145) a. The dog howled all night.
b. The cur howled all night.

When uttering (145b), the speaker expresses the same thought as when uttering (145a). But by using the term *cur* instead of the term *dog*, she conveys her dislike of the dog she is referring to. In (145), colourings affect connotations of a sentence, but not its truth-conditions.

Following the Fregean characterisation, colourings occur when different lexical items or expressions are truth-conditionally equivalent alternatives. Their connotations and use may differ, but they could be substituted to each other without changing the truth (or falsity) of a sentence. The use of active or passive voice in (146) does not affect the thought expressed by the sentence. (146a) and (146b) are true under the same conditions.

- (146) a. Caesar crossed the Rubicon.
b. The Rubicon was crossed by Caesar.

Sander (2019) uses Frege's list of examples to delineate four features of colourings. He calls the truth-conditionally equivalent expressions that trigger colourings *c-devices*. He provides a characterisation of *c-devices* in four points, that I gloss with example (145).

- i. Colourings are assertorically inert.

All *c-devices* are assertorically inert: If Σ is an assertoric sentence that contains a *c-device* which hints that q , then uttering Σ does not amount to asserting q .

Asserting *the cur howled all night* does not amount to asserting that one disapproves of the dog.

- ii. Colourings are alethically inert.

All *c-devices* are alethically inert: If Σ is an assertoric sentence with the propositional content p that contains a *c-device* which hints that q , then the truth-value of q does not affect the truth-value of p or the truth-value of Σ as a whole.

The disapproval of the dog has no bearing on the truth of the assertion *the cur howled all night*.

iii. Colourings are inferentially inert.

All c-devices are inferentially inert: If Σ_1 and Σ_2 are assertoric sentences with the propositional content p that differ only in colouring, then Σ_1 and Σ_2 are logically equivalent.

The dog howled all night and *the cur howled all night* are logically equivalent assertions. For this characterisation of logical equivalence, Sander relies on a truth-conditional, Fregean picture.

iv. Some colourings are communicatively inert: they do not communicate anything.

Some c-devices are communicatively inert: Not every c-device contains a hint (*Andeutung*) that q as part of its meaning.

Given these features of colourings, Sander denies that Frege intended colourings only as aesthetic phenomena, against Dummett (1981, p. 84sq). On the contrary, colourings may have content. They are adequate to certain situations. They also, sometimes, convey additional meanings by triggering inferences to their content. A prime example of a colouring with content is the difference between *and* and *but*.

- (147) a. Linda is tough and fair.
b. Linda is tough but fair.

(147a) and (147b) are true under exactly the same conditions: Linda is tough, Linda is fair. However, the use of *but* in (147b) conveys an additional content: that Linda being tough, and Linda being fair, stand in contrast.¹ But not all colourings hint at established, conventionalised contents. While the difference between *and* and *but* is shared among speakers of English – the content is grounded intersubjectively – some colourings have contents that are less conventionalised – the subjective difference in how an audience interprets *tempest* against *storm*. The less conventionalised contents are at play in the poetic use of colourings (Frege, 1983, p. 151sq; translated in Beaney, 1997, p. 240; Dummett, 1981, p. 85). But, be they conventionalised or not: some colourings have contents they communicate by means of hinting [*Andeutung*]. Horn (among others) likens this relation to Gricean conventional implicatures (Grice, 1975).

... the *Andeutung* relation, for a component of linguistic meaning that does not affect propositional content or touch what is true or false, is a direct precursor of Grice's conventional implicature. (Horn, 2013, p. 153)

¹I study the question of the contrast that the adversative marker *but* in chapter 3.

7.2.2 Conventional Implicatures

Conventional implicatures were born in neglect: in their first mention, Grice contrasts them with conversational implicatures.

In some cases the conventional meaning of the words used will determine what is implicated, besides helping to determine what is said. If I say (smugly), *He is an Englishman; he is, therefore, brave*, I have certainly committed myself, by virtue of the meaning of my words, to its being the case that his being brave is a consequence (follows from) his being an Englishman. But while I have said that he is an Englishman, and said that he is brave, I do not want to say that I have *said* (in the favored sense) that it follows from his being an Englishman that he is brave, though I have certainly indicated, and so implicated, that this is so. I do not want to say that my utterance of this sentence would be, strictly speaking, false should the consequence in question fail to hold. So some implicatures are conventional ... (Grice, 1975, p. 44sq)

As implicatures, conventional implicatures trigger defeasible², pragmatic, inferences. However, by contrast with conversational implicatures, conventional implicatures follow from an utterance in virtue of the words the speaker uses. That leads to four key differences between conventional and conversational implicatures.

First, conventional implicatures are not calculated from breaking a Gricean maxim of conversation. Instead, they are associated with the conventional meaning of a word.

Secondly, conventional implicatures are detachable.

- (148) a. He is an Englishman; he is, therefore, brave.
 b. He is an Englishman; and he is brave

In (148), sentences (148a) and (148b) are truth-conditionally equivalent: both are true if and only if the subject referred to by *he* is both an Englishman and brave. But in (148), an inference that the subject being brave follows from his being an Englishman follows from the sentence *He is an Englishman; he is, therefore, brave*. The same inference does not follow from the truth-conditionally equivalent sentence *He is an Englishman and he is brave*. Conventional implicatures depend

²Defeasibility of conventional implicatures is distinct from their cancellability. Conventional implicatures are hard to cancel: a speaker who follows up her claim that *Linda is tough, but fair* with *Not that toughness and fairness are usually in contrast* appears, at best, clumsy. However, the inference that follow from conventional implicatures is by nature pragmatic. As such, it can be revised without contradiction, conversely to entailments. A speaker can follow up *My damn bike has a flat again* with *I'll get it repaired, I love the damn thing* and undo the inference that she hates her bike.

on the lexical items that are used to express a proposition. By contrast, conversational implicatures are not detachable: they do not depend on the lexical items that express a proposition, but are triggered by the proposition being expressed in the first place. In the conversation in (149) (Grice, 1975), the conversational implicature that B expects C to be dishonest at his new job is triggered by the assertion of either (149b) or (149c).

- (149) a. So, how is C faring at his new job?
 b. Oh, very well. He gets on well with his colleagues and hasn't been arrested yet.
 c. Oh, very well. He gets on well with his colleagues and is still walking free.

Conversational implicatures are not triggered by specific lexical items, but by the utterance of a sentence in a certain context. The utterance of sentences with identical truth-conditions trigger the same conversational implicatures.

Thirdly, conventional implicatures are harder to target and cancel than conversational implicatures. While the speaker in (150) easily cancels the implicature that not all students passed the class, the speaker in (151) is, at best, clumsy, but most likely contradictory.

- (150) a. Some of the students passed the class.
 b. In fact, all of them did!
- (151) a. He is an Englishman; he is, therefore, brave.
 b. # But I don't mean to imply that Englishmen are brave!

Finally: conversational implicatures are truth-conditionally independent from the sentence they are associated with. If the implicature that not all students passed triggered by (150a) is false, the utterance remains true – as the successful cancellation shows. On the other hand, whether the truth of conventional implicatures is indifferent to the truth of their trigger utterance is debatable (Bach, 1999).

Unfortunately, except from setting them apart from conversational implicatures, Grice does not provide a principled characterisation of conventional implicatures. He gives examples: *but, therefore, even...* And the Gricean examples of conventional implicatures seem right at home in the Fregean category of colourings. In Fregean terms, conventional implicatures are differences in lexical items, or expressions, that do not affect the thought of a sentence, but nonetheless *hint* at an additional connotation. For example, that being brave in some way follows from being an Englishman.

Potts (2005) narrows down the category of Gricean conventional implicatures to implicatures that bring about commitments (p. 11); and he adds to the Gricean

examples supplemental expressions and expressives. According to Potts, these conventional implicatures convey additional commitments. By using the expressive *damn*, the speaker of (152) commits to a certain dislike of her bike at the moment of utterance.

(152) My damn bike has a flat again.

7.2.3 Distinguishing colourings and conventional implicatures

Sander (2019) distinguishes Fregean colourings and conventional implicatures. He classifies colourings in:

- i. Purely aesthetic phenomena
- ii. Colourings with content
 - a. communicative colourings or ‘hints’: *but, cur, unfortunately, etc.*
 - b. non-communicative colourings: *tu* vs. *vous, it is true that*, double negation, etc.

Communicative colourings, that bring about extra commitments, correspond to Potts’ conventional implicatures. On the other hand, according to Sander, non-communicative colourings with content and aesthetic colourings do not bring about commitments. The use of *it is true* to preface a statement allegedly does not bring about extra commitments from the speaker (or, if it brings them, they are not easy to subsume under one proposition). Thus, Sander sets colourings apart from conventional implicatures.

I argue against Sander’s distinction of communicative and non-communicative colourings. This distinction relies upon a picture of communication according to which a speaker communicates a content only if she means to commit to it. But when we take into account the fact that speaker intentions are reconstructed by an audience, we arrive at another notion of communication: where a speaker can communicate a content without necessarily *meaning to*, or committing to it in a strong sense.

7.3 Notions of communication

Non-communicative colourings, according to Sander, do not communicate their content. But this depends on his theory of communication, that is grounded on speakers intentions. He opposes Horn’s account of the linguistic meaning of *tu* in the sentence *Tu es souûl* (*you* [informal] *are drunk*):

Thus in affirming *Tu es soûl*, my belief that a certain social relationship obtains between us and that you are male is not part of the thought or of what is said; both propositions are indeed communicated, but what is said is simply that you're drunk. (Horn, 2013, p. 159)

... typically, *S* utters a sentence such as *Tu es soûl* (as opposed to *Vous êtes soûl*) because, among other things, there are certain social rules for using the word *tu* and because *S* intends to abide by these rules, but *S* does not thereby communicate to *H* (or to somebody else) that a certain social relationship obtains between *S* and *H*. Similarly, it would be odd to say that by using the masculine form *soûl* (vs. *soûle*) *S* is communicating to *H* that *H* is male or that *S* believes *H* to be male. (Sander, 2019, p. 386)

The difference between Sander and Horn stems from different understandings of the notion of communication. For the utterance *Tu es soûl* to be felicitous, certain conditions need to obtain. The addressee needs to be a man, and a certain social relationship needs to obtain between the speaker and the addressee. Now, according to Horn and Sander, these conditions are not part of the thought, or of what is actually *said* in the utterance. Indeed, the English translation '*You are drunk*' that preserves the thought of the utterance does not preserve the extra conditions of the French utterance. Where Horn and Sander disagree is on whether these conditions, which are not part of the utterance, are communicated.

Sander takes a hard line on Gricean communicative intentions: according to him, a speaker only communicates what she intends to communicate. He uses other terms for meanings that are associated to non-communicative colourings with content. When using *tu* instead of *vous*, the speaker 'displays her knowledge' of social rules and command of French, conveys the existence of a certain relation, etc. The term of communication is reserved to a content that the speaker meant to share with her audience – and that the audience appropriately recognises.

From a Gricean point of view, the term 'communication' should be roughly equivalent to 'non-natural meaning', and it seems obvious to me that in Horn's example the core condition for non-natural meaning is not satisfied: *S* can only mean or communicate something by 'tu' if *S* intends her use of 'tu' 'to produce some effect in an audience'. (Sander, 2019, p. 386)

Based on this picture of communication, Sander distinguishes between two categories of colourings with content. Some, according to him, hint at their content in a Fregean sense: these are communicative colourings with content, such as *but* vs. *and*, *therefore*, *cur*. He likens these colourings to conventional implicatures. On

the other hand, non-communicative colourings with content are, to him, Frege-specific, as they do not properly communicate their content, but only convey it. In this category, Sander classifies formal vs. informal address, double negation, *it is true that...* Because a speaker, when using formal voice, does not mean to communicate that a specific relationship obtains with her interlocutor, this content is not communicated, but only conveyed. On the other hand, presumably, when a speaker uses *but*, she means to communicate the existence of a contrast between the conjuncts of her utterance.³ Sander notes:

Since a certain type of social relationship is just what makes a sentence containing one of these two words appropriate or felicitous, correctly using such a sentence conveys the speaker's belief that a certain social relationship obtains between *S* and *H* (on the difference between conveying and communicating, see Stanley 2002: 327). Alternatively, we might say that, by using 'tu' or 'vous', a speaker displays her knowledge of the rules for speaking French felicitously. (Sander, 2019, p. 387)

However, a more intersubjective take on communicative intentions is possible and compatible with a Gricean perspective. This intersubjective picture of communicative intentions blurs the line between communicative and non-communicative colourings.

To establish this picture of communication, the first thing to do is establish that speaker's intentions are not transparent. An audience is not omniscient, and thus not automatically aware of what a speaker means by their utterance. Moreover, in most cases, speaker's intentions are not explicitly communicated to the audience. Instead, the audience recognises and guesses speaker's intentions on the basis of general use.

Explicitly formulated linguistic (or quasi-linguistic) intentions are no doubt comparatively rare. In their absence we would seem to rely on very much the same kinds of criteria as we do in the case of non-linguistic intentions where there is a general usage. An utterer is held to intend to convey what is normally conveyed (or normally intended to be conveyed), and we require a good reason for accepting that a particular use diverges from the general usage (e.g. he never knew or had forgotten the general usage). Similarly in non-linguistic cases: we are presumed to intend the normal consequences of our actions. (Grice, 1957, p. 387)

³A challenge to Sander's picture of Gricean communicative intentions is the question of the level of intent required for an intention to be communicated. In the case of *tu* vs *vous*, I certainly intend to convey that the use of *tu* is appropriate in a certain dialogue. It is just not my primary intention. However, it is also hard to say that when using *but*, the speaker has a primary goal of communicating the existence of a contrast, or even is aware of communicating the contrast.

The gist of it is: while speakers intentions do matter for communication, *intention recognition* is also crucial. In his later works, Grice focuses on speaker intentions, but I believe it is under the premise that such speaker intentions need to be recognised by the audience.

Bach and Harnish (1979), in a very traditional picture of Gricean intentions, count as crucial intention recognition for communication to actually take place.⁴ Their Speech Act Schema is the pattern according to which an audience infers speakers intentions from what they utter, under the communicative presumption.

The communicative presumption is the mutual belief prevailing in a linguistic community to the effect that whenever someone says something to somebody, he intends to be performing some identifiable illocutionary act. (Bach and Harnish, 1979, p. 12)

Speaker's intentions don't exist in a vacuum. They are recognised by an audience. So, while it is tempting to have speaker's intentions be law in how to interpret their utterances, there is always an element where the audience matter: the audience is the one who ascribes intentions to the speaker to illuminate their utterances.

Of course, the guesswork of intention recognition is principled: not every type of intention is up for grabs. Intention ascription follows common use of utterances and expressions. In the same way as an agent is expected to intend the normal consequences of their actions, a speaker is expected to intend the normal interpretation of their speech. If an agent lets go of a cup, they are assumed to intend the cup falling to the ground and breaking. If a speaker says *Please close the door*, they are assumed to intend that their addressee closes the door. To reconstruct the intentions of the speaker of an utterance, the audience needs to hold some reasonable assumptions: they assume that the speaker is competent, sincere, and respects the other Gricean maxims (Grice, 1975). In addition, the audience makes use of conventions, general use of expressions, etc. In this sense, the only speaker's intentions an audience comes to know are *the intentions that an audience rightfully ascribes to the speaker*. So what a speaker communicates is what her audience recognises she intends to communicate.

To further the argument: there are cases where the actual speaker intentions and the intentions that the audience ascribes to the speaker differ. But because the audience is justified in their intention ascription, we judge that the speaker

⁴Bach and Harnish also counter Searle's (1969) argument that Gricean effects of utterances are not produced by the recognition of an intention to produce them. For example, one might recognise that a speaker intends them to close the door, and yet refuse to do so. They solve this problem by distinguishing between illocutionary and perlocutionary effects (Austin, 1975).

the speaker's illocutionary act, whose identity he is trying to communicate, can succeed without the intended perlocutionary effect (if there is one) being produced. (Bach and Harnish, 1979, p. 14)

communicated what the audience interpreted, even though it differs from what the speaker meant. Here is an example:

- (153) I want to meet with Patricia for drinks. I call her: *Let's meet at the Descartes bar and go for drinks!* The reason why I suggest the Descartes is that, albeit closed, it is ideally situated for a meeting point – close to the metro, good parking spots, etc. Other nice bars are in the same street.

I make this proposition under the assumption that Patricia knows that the Descartes is closed. This is where I am wrong: she is not from the same neighbourhood, and doesn't know the Descartes is closed.

My intention to communicate a meeting point is recognised: Patricia correctly judged that I suggested the Descartes as a meeting point. However, she also ascribed me the intention to communicate that we would have drinks at the Descartes specifically. It was a perfectly reasonable intention to ascribe me given her knowledge state. Thus, I can be said to have, albeit unknowingly, communicated to Patricia that we would drink at the Descartes.

In a theory of communication fully focused on speaker's intentions, we would have to say that what I communicated in (153) is only that we would meet at the Descartes, and go somewhere else for drinks. That would make Patricia's interpretation of my speech (her assumption that we would go have drinks at the Descartes specifically) a mistaken assumption, and not a result of my speech. But considering that, in context, the utterance: *Let's meet at the Descartes and go for drinks!* is generally used to communicate that one wishes to have drinks at the Descartes specifically, Patricia's assumption that I intended her to understand that we would drink at the Descartes is fair game. In fact, I can be held to wish to communicate that we would have drinks at the Descartes. What exactly is communicated also depends on the intentions the audience ascribes to the speaker.

As seen earlier, not any interpretation is available to the audience. For example, if in (153) Patricia were to ascribe me the intention to convey that we were to go canoeing – under the assumption that *drinks* actually meant *canoeing* – she would not be ascribing intentions in good faith. And the speaker is also expected to play by the rules. For them to properly communicate their intentions, they should take into account what their audience can reasonably assume they wish to communicate. In (153), I, as a speaker, cannot expect Patricia to understand that by *drinks* I mean *canoeing*, unless we previously established such a convention. By taking me to be a competent speaker Patricia assumes that I am aware of the general use of certain terms and expressions, such as *go for drinks*. And, wish it or not, as a competent speaker, I communicate the general content associated with a term when I use it. This is why retracting my proposal to go for drinks,

and claiming I actually meant to suggest canoeing, would be in bad faith, and not be readily accepted.

That what the speaker communicates depends, in fact, on what the audience recognises the speaker intended to communicate considerably broadens the notion of communication. It includes multiple cases where the speaker didn't mean to communicate what she did. She made an off-colour joke, she claimed something that turned out to have repercussions she did not intend. In a narrow picture of communication, these cases count as *communication failure* where the audience misunderstood the speaker (and example 153 is interpreted as a communication failure). I propose, instead, that these are part of what the speaker communicated. To understand why, I turn to the commitment analysis of discourse (Brandom, 1983, 1998). According to Brandom, when a speaker makes a discourse move, she undertakes a certain amount of commitments and responsibilities with respect to the content of her discourse move. Suppose that her discourse move has unintended consequences: say, it triggered a conversational implicature that was offensive to her audience. While her audience is mistaken into ascribing the speaker the intention to offend, the audience is *justified* in doing so, following ordinary use of expressions. In a narrow picture of communication, the offence was not communicated, as the speaker did not intend it. In the hearer-oriented picture of communication I defend, the offence was communicated: the speaker can be held responsible for it, and the audience can demand a retraction on the basis of their justified ascription. The existence of retractions, by which a speaker can acknowledge that her utterance no longer corresponds to her intentions (Caponetto, 2020; MacFarlane, 2014; Marques, 2018) is crucial to understand speaker's responsibilities towards unintentional effects of her utterances.

When we integrate intention ascription and intention recognition in our picture of communication, what we get is an intersubjective structure that relies on constant feedback loops.

- i. A speaker makes an utterance, by which she means to communicate content *p*.
- ii. But for this to work, the audience has to attribute to the speaker the intention to communicate content *p* on the basis of her utterance.
- iii. So the competent speaker will use lexical items that are generally (in the community), used to communicate *p*. For example: *Beware, there are bees!* is generally used to communicate to an audience that there are bees, and they should be careful in the area.
- iv. When the speaker uses lexical items, the audience reconstructs their communicative intention on the basis of what the lexical items are generally used to communicate. For example, the audience interprets that the speaker uses

Beware to communicate a warning, as it is generally used to warn others from a danger.

- v. Maybe the speaker actually meant to warn their audience to look out for the pretty bees.
- vi. But crucially, v. only matters if the speaker manages to make it known that they will henceforth use *Beware* as a general attention grabbing item, instead of a warning.

A note on the last point: speakers can change the general use of words (there is a creative use of language, such as reclaiming slurs). They do so only to the extent that their intention to use an expression in a way that diverges from the already established general use is recognised by, and sometimes made explicit to, the audience.

- (154) A professor starts a botanic lesson on berries. She declares that she will henceforth use the botanical definition of *berry*: a fruit produced from the ovary of a single flower in which the outer layer of the ovary wall develops into an edible fleshy portion (pericarp). This use of the word *berry* will cover many fruits that are not known as berries in the common use: grapes, tomatoes, cucumbers, eggplants, bananas, and chilli peppers.

But a speaker who attempts to use an expression in a way that differs from the general use without establishing her intentions runs into risks of misunderstandings. The professor in (154) would (reasonably) cause some confusion by stating that strawberries are not actually berries if she didn't clarify her intentions in using the term *berry* first.

This also brings forward an interesting consequences of accounting for ascribed intentions in communicated content. In a speaker-oriented view of communication, what is communicated is what the speaker intended. Communicative failure occurs when the audience does not ascribe the right intentions to the speaker. In our broader view of communication, what is communicated is what the audience ends up recognising. Communicative failure still occurs, under similar circumstances, when what the audience recognises is not what the speaker intended. But, crucially, there is a shift in responsibility: it is not the audience who failed to recognise speaker's intentions. Instead, it is the role of the speaker to clarify herself so that she can communicate what she intended to communicate.

This picture is, of course, schematic. The interactional nature of discourse and conversation refines it. Yes, the speaker communicates the content her audience ascribes her the intention to communicate, in the sense that she is responsible for the effect of her utterances even if her intention was mistaken by the audience. But the existence of a conversational feedback loop also makes it a responsibility for the audience to question their grasp of speaker's intentions, and if needed modify it based on further utterances. The defeasibility of pragmatic inferences

(Asher and Lascarides, 2003) is crucial here: it allows the audience to refine and revise their picture of speaker's intentions along the way.

Gricean communicative intentions still determine the interpretation of an utterance. But the communicative intentions are not only what the speaker intends to communicate, but also *what the audience recognises* that the speaker intends. This leads to a broad understanding of communication: not only what the speaker says, or means, is communicated. Additional inferences, based on the speaker's utterance, are also part of what the speaker communicates.

7.4 Communicating with colourings

When we consider intention ascription in communication, what a speaker communicates is not limited to what she intends to communicate. Instead, what a speaker communicates are the communicative intentions that her audience rightfully ascribes her on the basis of her utterance and reasonable inferences.

How does that apply to the category of non-communicative colourings with content? Recall that the main argument Sander provides for the existence of this category is that there are colourings with content (e.g. *tu* vs *vous* in French) that a speaker can use without communicating their content. The reason, according to Sander, is that when using such colourings, the speaker does not usually intend to communicate their content. When a speaker uses *tu* instead of *vous* in a conversation, she does not mean to communicate that a certain relationship obtains with her interlocutor.

I argue that under a theory of communication that takes into account intention recognition, these colourings actually communicate their associated content. To streamline the argument: since intention recognition matters, the speaker has to take into account the general use of an expression. She knows that she will be held by her audience to intend what is generally intended by the use of this expression. In particular, if a colouring conveys or hints at a certain content, it is because it is generally used to communicate this content. So when a speaker uses this colouring, she knows that she will be ascribed (and reasonably so) the intention to communicate what this colouring is generally used to communicate.

To focus on an example: a competent speaker knows that *tu*, in its general use, conveys a certain sense of familiarity and equality, distinct from the more distant *vous*. She is aware that other competent French speakers know the distinct uses of *vous* and *tu*; and that other competent French speakers will ascribe her the intention to convey the familiarity and equality that generally comes with the use of *tu*. Therefore, she should only use *tu* in a situation where it is adequate for her to convey such meaning. But what allows us to say that the speaker actually communicates that a relation obtains, and not only displays her knowledge of the rules of French?

Recall that colourings, in their first Fregean approximation, are expressions,

or constructions, that may convey an additional content, but do not affect the truth of a sentence. So colourings, almost by definition – such as *tu* vs. *vous*, *and* vs. *but*, passive vs. active voice – constitute alternative ways of dressing up the same sentence.

7.4.1. DEFINITION. Two colourings c_1 and c_2 are alternatives when c_1 and c_2 are expressions that can be substituted to one another in an utterance without affecting the thought it expresses.

(155) formal and informal address are alternatives.

And what we know about alternatives is that we may use them to trigger conversational implicatures (Grice, 1975):

(156) a. TEACHER: *Anna, je te prierai de me vouvoyer.*

Anna, I ask you [informal] to address me as vous [you, formal].

b. ANNA: *Et je te prierai de me vouvoyer.*

And I ask you [informal, focus] to address me as vous [you, formal].

In (156), Anna uses colourings to produce a conversational implicature. By using a *tu* address after being told by her teacher that a *tu* address was inappropriate, she implicates her refusal to use a formal address until the teacher also does it. She also, more generally, communicates what the *tu* address generally conveys, which is a sense of familiarity and proximity; this is the main reason for the sassiness of the example, that after being told to use the formal address, Anna still uses the familiar *tu* while echoing the teacher. Thus, she emphasises the unequal treatment: that the teacher requires a formal address while addressing her informally.

For a competent speaker, a sentence with a certain colouring is an immediate alternative to another sentence expressing the same thought, with a different colouring. When a competent English speaker uses *but*, she is aware that she could have used another conjunction, such as *and*, and chose not to. Similarly, since *tu* and *vous* are alternatives, the competent French speaker who uses *tu* knows that she could have used the other form of address. The same goes for other examples of so-called ‘non-communicative’ colourings: double negation, *it is true that*, etc. When using a double negation such as (157), the speaker knows that there is a more concise alternative. So she is also aware that she will be taken to implicate something by her use of an unnecessarily convoluted formulation.

(157) Marta is not not a lawyer.

What she will be taken to implicate specifically depends on theories of negation at play in the conversation, and whether the participants take a double negation to actually equate with an affirmation (Dummett, 1981; Horn, 2001; Rumfitt,

2000). But it does seem that the speaker communicates something different from the content of *Marta is a lawyer*. In my opinion, she would communicate something along the lines of *Marta is in a law-adjacent field*, or *Marta is a very bad lawyer*.

A competent speaker is aware that colourings are alternatives to another formulation; and she is aware that when she uses one, she conveys that she did not choose the other. So it is very easy for a speaker and an audience to generate conversational implicatures by drawing attention to the use of a colouring. The general rule is:

7.4.2. DEFINITION. Suppose colourings c_1 and c_2 are alternatives. They may hint at different contents. When a competent speaker uses c_1 , she conveys the content associated with c_1 . She additionally conveys that she did not choose to use c_2 in the context.

The competent speaker is aware of the rules of use of c_1 and c_2 respectively. These rules can correspond to what Sander distinguishes as communicative contents, like being aware that the use of *but* instead of *and* conveys the presence of a contrast between the conjuncts of a sentence. But they can also correspond to the contents of non-communicative colourings, like being aware that the use of *tu* is adequate in a conversation struck in a casual context, between people that are not concerned by hierarchical relationships.

The competent speaker is aware that, when she utters a sentence with a colouring c_1 , she will be ascribed the intention to convey that the sentence with the alternative colouring c_2 was inappropriate, or did not convey her intentions accurately. The parsimonious explanation is that she wishes to communicate the content generally associated with the colouring c_1 she uses. So, in the intersubjective picture of communication I delineated in section 7.3, the speaker will communicate the content associated with the colouring she used.

The content associated with alternatives colourings might be more or less close. The choice of *but* instead of *and* is more significant than the choice of *however* instead of *but*. In the latter type of cases, where there is less distinction between the content associated with the alternatives, the hearer may infer indifference of the speaker between two formulations.⁵ I believe this type of cases, where there is no significant content difference between alternative colourings, falls under Sander's category of aesthetic colourings; that is, colourings that do not have a particular content. When colourings differ in content, this content ends up being part of what is communicated by the speaker, as it is part of what the hearer can legitimately recover from the speaker's utterance and hold the speaker responsible for.

Finally, I want to ease some possible worries concerning the intersubjective notion of communication: having a wider notion of communication, that also

⁵I thank a helpful referee for pointing out the issue of indifferent elections.

covers implicated content, or content that is hinted at by certain constructions, does not come with a strengthening of speakers commitments. Why? because we maintain a distinction between content that is merely communicated, and content that is communicated and asserted. Speakers have responsibilities for the contents they communicate. But the heavy justificatory responsibilities come with asserting content. When a content is supposed, or implicated, or conveyed, the speaker is still assumed to communicate that they subscribe to it. But they are not held responsible to the same degree than when they make assertions.⁶

When a speaker communicates a content by way of conversational implicature, they affect the conversation in a different way than when they assert the same content. In (158b), the speaker uses a conversational implicature to communicate that they believe that C is likely to commit a crime at his new job. In (158c), they assert the content.

- (158) a. So, how is C faring at his new job?
 b. Oh, very well. He gets on well with his colleagues and hasn't been arrested yet.
 c. Oh, very well. I mean, it is only a matter of time that he commits a crime, but you know.

The effect on the conversation is different between communicating the same content with a conversational implicature than with an assertion. When they assert (158c), the speaker warrants their audience to repeat their claim, and to attribute the claim to them. *Speaker told me that C will probably commit a crime at his new job.* When they *implicate* the same content, while they still communicate it to their audience, they are not responsible for it in the same way. In fact, conversational implicatures are often used to communicate content *off the record*.

Similarly, when a speaker communicates a content by way of presupposition, or colourings, they endorse different commitments that when they communicate a content by way of assertion. The inferences from the use of colourings to their content are pragmatic, and hence defeasible, conversely to the inference from an assertion to a speaker being committed to the content of their assertion. They are also harder to target for a justificatory challenge. Speakers get to wiggle out of commitments to the content of their colourings (159). Speakers don't get to wiggle out that easily out of commitments associated with assertions: they have to retract them, with the issues that retractions comprise (160) (chapter 5).

- (159) She is poor, but honest. I mean, no moral judgment associated with poverty, that's terrible late capitalism talk.

⁶MacFarlane (2011b) compares different theories of assertion and surveys the associated responsibilities that the speaker bears (Brandom, 1983, 1998; Lackey, 2007; Stalnaker, 1978; Weiner, 2005; Williamson, 1996).

- (160) a. # Poor people are generally dishonest. But I don't want to associate a moral judgment with poverty.
- b. Poor people are generally dishonest. Uh, scratch that, that's a terrible thing to say.

The fact that colourings communicate their content does mean that speakers endorse commitments associated with the use of colourings, in the same way as they endorse commitments in a theory of conventional implicatures (Potts, 2005). But these commitments are not of the same nature as the commitments speakers endorse when they make assertions.

7.5 Conclusion

The argument to distinguish Fregean colourings from Gricean implicatures is that some colourings do not communicate the content that is associated with them. This argument relies on a notion of communication according to which a speaker communicates a content only if she intends to communicate such content. We can deem this position a hard line, speaker focused, take on Gricean communicative intentions.

However, Gricean communicative intentions are not transparent. Instead, they are known insofar as an audience recognises them in a speaker. Recognising intentions involves ascribing to a speaker communicative intentions based on what they say, under some reasonable assumptions. For example, that a speaker uses expressions according to a general use; or that she indicates if she intends to stray from the general use. When we take into account intention recognition in the picture of communication, we have to assume that a competent speaker means to convey the content associated to a certain colouring when she uses it. This is because colourings are alternatives to other possible ways of expressing the same thought in a sentence. So the competent speaker is ascribed the intention to communicate the content associated with a certain colouring.

In a theory of communication that takes intention recognition into account, the argument that some colourings do not communicate their content does not hold. So the distinction, on these grounds, between colourings with content and conventional implicatures collapses. If we take seriously intention recognition, competent speakers are supposed to intend the inferences normally associated with a colouring. That allows to treat colourings with content as communicative devices. Nonetheless, the communicative responsibilities associated with the communicative use of colourings are different from the responsibilities associated with proper assertions. This maintains a distinction between content that is asserted, and content that is merely communicated.

The running theme of this thesis has been speech acts that occur in situations of disagreement. In conversation, speakers use *rejections* and *weak assertions* as a way to preemptively block information they disagree with from entering the conversation. They use *retractions* to undo the effect of utterances they no longer agree with. Characterising these speech acts brought to light issues with idealised models of communication. When we assume that speakers contributions are accepted unless they are explicitly rejected, this assumption prevents us for accounting fully for disagreements in discourse settings. Therefore, I introduced directions to build new models of communication that take into account disagreement and audience input.

8.1 What we learned

The goal of this dissertation was to explore speakers' disagreements from different perspectives.

In chapter 2, I focus on the speech act of rejection. When speakers say *No*, to a propositional question or to another speech act, they express their disagreement with it. By doing so, they prevent the effect of the speech act from being added to the pool of information agreed upon by the participants in a conversation. Rejection has the effect of blocking contributions to a conversation. However, studying rejection as a primitive, self-standing speech act also supposes finding out its felicity conditions and constitutional norms. I start from the idea that the rejection of a statement is incompatible with its assertion. Thus, I build possible norms for rejection that are dual to norms that have been proposed for assertion. The candidate norms of rejection are a Truth Norm, a Belief Norm a Knowledge Norm, and a Reasonable or Justified Belief Norm. These norms link bare rejections to the epistemic attitude of the speaker. I then compare the different norms of rejection to assess the constitutive conditions for rejection. Surprisingly, among epistemic norms, a Knowledge Norm fares best. When examining differ-

ent scenarios of rejection, it turns out that speakers may reject statements they believe, are justified in believing, think reasonable to believe; but when speakers reject something they know, they appear blameworthy.

In chapter 3, I apply the study of rejection to determining the contribution of a linguistic item: the adversative marker *but*. I account for the pragmatic processing in the different accounts of *but* in terms of indirect speech acts: assertion and rejection. Doing so not only covers the different uses of the adversative markers; it also restricts pragmatic processing to exclude these uses of *but* speakers deem infelicitous. The analysis of the contribution of *but* in terms of indirect speech acts is in line with an inferential expressivist programme: it allows us to account for the meaning of *but* in terms of speaker's attitudes it expresses, and inferences the use of *but* permits.

Chapter 4 focused on assessing the relation that Incurvati and Schlöder (2019) draw between the lexical items *perhaps* and *might* and weak assertion. I did so by gathering linguistic examples from corpora for close reading to determine *perhaps* and *might* embedding behaviour. This behaviour is consistent with Incurvati and Schlöder's hypothesis that *perhaps* is a force marker for weak assertion, and *might* the corresponding modal.

Chapter 5 characterises the speech act of retraction. When speakers disagree with one of their previous utterances, they might attempt to undo its illocutionary effects by means of a retraction. In this characterisation, we insist on treating a retraction as a proposal to undo the illocutionary effects of an utterance. We then model the way that speakers undo the effects of a previous utterance. Studying retractions opens new questions concerning power imbalances in conversation: in some contexts, speakers can dispense with an explicit agreement from the audience for their retraction to take effect.

The notion of power imbalance comes at play in chapter 6. An idealised picture of conversation assumes that all participants have equal power in the conversation. It also assumes that participants cooperate in a goal to maximise information. In this idealised picture, silence can be interpreted as tacit acceptance. After all, if participants wanted to dissent, they would. However, once we bring in power imbalances, not all participants are free to explicitly dissent whenever they want to. In these cases, silence can sometimes be an expression of defiance, or dissent, where participants are not free to reject. Because of this, I propose a new interpretation of silence, according to which silence expresses a default attitude provided by the perceived cooperativeness of the conversation. This gives different type of conversational settings, where speech acts have different update behaviours: they are not accepted by default anymore. I conclude this chapter by explaining why silencing is still a problem if it does not provide a presumption of assent. The adverse effects of silencing comes from depriving participants of agency.

In chapter 7, I interrogated the Fregean category of colourings. According to Sander (2019), some colourings do not communicate the content that they

convey. This stance holds under a notion of communication where what is communicated by a speaker's utterance depends only on what the speaker *intended* to communicate. However, I build an alternative intersubjective picture of communication, where speakers can be held to communicate something they didn't intend to communicate. Because contents are generally associated with certain utterances, competent speakers who perform these utterances can be interpreted by audiences as intending to communicate the generally associated content. In this sense, when speakers use colourings, they communicate the content of these colourings even when they don't intend it.

8.2 Further Prospects

The research in this dissertation focused, one study after the other, on specific speech acts and lexical items. However, these six essays orbit the same theme: how to theorise and model conversations in contexts where speakers disagree, have different goals, or have different status. Fortunately, the moment is ripe for tackling political and social themes within philosophy of language (Khoo and Sterken, 2021). And inferential expressivism, as an approach to language that characterises meanings in terms of inferences from speakers' attitudes, constitutes an ideal theoretical standpoint to approach these themes.

Speech acts of disagreement Two chapters of this dissertation closely analysed non-assertoric speech acts that occur in contexts of disagreement. Speakers use *rejection* to express (sometimes preemptively) their disagreement with an informational content or the performance of a speech act. Speakers use *retractions*, *a posteriori*, to express that they disagree with a previous speech act they are responsible for. These are not the only two speech acts that occur in contexts of disagreements. Caponetto (2017) characterises a whole category of speech acts of *refusals* that speakers use to decline illocutionary proposals. The close study of different types of refusals, not only in terms of the conversational updates they block, but also in terms of the deontic commitments they introduce and the norms they obey, offers many exciting prospects for the analysis of conflictual conversations.

Speech acts and lexical items. Two chapters in this dissertation closely analysed the behaviour of specific lexical items in terms of speech acts. I explain the behaviour of the adversative marker *but* in terms of an assertion-denial speech structure. I relate *perhaps* and *might* to the speech act of weak assertion.

I look forward to an application of non-assertoric speech acts to other linguistic items. The analysis of *slurs*, at the interface of expressivist perspectives on meaning and political philosophy of language, is one such example. Davis and McCready (2020), for example, analyse the expressive component of *slurs* as

invoking a complex of facts and social attitudes concerning the insulted group (Davis and McCready, 2020).

I believe that a similar account could be given for positive identity terms. For example, when an individual performs a speech act of *coming out*, they join a group of individuals sharing an identity. This identity refers to a complex of facts and social attitudes concerning the group. Moreover, when an individual comes out, they propose changes to the informational and deontic structure of the conversation. They challenge a previous belief of their audience concerning their identity, and propose changes of audience behaviour with respect to them. The speech act of *coming out* may be what is at play in the expressive, community building component of identity terms (such as *lesbian, gay, trans*, etc). The speech act of *coming out* may also underlie the ability for marginalised groups to reclaim oppressive slurs. This is one possible example where a fine-grained study of specific non-assertoric speech acts may prove crucial in analysing the meaning of specific lexical items.

Towards non-ideal theory. The last take away from the work I have done during my PhD was the possibility of doing theoretical philosophy of language with non-ideal contexts, and non-ideal models. In fact, it was crucial for me, in order to study speakers attitudes in disagreement, to consider contexts where speakers do not cooperate, have different goals, or have different power. Therefore, I delineated some ways in which we can build non-ideal models of communication. I believe these open new prospects to work on speech acts and their effects while keeping in mind their possible failure and behaviour in non-ideal contexts.

Bibliography

- Asher, Nicholas (2008). “Troubles on the right frontier”. In: *Constraints in Discourse*. Ed. by Benz, Anton and Kühnlein, Peter. John Benjamins, pp. 29–52.
- Asher, Nicholas and Lascarides, Alex (2001). “Indirect Speech Acts”. In: *Synthese* 128.1, pp. 183–228.
- (2003). *Logics of Conversation*. Cambridge University Press.
- Asher, Nicholas and Vieu, Laure (2005). “Subordinating and Coordinating Discourse Relations”. In: *Lingua* 115 (4), pp. 591–610.
- Austin, John L. (1975). *How to do things with words*. Ed. by Urmson, J. O. and Sbisà, Marina. 2nd ed.
- Bach, Kent (1999). “The Myth of Conventional Implicature”. In: *Linguistics and Philosophy* 22, pp. 367–421.
- Bach, Kent and Harnish, Robert M. (1979). *Linguistic communication and speech acts*. The MIT Press.
- Baltag, Alexandru and Renne, Bryan (2016). “Dynamic Epistemic Logic”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Zalta, Edward N. Winter 2016. Metaphysics Research Lab, Stanford University.
- Beaney, Michael, ed. (1997). *The Frege Reader*. Blackwell.
- Beaver, David (2001). *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications.
- Beaver, David and Krahmer, Emiel (2001). “A Partial Account of Presupposition Projection”. In: *Journal of Logic, Language and Information* 10.2, pp. 147–182.

- Beaver, David, Roberts, Craige, et al. (2017). “Questions under Discussion: Where Information structure meets projective content”. In: *Annual Review of Linguistics* 3, pp. 265–84.
- Benton, Matthew A. and van Elswyk, Peter (2020). “Hedged Assertion”. In: *The Oxford Handbook of Assertion*. Ed. by Goldberg, Sanford C. Oxford University Press.
- Berto, Franz (2015). “A modality called ‘Negation’”. In: *Mind* 124.495, pp. 761–793.
- Bird, Alexander (2002). “Illocutionary Silencing”. In: *Pacific Philosophical Quarterly* 83 (1).
- Blackburn, Simon (1984). *Spreading the Word: Groundings in the Philosophy of Language*. Clarendon Press.
- Brandom, Robert (1983). “Asserting”. In: *Noûs* 17.4, pp. 637–650.
- (1998). *Making it Explicit*. Harvard University Press.
- Brown, Jessica and Cappelen, Herman, eds. (2011). *Assertion*. Oxford University Press.
- Brown, Penelope and Levinson, Stephen C. (1999). “Politeness: some universals in language usage”. In: *The Discourse Reader*. Ed. by Jaworski, Adam and Coupland, Nicolas. 2nd ed. Routledge, pp. 311–22.
- Bussière-Caraes, Lwenn (2021). “Silence, Dissent and Common Ground”. In: *Proceedings of the ESSLLI Student Session 2021*. Ed. by Pedersen, Mina Young and Pavlova, Alexandra. 32nd European Summer School in Logic, Language and Information. online.
- (2022). “Communicating with Colourings”. In: *Philosophical Approaches to Language and Communication*. Ed. by Stalmaszczyk, Piotr and Hinton, Martin. Vol. 2. Studies in Philosophy of Language and Linguistics. Peter Lang, pp. 151–70.
- Camp, Elisabeth (2018). “Insinuation, Common Ground, and the Conversational Record”. In: *New Work on Speech Acts*. Ed. by Fogal, Daniel, Harris, Daniel W., and Moss, Matt. Oxford University Press.
- Caponetto, Laura (2017). “On Silencing, Authority, and the Act of Refusal”. In: *Rivista di Estetica* 64 (Discrimination in Philosophy), pp. 35–52.
- (2020). “Undoing Things with Words”. In: *Synthese* 197, pp. 2399–2414.
- Cappelen, Herman (2011). “Against Assertion”. In: *Assertion*. Ed. by Brown, Jessica and Cappelen, Herman. Oxford University Press, pp. 21–47.

- Carnap, Rudolf (1934). *Logische Syntax der Sprache*. Vienna: Springer. *The Logical Syntax of Language*. Trans. by Smeaton, Amethe. London: Routledge, 1937.
- Carroll, Lewis (1895). “What the Tortoise said to Achilles”. In: *Mind* 4.14, pp. 278–280.
- Davis, Christopher and McCready, Elin (2020). “The Instability of Slurs”. In: *Grazer Philosophische Studien* 97.1, pp. 63–85.
- DeRose, Keith (2002). “Assertion, Knowledge, and Context”. In: *The Philosophical Review* 111.2, pp. 167–203.
- DiAngelo, Robin (2012). “Nothing to add: A Challenge to White Silence in Racial Discussions”. In: *Understanding and Dismantling Privilege* 2, pp. 1–17.
- Dickie, Imogen (2010). “Negation, anti-realism, and the denial defence”. In: *Philosophical Studies*.
- Draper, Paul (2017). “Atheism and Agnosticism”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Zalta, Edward N. Fall 2017. Stanford University: Metaphysics Research Lab.
- Dummett, Michael (1981). *Frege: Philosophy of Language*. Duckworth.
- (1991). *The Logical Basis of Metaphysics*. Harvard University Press.
- Dutilh Novaes, Catarina (2020). *The Dialogical Roots of Deduction: Historical, Cognitive, and Philosophical Perspectives on Reasoning*. Cambridge: Cambridge University Press.
- Egbert, Jesse, Biber, Douglas, and Gray, Bethany (2022). *Designing and Evaluating Language Corpora: A Practical Framework for Corpus Representativeness*. Cambridge: Cambridge University Press.
- Egbert, Jesse, Larsson, Tove, and Biber, Douglas (2020). *Doing Linguistics with a Corpus: Methodological Considerations for the Everyday User*. Elements in Corpus Linguistics. Cambridge: Cambridge University Press.
- Engel, Pascal (2008). “In What Sense Is Knowledge the Norm of Assertion?” In: *Grazer Philosophische Studien* 77, pp. 45–59.
- Farkas, Donka F. and Bruce, Kim B. (2010). “On Reacting to Assertions and Polar Questions”. In: *Journal of Semantics* 27.1, pp. 81–118.
- Foucault, Michel (1966). *Les mots et les choses. Une Archéologie des Sciences Humaines*. Gallimard.
- (1972). *Histoire de la Folie à l'Âge Classique*. 2nd ed. Gallimard.

- Frege, Gottlob (1879). "Begriffsschrift, selections". In: *The Frege Reader*. Ed. by Beaney, M. Oxford: Blackwell, pp. 47–78.
- (1892). "Über Sinn und Bedeutung". In: *Zeitschrift für Philosophie und Philosophische Kritik* 100, pp. 25–50.
- (1918a). "Der Gedanke". In: *Beiträge zur Philosophie des deutschen Idealismus I*, pp. 58–77. Trans. by Geach, Peter and Stoothoff, Robert H. as "Thought" (Blackwell, 1997).
Quinton, A. M. and M. as "The Thought: A Logical Inquiry" (Oxford University Press, 1956).
- (1918b). "Die Verneinung". In: *Beiträge zur Philosophie des deutschen Idealismus*. Vol. I, pp. 143–57. Repr. Beaney, Michael, ed. *The Frege Reader*. Blackwell, 1997.
- (1956). "The Thought: A Logical Inquiry". Trans. by Quinton, A. M. and M. In: *Mind* 65.259, pp. 289–311.
- (1983). "Logik". In: *Nachgelassene Schriften*. Ed. by Hermes, H. et al. Hamburg: Meiner, pp. 137–163.
- Gentzen, Gerhard (1934). "Untersuchungen über das logische Schließen. I". In: *Mathematische Zeitschrift* 39, pp. 176–210. "Investigations into logical deduction". In: *American Philosophical Quarterly* 1.4 (1964), pp. 288–306.
- Gettier, Edmund L. (1963). "Is Justified True Belief Knowledge?" In: *Analysis* 23.6, pp. 121–123.
- Gibbard, Allan (2003). *Thinking how to live*. Harvard University Press.
- Goldberg, Sanford (2015). *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford: Oxford University Press.
- (2016). "Arrogance Silence and Silencing". In: *Proceedings of the Aristotelian Society* 90, pp. 93–112.
- (2018). "Dissent: Ethics and Epistemology". In: *Voicing dissent*. Ed. by Johnson, Casey Rebecca. London: Routledge.
- (2020a). "Conversational silence". In: *Conversational Pressure*. Oxford University Press, pp. 151–186.
- (2020b). "Silence misconstrued: The Double Harm of Silencing". In: *Conversational Pressure*. Oxford University Press, pp. 151–186.
- ed. (2020c). *The Oxford Handbook of Assertion*. Oxford: Oxford University Press.
- Grice, Paul (1957). "Meaning". In: *The Philosophical Review* 66.3, pp. 377–388.

- (1961). “The Causal Theory of Perception”. In: *Proceedings of the Aristotelian Society* 35, pp. 121–152.
- (1975). “Logic and Conversation”. In: *Studies in the Way of Words*. Ed. by Grice, Paul. Harvard University Press, pp. 41–58.
- Groenendijk, Jeroen and Stokhof, Martin (1991). “Dynamic Predicate Logic”. In: *Linguistics and Philosophy* 14, pp. 39–100.
- Groenendijk, Jeroen, Stokhof, Martin, and Veltman, Frank (1996). “Coreference and Modality”. In: *The Handbook of Contemporary Semantic Theory*. Ed. by Lappin, S. Oxford: Blackwell, pp. 179–213.
- Horn, Lawrence R. (2001). *A Natural History of Negation*. The David Hume Series. Originally published: Chicago: University of Chicago Press, 1989. Stanford: CSLI Publications.
- (2013). “I Love Me Some Datives: Expressive Meaning, Free Datives, and F-Implicature”. In: *Beyond Expressives: Explorations in Use-Conditional Meaning*. Ed. by Gutzmann, D. and Gärtner, H.-M. Leiden: Brill, pp. 151–199.
- Hume, David (1748). *An Enquiry concerning Human Understanding*.
- Incurvati, Luca and Schlöder, Julian J. (2017). “Weak Rejection”. In: *Australasian Journal of Philosophy* 95.4, pp. 741–760.
- (2019). “Weak Assertion”. In: *Philosophical Quarterly* 69.277, pp. 741–770.
- (2021). “Inferential Expressivism and the Negation Problem”. In: *Oxford Studies in Metaethics*. Ed. by Shafer-Landau, Russ. Vol. 16. Oxford University Press.
- Izutsu, Mitsuko Narita (2008). “Contrast, concessive, and corrective: Toward a comprehensive study of opposition relations”. In: *Journal of Pragmatics* 40.4, pp. 646–675.
- Jasinskaja, Katja (2012). “Correction by adversative and additive markers”. In: *Lingua* 122.15. SI: Additivity and Adversativity, pp. 1899–1918.
- Jasinskaja, Katja and Zeevat, Henk (2008). “Explaining Additive, Adversative and Contrast Marking in Russian and English”. In: *Revue de Sémantique et Pragmatique* 24, pp. 65–91.
- Kaplan, David (1978). “Dthat”. In: *Syntax and Semantics* 9, pp. 221–43.
- (1979). “On the Logic of Demonstratives”. In: *Journal of Philosophical Logic* 8.1, pp. 81–98.

- Khoo, Justin and Sterken, Rachel Katharine, eds. (2021). *The Routledge Handbook of Social and Political Philosophy of Language*. Routledge Handbooks in Philosophy. Routledge.
- Knobe, Joshua and Yalcin, Seth (2014). “Epistemic modals and context: Experimental data”. In: *Semantics and Pragmatics* 7.10, pp. 1–21.
- Koyré, Alexandre (1957). *From the Closed World to the Infinite Universe*. Baltimore: John Hopkins Press.
- Krabbe, Eric C. W (2001). “The problem of Retraction in Critical Discussion”. In: *Synthese*. New Perspectives in Dialogical Logic 127.1/2, pp. 141–159.
- Krifka, Manfred (2014). “Embedding Illocutionary Acts”. In: *Recursion: Complexity in Cognition*. Ed. by Roeper, Tom and Speas, Margaret. Cham: Springer International Publishing, pp. 59–87.
- Kripke, Saul A. (2017). “‘And’ and ‘But’: A Note”. In: *Thought* 6, pp. 102–105.
- Lackey, Jennifer (2007). “Norms of Assertion”. In: *Noûs* 41.4, pp. 594–626.
- (2018). “Silence and Objecting”. In: *Voicing dissent*. Ed. by Johnson, Casey Rebecca. London: Routledge.
- Langton, Rae (1993). “Speech Acts and Unspeakable Acts”. In: *Philosophy and Public Affairs* 22 (4), pp. 293–330.
- Lasersohn, Peter (2005). “Context Dependence, Disagreement and Predicates of Personal Taste”. In: *Linguistics and Philosophy* 28.6, pp. 643–86.
- Lederman, Harvey (2018). “Uncommon Knowledge”. In: *Mind* 127.508, pp. 1069–1105.
- Leech, G.N. and Short, M.H. (2007). *Style in fiction*. 1st ed. 1981. Longman.
- Levinas, Emmanuel (1984). *Éthique et Infini: Dialogues avec Philippe Nemo*. Fayard.
- Lewis, David (1979). “Scorekeeping in a Language Game”. In: *Journal of Philosophical Logic* 8, pp. 339–59. Repr. *Philosophical Papers*. Vol. 1. Oxford: Oxford University Press, 1983.
- (1980). “Index, context, and content”. In: *Philosophy and Grammar*. Ed. by Kanger, Stig and Ohman, Sven. Dordrecht: Reidel. Repr. *Papers in Philosophical Logic*. Cambridge: Cambridge University Press, 1998.
- (1983). *Philosophical Papers*. Vol. 1. Oxford: Oxford University Press.
- (1986). *On the plurality of Worlds*. Oxford: Blackwell.
- (1998). *Papers in Philosophical Logic*. Cambridge: Cambridge University Press.

- (2002). *Convention: A Philosophical Study*. Wiley.
- Lyotard, Jean-François (1984). *Le Différend*. Critique. Les Éditions de Minuit.
- MacFarlane, John (2009). “Nonindexical contextualism”. In: *Synthese* 166.2, pp. 231–50.
- (2011a). “Epistemic modals are assessment-sensitive”. In: *Epistemic Modality*. Ed. by Weatherson, B. and Egan, A. Oxford: Oxford University Press.
- (2011b). “What is Assertion?” In: *Assertion*. Ed. by Brown, Jessica and Cappelen, Herman. Oxford University Press.
- (2014). *Assessment sensitivity. Relative truth and its applications*. Clarendon Press.
- Maitra, Ishani (2011). “Assertion, Norms and Games”. In: *Assertion*. Ed. by Brown, Jessica and Cappelen, Herman. Oxford University Press, pp. 277–96.
- Malinowski, Grzegorz (1993). *Many-Valued Logics*. Oxford: Clarendon Press.
- Mandelkern, Matthew and Dorst, Kevin (2021). “Assertion is weak”. In: *Philosopher’s Imprint*. forthcoming.
- Marques, Teresa (2014). “Relative Correctness”. In: *Philosophical Studies* 167.2, pp. 361–73.
- (2018). “Retractions”. In: *Synthese* 195, pp. 3335–3359.
- McGowan, Mary Kate et al. (2011). “A partial Defense of Illocutionary Silencing”. In: *Hypathia* 26 (1).
- McKinnon, Rachel (2013). “The Supportive Reasons Norm of Assertion”. In: *American Philosophical Quarterly* 50.2, pp. 121–35.
- (2015). *The Norms of Assertion: Truth, Lies, and Warrant*. Palgrave Macmillan.
- Mikkola, Mari (2011). “Illocution, silencing and the act of refusal”. In: *Pacific Philosophical Quarterly* 92 (3).
- Montminy, Martin (2013). “The Single Norm of Assertion”. In: *Perspectives on Pragmatics and Philosophy*. Ed. by Capone, Alessandro, Piparo, Franco Lo, and Carapezza, Marco. Cham: Springer International Publishing, pp. 35–52.
- Moss, Sarah (2018). *Probabilistic Knowledge*. OUP Oxford.
- Pagin, Peter (2016). “Problems with Norms of Assertion”. In: *Philosophy and Phenomenological Research* 93.1, pp. 178–207.

- Pagin, Peter and Marsili, Neri (2021). “Assertion”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Zalta, Edward N. Winter 2021. Stanford University: Metaphysics Research Lab.
- Pettit, Philip (2002). “Enfranchising Silence”. In: *Rules, Reasons, Norms*. Clarendon Press, pp. 367–78.
- Portner, Paul (2004). “The semantics of imperatives within a theory of clause types”. In: *Proceedings of semantics and linguistic theory*. Ed. by Watanabe, K. and Young, R. B. Vol. 14. CLC Publications.
- (2007). “Imperatives and modals”. In: *Natural Language Semantics* 15, pp. 351–383.
- Potts, Christopher (2005). *The Logic of Conventional Implicatures*. Oxford Studies in Theoretical Linguistics. Oxford University Press.
- Prawitz, Dag (1965). *Natural Deduction: A Proof-Theoretical Study*. Stockholm: Almqvist & Wicksell.
- (1971). “Ideas and Results in Proof Theory”. In: *Proceedings of the Second Scandinavian Logic Symposium*. Ed. by Fenstad, J. E. Vol. 63. Studies in Logic and the Foundations of Mathematics. Amsterdam: Elsevier, pp. 235–307.
- Price, Huw (1983). “Sense, Assertion, Dummett and Denial”. In: *Mind* 92.366, pp. 161–173.
- Price, Huw et al. (2013). *Expressivism, Pragmatism and Representationalism*. Ed. by Price, Huw. Cambridge University Press.
- Priest, Graham (2006). *In Contradiction*. Oxford University Press.
- Quine, Willard V. O. (1948). “On What there Is”. In: *The Review of Metaphysics* 2.5, pp. 21–38. Repr. *From a Logical point of view*. Harvard University Press, 1953.
- Ripley, Dave (2015). “Embedding denial”. In: *Foundations of Logical Consequence*. Ed. by Caret, Colin and Hjortland, Ole. Oxford University Press, pp. 289–309.
- Roberts, Craige (1996). “Information structure in Discourse: Towards an integrated formal theory of Pragmatics”. In: *OSU working papers in Linguistics* 49: Papers in Semantics, pp. 91–136.
- Ross, Jacob and Schroeder, Mark (2013). “Reversibility or disagreement”. In: *Mind* 122.485, pp. 43–84.
- Rumfitt, Ian (2000). ““Yes” and “No””. In: *Mind* 109.436, pp. 781–823.

- Saebø, Kjell Johan (2004). “Conversational Contrast and Conventional Parallel: Topic Implicatures and Additive Presuppositions”. In: *Journal of Semantics* 21.2, pp. 199–217.
- Sander, Thorsten (2019). “Two misconstruals of Frege’s theory of colouring”. In: *The Philosophical Quarterly* 69.275, pp. 374–392.
- Saville-Troike, Muriel (1995). “The place of silence in an integrated theory of communication”. In: *Perspective on Silence*. Ed. by Tannen, Deborah and Saville-Troike, Muriel. 2d. Ablex Publishing Corporation, pp. 3–18.
- Sbisà, Marina (2018). “Varieties of Speech Act Norms”. In: *Normativity and Variety of Speech Actions*. Ed. by Witek, Maciej and Witczak-Plisiecka, Iwona. Vol. 112. Poznań Studies in the Philosophy of the Sciences and the Humanities, pp. 23–50.
- Schlöder, Julian J. (2018). “Assertion and Rejection”. ILLC Dissertation Series DS-2018-07. PhD thesis. University of Amsterdam: Institute for Logic, Language and Computation (ILLC).
- Searle, John R. (1969). *Speech Acts: an essay in the philosophy of language*. Cambridge University Press.
- (1975). “Indirect Speech Acts”. In: *Syntax and Semantics 3: Speech Acts*. Ed. by Cole, Peter and Morgan, Jerry L., pp. 59–82.
- Sextus Empiricus (1996). *The Skeptic Way: Sextus Empiricus’s Outlines of Pyrrhonism*. Trans. by Mates, Benson. Oxford: Oxford University Press.
- Smiley, Timothy (1996). “Rejection”. In: *Analysis* 56.1, pp. 1–9.
- Stalnaker, Robert C. (1978). “Assertion”. In: *Syntax and Semantics*. Vol. 9. New York Academic Press, pp. 315–332.
- (1984). *Inquiry*. Cambridge University Press.
- (1999). *Context and Content*. Oxford University Press.
- (2002). “Common Ground”. In: *Linguistics and Philosophy* 25, pp. 701–721.
- Tanesini, Alessandra (2016). “‘Calm down, dear’: intellectual arrogance, silencing and ignorance”. In: *Aristotelian Society Supplementary Volume* 90 (1), pp. 71–92.
- (2018). “Eloquent Silences: Silence and Dissent”. In: *Voicing Dissent: The Ethics and Epistemology of Making Disagreement Public*. Ed. by Johnson, C. R. Routledge, pp. 109–128.
- Tonhauser, Judith et al. (2013). “Towards a taxonomy of projective content”. In: *Language* 89.1, pp. 66–109.

- Turri, John (2013). “The test of truth: An experimental investigation of the norm of assertion”. In: *Cognition* 129.2, pp. 279–291.
- (2016). “Knowledge and Assertion in ‘Gettier’ Cases”. In: *Philosophical Psychology* 29.5, pp. 759–775.
- Umbach, Carla (2005). “Contrast and Information Structure: A Focus-Based Analysis of but”. In: *Linguistics* 43.1, pp. 207–232. DOI: 10.1515/ling.2005.43.1.207.
- Veltman, Frank (1990). “Defaults in Update Semantics I”. In: *Conditionals, Defaults and Belief Revision*. Ed. by Kamp, H. Edinburgh: Centre for Cognitive Science, pp. 28–64.
- (1996). “Defaults in Update Semantics”. In: *Journal of Philosophical Logic* 25, pp. 221–261. Repr. Grim, P., Baynes, K., and Mar, G., eds. *The Philosophers Annual*. Vol. XIX. Atascadero: Ridgeview Publishing Company, 1998.
- Vermaire, Matthew (2020). “Retraction and testimonial justification: a new problem for the assurance view”. In: *Philosophical Studies* 177, pp. 3959–3972.
- Wales, Katie (2011). *A Dictionary of Stylistics*. 3rd ed. Longman Linguistics. Routledge.
- Weiner, Matthew (2005). “Must we know what we say?” In: *The Philosophical Review* 114.2, pp. 227–251.
- Williamson, Timothy (1996). “Knowing and Asserting”. In: *The Philosophical Review* 105.4, pp. 489–523.
- (2002). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Winterstein, Grégoire (2012). “What but-sentences argue for: An argumentative analysis of but”. In: *Lingua* 122.15. SI: Additivity and Adversativity, pp. 1864–1885.
- Wittgenstein, Ludwig (1953). *Philosophical Investigations*. Ed. by Anscombe, G. E. M. and Rhes, Rush. Trans. by Anscombe, G. E. M. Blackwell.
- (2009). *Philosophical Investigations*. Ed. by Hacker, P. M. S. and Schulte, Joachim. Trans. by G. E. M. Anscombe P. M. S. Hacker, Joachim Schulte. 4th ed. Wiley-Blackwell.

List of Figures

2.1	Updating the Common Ground with an assertion that R	16
4.1	Updating the Common Ground with a weak assertion <i>Perhaps P</i>	62
5.1	Updating the Common Ground with an assertion that R	96
5.2	Retracting an assertion that R , and asserting that <i>not R</i>	97
5.3	Toy example: assertion that R , assertion that B	98
5.4	Toy example: retracting an assertion that R after an assertion that B	99
5.5	Toy example: the effect of retractions on Common Ground and Conversational Record	101

List of Tables

2.1	Classification of theories of Assertion	13
4.1	Numerical results	69

No means No! Speech Acts in Conflict

This dissertation gathers a series of studies on speech acts that appear in contexts of disagreement, and non-cooperative conversations. Seminal works in philosophy of language often take for granted idealised models of conversations to analyse speech acts. They focus on contexts where participants share common goals, and speech acts fulfil their effect according to plan. By contrast, this dissertation asks: what happens when things don't go according to plan? In situations of disagreement, speakers use specific speech acts to prevent conversational moves from being made. Incurvati and Schlöder's *weak rejection* (2017) and *weak assertion* (2019) are such speech acts. When they wish to renege on their commitments, speakers use *retractions*, that cancel the effect of a previous utterance. The study of speech acts that express disagreement can be applied to pragmatic inferences that some linguistic items trigger. It also leads naturally to studying speech acts in non-ideal contexts.

Chapter 2 delineates the norms and normative effects of rejection on conversations. I start with analysing rejection in Stalnakerian terms, as the speech act by which a speaker blocks an update to the conversation. Rejecting P is not equivalent to asserting $\neg P$, but it prevents the assertion of P (Incurvati and Schlöder, 2017). Analysing the effect of rejections on the informational structure of conversation in stalnakerian terms explains this essential effect of rejection. It also explains the 'messiness' of rejection (Dickie, 2010; Price, 1983), due to the similarity between rejection of utterances on semantic grounds, and rejection of utterances that violate other norms of conversation (relevance, politeness, etc). But treating assertion and rejection as activities on all fours with one another requires determining the norms for the rejection of a statement. What does a speaker need to block a potential assertion? For each possible norm of assertion, I build a complementary norm for rejection. I then evaluate these different norms, and argue for a Knowledge Norm of assertion and rejection (Williamson, 1996).

Chapter 3 applies a speech acts analysis of assertion and rejection to a semantic problem: adversative markers. To find the relevant contrast that the adversative marker *but* contributes to a sentence, speakers rely on pragmatic inferences (Jasinskaja, 2012; Winterstein, 2012). But pragmatic processing needs restrictions; or it would make any *but*-sentence acceptable, by finding some tenuous contrast. I explain the pragmatic inferences speakers rely on in a principled manner. The yes-no polarity in *but*-sentences can be understood as a contrast between speaker attitudes, and thus, indirect speech acts of assertion and denial. This foray into applying speech acts to a problem at the interface of pragmatics and semantics provides a method to analyse some parts of meaning as inferences from speakers' attitudes and commitments.

Chapter 4 focuses on the lexical items corresponding to the speech act of *weak assertion*: *perhaps* and *might*. Weak assertion is the speech act complementary to rejection (Incurvati and Schlöder, 2019). While a weak rejection blocks an assertion without proposing a contradictory update to the conversation, a weak assertion blocks a strong rejection without proposing a contradictory update to the conversation. In other words: a weak assertion leaves open a possibility in the conversation. By means of corpus examples, chapter 4 provides linguistic evidence for the speech act of weak assertion, triggered by the linguistic marker *perhaps*. It also examines *perhaps* embedding behaviour, and its articulation with the epistemic possibility modal *might*.

The material in chapter 5 is based on joint work with Luca Incurvati, Giorgio Sbardolini and Julian Schlöder. In this work, we focus on the speech act of *retraction*. When a speaker wishes to cancel certain commitments they made, they can attempt to take back their utterance (Caponetto, 2020; Marques, 2018). We provide an account of retractions as a proposal to update conversational information. This account is explained *via* a Common Ground model of conversation. If the retraction proposal is accepted, the participants update the Common Ground by going back on the conversational record and cancelling the illocutionary effect of the target utterance. As such, retractions have a significant cost, and require investigating subsequent speech acts that relied on the retracted utterance. This chapter stresses the proposal aspect of retractions, that need to be accepted by participants to affect the Common Ground. In addition, it covers problematic cases of retractions, and links them to speakers commitments.

Chapter 6 applies the idea that contexts can be non-ideal, such that speaker and audience have different, or even opposing, conversational goals, to the picture of silence. A traditional picture of speech acts (Stalnaker, 1978) views silence as a default assent response, where the audience lets the speaker's discourse moves go through. However, works on political philosophy of language show how silence can express dissent (Pettit, 2002; Tanesini, 2018). I present a more nuanced picture, where silence expresses a default attitude (Saville-Troike, 1995) – which can be, but is not limited to, assent – that the speaker attributes to the audience on the grounds of the perceived cooperation level of the conversation. For example,

if the speaker deems the conversation non-cooperative, she will attribute to her audience a default attitude of dissent: unless the audience expresses assent, she will take her discourse moves to be rejected. This more fine-grained picture of conversation enables us to give a more precise picture of the effects of speech acts, that vary depending on whether the context is cooperative or not.

Chapter 7 sketches a broader notion of communication and applies it to the characterisation of Fregean colourings. In a speaker-oriented view of communication, the speaker communicates what they meant to communicate (Searle, 1969). That also means that, sometimes, the audience may make inferences, from the speaker's utterances, that the speaker did not intend. These unintended consequences of one's utterance only count as communicated content under a theory of communication that takes the audience's interpretation of the speaker's utterances into account. I sketch such a hearer-oriented theory of communication, and examine its consequences. When speakers use linguistic items that trigger specific inferences (conventional implicatures (Grice, 1975); Fregean colourings (Frege, 1879; Sander, 2019)), the audience draws these inferences, but the speaker can be held responsible for them. As such, they are communicated under an hearer-oriented picture of conversation.

With these essays, I open some avenues to study dissent and disagreement within and beyond the field of speech acts theory. By showing both how speech acts behave in conflictual settings, and impact our pragmatic understanding of linguistic items, I integrate agonistic practices of discourse in our understanding of conversation.

Nee is nee!

Taalhandelingen in conflict

Dit proefschrift bundelt een reeks studies die gaan over taalhandelingen in situaties van onenigheid en in niet-coöperatieve gesprekken. Toonaangevend werk in de taal filosofie gaat vaak uit van geïdealiseerde gespreksmodellen voor de analyse van taalhandelingen. Het concentreert zich op situaties waarin gespreksdeelnemers gemeenschappelijke doelen hebben, en waarin de taalhandelingen zoals bedoeld hun bijdrage leveren. In contrast daarmee stelt dit proefschrift de vraag: wat gebeurt er wanneer de dingen niet volgens het boekje verlopen? In situaties van onenigheid maken sprekers gebruik van bepaalde taalhandelingen om zekere conversationele zetten te verhinderen. De taalhandelingen *zwakke afwijzing* (2017) en *zwakke bewering* (2019) van Incurvati and Schlöder zijn voorbeelden daarvan. Wanneer sprekers hun conversationele aansprakelijkheden (*commitments*) willen herzien, nemen ze die terug door middel van retracties (*retractions*) die de effecten van een eerdere uiting ongedaan maken. Onderzoek naar taalhandelingen die onenigheid uitdrukken verheldert het soort pragmatische gevolgtrekkingen waar bepaalde talige elementen aanleiding toe geven, en zet op natuurlijke wijze aan tot het bestuderen van taalhandelingen in niet geïdealiseerde contexten.

Hoofdstuk 2 geeft een afbakening van de normen en de normatieve effecten van afwijzing (*rejection*) in gesprekken. Ik begin met een analyse van afwijzing in stalnakeriaanse termen, als een taalhandeling waarmee een spreker een voorgestelde voorzet in een gesprek blokkeert. Het afwijzen van P is niet equivalent aan het beweren van $\neg P$, maar het voorkomt de bewering van P (Incurvati and Schlöder, 2017). Als we het effect van afwijzingen op de informationele structuur van gesprekken analyseren in stalnakeriaanse termen dan is dit essentiële effect van afwijzing verklaard. Het verklaart ook de ‘rommeligheid’ (*‘messiness’*) van afwijzing (Dickie, 2010; Price, 1983), gezien de overeenkomst tussen het afwijzen van uitingen op semantische gronden en het afwijzen van uitingen die an-

dere gespreksnormen schenden, zoals relevantie, beleefdheid, etc. De behandeling van bewering en afwijzing als gelijkaardige activiteiten vraagt echter ook om een bepaling van de voorschriften voor het afwijzen van een uitspraak. Wat vereist het voor een spreker een mogelijke bewering te blokkeren? Voor iedere mogelijke voorwaarde voor bewering construeer ik een complementair voorschrift voor afwijzing. Vervolgens evalueer ik deze verschillende normen, en verdedig ik een Kennis Voorschrift (*Knowledge Norm*) van bewering en afwijzing (Williamson, 1996).

Hoofdstuk 3 past een taalhandelingsanalyse van bewering en afwijzing toe op een semantisch probleem: het probleem van adversatieve markering. Voor het bepalen van het relevante contrast dat de Engelse adversatieve markering *but* (Nederlands: *maar*) aan een zin toevoegt, moeten sprekers een beroep doen op pragmatische inferenties (Jasinskaja, 2012; Winterstein, 2012). Maar zulke pragmatische verwerking moet ingetoomd worden; anders zou elke *but*-zin acceptabel zijn, omdat zomaar een onbeduidend contrast altijd wel te vinden is. Ik verklaar de pragmatische inferenties waar sprekers van uitgaan op een principiële wijze. De ja-nee polariteit in *but*-zinnen kan worden begrepen als een contrast tussen spreker-houdingen (*attitudes*), en zodoende ook als indirecte taalhandelingen van bewering en afwijzing. Deze toepassing van taalhandelingen op een probleem in het grensgebied van de pragmatiek en semantiek levert een nieuwe methode op voor de analyse van onderdelen van betekenis als inferenties vanuit spreker-houdingen (*attitudes*) en -aansprakelijkheden (*commitments*).

Hoofdstuk 4 richt zich op de lexicale elementen die gepaard gaan met de taalhandeling van *zwakke bewering*: *perhaps* (Nederlands: *wellicht*) en *might* (Nederlands: *kan / zou kunnen*). Zwakke bewering is als taalhandeling complementair aan afwijzing (Incurvati and Schlöder, 2019). Terwijl een zwakke afwijzing een bewering blokkeert zonder een met die bewering tegenstrijdige update van het gesprek voor te stellen, blokkeert een zwakke bewering een sterke afwijzing zonder een met die afwijzing tegenstrijdige update van het gesprek voor te stellen. Met andere woorden: een zwakke bewering laat, net als een zwakke afwijzing, mogelijke updates van het gesprek open. Door middel van corpusvoorbeelden draagt hoofdstuk 4 taalkundig bewijsmateriaal aan voor de taalhandeling van zwakke bewering teweeggebracht door het talige element *perhaps*. Ook bestudeert het hoe *perhaps* zich inbedt in verschillende contexten, alsmede de articulatie ervan met behulp van de uitdrukking voor epistemische mogelijkheid *might*.

Het materiaal in hoofdstuk 5 is gebaseerd op gezamenlijk werk uitgevoerd met Luca Incurvati, Giorgio Sbardolini en Julian Schlöder. Dit werk is gericht op de taalhandeling van retractie (*terugtrekking*). Als sprekers bepaalde eerder gedane toezeggingen ongedaan willen maken, kunnen ze proberen om hun eerdere uitingen terug te trekken (Caponetto, 2020; Marques, 2018). We representeren retracties als voorstellen voor een aanpassing van de conversationele informatie. Deze weergave wordt uitgelegd via een Gemeenschappelijke Grond (*Common Ground*) model van gesprekken. Als een voorstel tot retractie is geaccepteerd,

dan stellen de deelnemers de Gemeenschappelijke Grond bij door terug te gaan op de gespreksgeschiedenis en door het ongedaan maken van het illocutionaire effect van de doeluiting. Retracties zijn, als zodanig, tamelijk ingrijpend, en vragen om een herziening van taalhandelingen die op de teruggetrokken uitingen kunnen hebben voortgebouwd. Dit hoofdstuk benadrukt het voorstel-aspect van retracties, die geaccepteerd moeten worden om effect te hebben op de Gemeenschappelijke Grond. Het beschrijft tevens problematische gevallen van retracties en relateert die aan de conversationele aansprakelijkheden (*commitments*) van sprekers.

Hoofdstuk 6 past het idee dat contexten ook niet-ideaal kunnen zijn, en dat de spreker en toehoorder verschillende, of zelfs tegengestelde, gespreksdoelen kunnen hebben, toe op ons beeld van stilzwijgen. In een gebruikelijk beeld van taalhandelingen (Stalnaker, 1978) wordt stilzwijgen geschetst als een standaard instemmingsreactie, waarbij de toehoorder de zinnen van de spreker in het gesprek laat passeren. Werk over de politieke filosofie van taal laat echter zien hoe stilte juist afkeuring kan uitdrukken (Pettit, 2002; Tanesini, 2018). Ik presenteer een genuanceerder beeld, waarbij stilte een standaardattitude uitdrukt (Saville-Troike, 1995) — dat instemming kan zijn, maar zich niet daartoe beperkt — die de spreker toeschrijft aan de toehoorder op basis van het ingeschatte coöperatieve gehalte van het gesprek. Als de spreker het gesprek als non-coöperatief beschouwt, zal ze haar toehoorder een standaardattitude van afwijzing toeschrijven; als de toehoorder geen instemming aangeeft, zal de spreker haar zinnen in het gesprek als afgewezen beschouwen. Dit verfijndere beeld van gesprekken stelt ons in staat om een exacter beeld van de effecten van taalhandelingen te creëren, die variëren met de coöperativiteit van de context.

Hoofdstuk 7 schetst een breder begrip van communicatie en gebruikt dit in de karakterisering van Fregeaanse schakering (*Färbung, colouring*). Volgens een op de spreker gericht begrip van communicatie communiceren sprekers wat ze bedoelen te communiceren (Searle, 1969). Dit houdt ook in dat de toehoorder, soms, uit de uitingen van de spreker conclusies kan trekken die de spreker niet bedoeld had. Zulke niet bedoelde implicaties van een uiting kunnen alleen als gecommuniceerde inhoud gelden binnen een theorie van communicatie die rekening houdt met de interpretatie van een uiting van de spreker door de toehoorder. Ik schets zo'n op de toehoorder gerichte theorie van communicatie, en bestudeer de consequenties daarvan. Wanneer sprekers talige elementen gebruiken die specifieke inferenties oproepen (conventionele implicaturen (Grice, 1975); Fregeaanse schakeringen (Frege, 1879; Sander, 2019)), dan maakt de toehoorder die inferenties, maar kan de spreker daarvoor verantwoordelijk worden gehouden. Zo worden ze gecommuniceerd in een op de toehoorder gericht begrip van gesprekken. De bijdragen in dit proefschrift leveren enkele aanzetten tot de studie van afwijzing en onenigheid binnen en buiten het gebied van de taalhandelingstheorie. Door zowel te laten zien hoe taalhandelingen in conflictsituaties werken, en hoe ze ons pragmatisch begrip van talige elementen beïnvloeden, breng ik strijdige gesprek-

spraktijken in ons begrip van conversaties onder één noemer bijeen.

Titles in the ILLC Dissertation Series:

- ILLC DS-2016-01: **Ivano A. Ciardelli**
Questions in Logic
- ILLC DS-2016-02: **Zoé Christoff**
Dynamic Logics of Networks: Information Flow and the Spread of Opinion
- ILLC DS-2016-03: **Fleur Leonie Bouwer**
What do we need to hear a beat? The influence of attention, musical abilities, and accents on the perception of metrical rhythm
- ILLC DS-2016-04: **Johannes Marti**
Interpreting Linguistic Behavior with Possible World Models
- ILLC DS-2016-05: **Phong Lê**
Learning Vector Representations for Sentences - The Recursive Deep Learning Approach
- ILLC DS-2016-06: **Gideon Maillette de Buy Wenniger**
Aligning the Foundations of Hierarchical Statistical Machine Translation
- ILLC DS-2016-07: **Andreas van Cranenburgh**
Rich Statistical Parsing and Literary Language
- ILLC DS-2016-08: **Florian Speelman**
Position-based Quantum Cryptography and Catalytic Computation
- ILLC DS-2016-09: **Teresa Piovesan**
Quantum entanglement: insights via graph parameters and conic optimization
- ILLC DS-2016-10: **Paula Henk**
Nonstandard Provability for Peano Arithmetic. A Modal Perspective
- ILLC DS-2017-01: **Paolo Galeazzi**
Play Without Regret
- ILLC DS-2017-02: **Riccardo Pinosio**
The Logic of Kant's Temporal Continuum
- ILLC DS-2017-03: **Matthijs Westera**
Exhaustivity and intonation: a unified theory
- ILLC DS-2017-04: **Giovanni Cinà**
Categories for the working modal logician
- ILLC DS-2017-05: **Shane Noah Steinert-Threlkeld**
Communication and Computation: New Questions About Compositionality

- ILLC DS-2017-06: **Peter Hawke**
The Problem of Epistemic Relevance
- ILLC DS-2017-07: **Aybüke Özgün**
Evidence in Epistemic Logic: A Topological Perspective
- ILLC DS-2017-08: **Raquel Garrido Alhama**
Computational Modelling of Artificial Language Learning: Retention, Recognition & Recurrence
- ILLC DS-2017-09: **Miloš Stanojević**
Permutation Forests for Modeling Word Order in Machine Translation
- ILLC DS-2018-01: **Berit Janssen**
Retained or Lost in Transmission? Analyzing and Predicting Stability in Dutch Folk Songs
- ILLC DS-2018-02: **Hugo Huurdeman**
Supporting the Complex Dynamics of the Information Seeking Process
- ILLC DS-2018-03: **Corina Koolen**
Reading beyond the female: The relationship between perception of author gender and literary quality
- ILLC DS-2018-04: **Jelle Bruineberg**
Anticipating Affordances: Intentionality in self-organizing brain-body-environment systems
- ILLC DS-2018-05: **Joachim Daiber**
Typologically Robust Statistical Machine Translation: Understanding and Exploiting Differences and Similarities Between Languages in Machine Translation
- ILLC DS-2018-06: **Thomas Brochhagen**
Signaling under Uncertainty
- ILLC DS-2018-07: **Julian Schlöder**
Assertion and Rejection
- ILLC DS-2018-08: **Srinivasan Arunachalam**
Quantum Algorithms and Learning Theory
- ILLC DS-2018-09: **Hugo de Holanda Cunha Nobrega**
Games for functions: Baire classes, Weihrauch degrees, transfinite computations, and ranks

- ILLC DS-2018-10: **Chenwei Shi**
Reason to Believe
- ILLC DS-2018-11: **Malvin Gattinger**
New Directions in Model Checking Dynamic Epistemic Logic
- ILLC DS-2018-12: **Julia Ilin**
Filtration Revisited: Lattices of Stable Non-Classical Logics
- ILLC DS-2018-13: **Jeroen Zuiddam**
Algebraic complexity, asymptotic spectra and entanglement polytopes
- ILLC DS-2019-01: **Carlos Vaquero**
What Makes A Performer Unique? Idiosyncrasies and commonalities in expressive music performance
- ILLC DS-2019-02: **Jort Bergfeld**
Quantum logics for expressing and proving the correctness of quantum programs
- ILLC DS-2019-03: **Andras Gilyen**
Quantum Singular Value Transformation & Its Algorithmic Applications
- ILLC DS-2019-04: **Lorenzo Galeotti**
The theory of the generalised real numbers and other topics in logic
- ILLC DS-2019-05: **Nadine Theiler**
Taking a unified perspective: Resolutions and highlighting in the semantics of attitudes and particles
- ILLC DS-2019-06: **Peter T.S. van der Gulik**
Considerations in Evolutionary Biochemistry
- ILLC DS-2019-07: **Frederik Mollerstrom Lauridsen**
Cuts and Completions: Algebraic aspects of structural proof theory
- ILLC DS-2020-01: **Mostafa Dehghani**
Learning with Imperfect Supervision for Language Understanding
- ILLC DS-2020-02: **Koen Groenland**
Quantum protocols for few-qubit devices
- ILLC DS-2020-03: **Jouke Witteveen**
Parameterized Analysis of Complexity
- ILLC DS-2020-04: **Joran van Apeldoorn**
A Quantum View on Convex Optimization

- ILLC DS-2020-05: **Tom Bannink**
Quantum and stochastic processes
- ILLC DS-2020-06: **Dieuwke Hupkes**
Hierarchy and interpretability in neural models of language processing
- ILLC DS-2020-07: **Ana Lucia Vargas Sandoval**
On the Path to the Truth: Logical & Computational Aspects of Learning
- ILLC DS-2020-08: **Philip Schulz**
Latent Variable Models for Machine Translation and How to Learn Them
- ILLC DS-2020-09: **Jasmijn Bastings**
A Tale of Two Sequences: Interpretable and Linguistically-Informed Deep Learning for Natural Language Processing
- ILLC DS-2020-10: **Arnold Kochari**
Perceiving and communicating magnitudes: Behavioral and electrophysiological studies
- ILLC DS-2020-11: **Marco Del Tredici**
Linguistic Variation in Online Communities: A Computational Perspective
- ILLC DS-2020-12: **Bastiaan van der Weij**
Experienced listeners: Modeling the influence of long-term musical exposure on rhythm perception
- ILLC DS-2020-13: **Thom van Gessel**
Questions in Context
- ILLC DS-2020-14: **Gianluca Grilletti**
Questions & Quantification: A study of first order inquisitive logic
- ILLC DS-2020-15: **Tom Schoonen**
Tales of Similarity and Imagination. A modest epistemology of possibility
- ILLC DS-2020-16: **Ilaria Canavotto**
Where Responsibility Takes You: Logics of Agency, Counterfactuals and Norms
- ILLC DS-2020-17: **Francesca Zaffora Blando**
Patterns and Probabilities: A Study in Algorithmic Randomness and Computable Learning
- ILLC DS-2021-01: **Yfke Dulek**
Delegated and Distributed Quantum Computation
- ILLC DS-2021-02: **Elbert J. Booij**
The Things Before Us: On What it Is to Be an Object

- ILLC DS-2021-03: **Seyyed Hadi Hashemi**
Modeling Users Interacting with Smart Devices
- ILLC DS-2021-04: **Sophie Arnoult**
Adjunction in Hierarchical Phrase-Based Translation
- ILLC DS-2021-05: **Cian Guilfoyle Chartier**
A Pragmatic Defense of Logical Pluralism
- ILLC DS-2021-06: **Zoi Terzopoulou**
Collective Decisions with Incomplete Individual Opinions
- ILLC DS-2021-07: **Anthia Solaki**
Logical Models for Bounded Reasoners
- ILLC DS-2021-08: **Michael Sejr Schlichtkrull**
Incorporating Structure into Neural Models for Language Processing
- ILLC DS-2021-09: **Taichi Uemura**
Abstract and Concrete Type Theories
- ILLC DS-2021-10: **Levin Hornischer**
Dynamical Systems via Domains: Toward a Unified Foundation of Symbolic and Non-symbolic Computation
- ILLC DS-2021-11: **Sirin Botan**
Strategyproof Social Choice for Restricted Domains
- ILLC DS-2021-12: **Michael Cohen**
Dynamic Introspection
- ILLC DS-2021-13: **Dazhu Li**
Formal Threads in the Social Fabric: Studies in the Logical Dynamics of Multi-Agent Interaction
- ILLC DS-2022-01: **Anna Bellomo**
Sums, Numbers and Infinity: Collections in Bolzano's Mathematics and Philosophy
- ILLC DS-2022-02: **Jan Czajkowski**
Post-Quantum Security of Hash Functions
- ILLC DS-2022-03: **Sonia Ramotowska**
Quantifying quantifier representations: Experimental studies, computational modeling, and individual differences

ILLC DS-2022-04: **Ruben Brokkelkamp**

How Close Does It Get?: From Near-Optimal Network Algorithms to Suboptimal Equilibrium Outcomes