



The Semantics and Pragmatics of slurs and thick terms

Bianca Cepollaro

► To cite this version:

Bianca Cepollaro. The Semantics and Pragmatics of slurs and thick terms. Philosophy. Université Paris sciences et lettres; Scuola normale superiore (Pise, Italie), 2017. English. NNT: 2017PSLEE003 . tel-01508856

HAL Id: tel-01508856

<https://theses.hal.science/tel-01508856>

Submitted on 14 Apr 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE DE DOCTORAT

de l'Université de recherche Paris Sciences et Lettres
PSL Research University

Préparée dans le cadre d'une cotutelle entre
Scuola Normale Superiore, Pisa
et École Normale Supérieure, Paris

La sémantique et la pragmatique des termes d'offense et des termes
éthiques épais

Ecole doctorale n°540

ÉCOLE TRANSDISCIPLINAIRE LETTRES/SCIENCES

Spécialité Philosophie

**Soutenue par Bianca
CEPOLLARO
Le 20 janvier 2017**

Dirigée par
Pier Marco BERTINETTO
et
Isidora STOJANOVIC

COMPOSITION DU JURY :

Mme. JESHION Robin
University of South California,
Rapporteur

M. VÄYRYNEN Pekka
University of Leeds, Rapporteur

Mme. BIANCHI Claudia
Università Vita-Salute San Raffaele,
Membre du jury

Mme. SBISÀ Marina
Università degli Studi di Trieste, Membre
du jury



The semantics and pragmatics of slurs and thick terms

Bianca Cepollaro

Abstract

In this thesis I develop a uniform account of slurs and thick terms in terms of presuppositions. I argue that slurs and thick terms – that were mainly studied by different disciplines – belong to the same class of ‘hybrid evaluatives’. My work is a contribution in filling the gap between the research on thick terms on the one hand and the research on slurs on the other, by showing that the mechanism underlying slurs and thick terms is one and the same and that the phenomenal differences that one can observe depend on the peculiarities of their descriptive content. In Part I, I define what counts as ‘hybrid evaluative’ and present my presuppositional account. Part II is a critical review of some alternative theories that have been put forward to account for slurs and thick terms respectively. In Part III, I propose a uniform account of two particular uses of slurs and thick terms that have been treated as independent so far, namely Appropriation of Slurs and Variability of thick terms.

Acknowledgements

Many thanks to Inés Crespo, Christopher Gauker, Michael Glanzberg, Robin Jeshion, Max Kölbel, Uriah Kriegel, Robert May, Geoffrey Nunberg, Mihaela Popa-Wyatt, François Recanati, Jennifer Saul, Marina Sbisa, Philippe Schlenker, Andrés Soria Ruiz, Benjamin Spector, Brent Strickland, Tristan Thommen, Giuliano Torrenco, Pekka Väyrynen and Julia Zakkou for engaging discussions and very helpful suggestions.

I wish to express my deepest gratitude to my supervisors Pier Marco Bertinetto and Isidora Stojanovic for being inspiring and supporting throughout.

I am especially grateful to Claudia Bianchi, who is always teaching me something.

Many thanks also to the audience and commentators in the workshops and seminars in Pisa, Malta, Bucharest, Gargnano, Paris, Trieste, Milan, Barcelona, Storrs and Dubrovnik for their precious insights and for the good time.

Special thanks to my friends and colleagues in Pisa, Paris, Milan and all around the world: I have been incredibly fortunate to meet each of you.

Heartfelt thanks to my closest friends, to my family, to my parents, Biagio and Francesca, to my brother, Carlo and to Gio.

INTRODUCTION	7
 PART I. THE PRESUPPOSITIONAL ACCOUNT OF HYBRID EVALUATIVES. 13	
 1. What counts as a hybrid evaluative	14
1.1 What slurs are about.....	17
1.2 Distinguishing thick terms from thin terms	20
1.3 Previous suggestions for a uniform account	24
 2. The projective behavior of slurs and thick terms	27
2.1 Projection.....	28
2.2 Rejection.....	32
2.3 An alternative explanation of projection	36
 3. The semantics of hybrid evaluatives	38
3.1 The evaluative content	39
3.1.1 Gradability.....	40
3.1.2 Multidimensionality	42
3.1.3 How to interpret the evaluative content	44
3.1.4 Inter-variation and intra-variation.....	46
3.2 The descriptive content	48
3.3 When things go wrong: presuppositional failure	52
3.3.1 Reference and extension.....	52
3.3.2 What is exactly presupposed	53
3.3.3 Failure.....	55
3.3.4 A theory of value	60
 4. The Dynamics of hybrid evaluatives: complicity, propaganda, negotiation.....	62
4.1 Slurs in conversation.....	63
4.1.1 The shortcut of discrimination.....	73
4.2 Thick terms: negotiation and conceptual ethics	76
 5. Defending the presuppositional account	82
5.1 Slurs and presuppositions	82
5.2 Thick terms and presuppositions.....	90
 6. Conclusion: the residual differences between slurs and thick terms	93

6.1 The semantics	94
6.2 The projective behavior	96
PART II. RIVAL THEORIES	102
7. Truth-conditional theories of slurs: It's just a matter of semantics	102
7.1. Denying some projection data	106
7.2 Explaining <i>some</i> projection data: derogation and offense.....	114
8. Kyle's semantic analysis of thick terms	118
8.1 Repeatability and redundancy	124
9. Deflationary accounts of slurs: It's just a matter of pragmatics	126
9.1 Anderson and Lepore: violating prohibitions and taboos	126
9.2 Bolinger: co-occurrence expectations and contrastive preferences.....	131
9.3 Nunberg: markedness, affiliation and manner implicatures	138
9.4 A note on markedness	142
9.5 Conclusion	146
10. Deflationary accounts of thick terms	146
11. Conventional Implicature: an alternative hybrid account.....	154
11.1 Pott's proposal.....	154
11.2 Challenges.....	157
11.3 Conclusions.....	165
PART III. WHEN THE EVALUATION CHANGES	167
12. The impermanence of evaluation	168
12.1 Appropriation of slurs	168
12.2 Variability of thick terms.....	170
13. The echoic account	171
13.1 Merely attributive and echoic: a useful labelling	175
13.2 Appropriation and Variability: residual differences.....	180
13.3 What's left of variability	181
14. Echo and perspectives	182
CONCLUSION.....	185

APPENDIX.....	186
SUMMARY.....	188
RESUME.....	194
REFERENCES.....	202

“The word ‘nigger’ incapsulates the standards of the society, and, if we were confined to it, we could not break free of those standards. But fortunately we are not so confined; our language, as we have it, can be a vehicle for new ideas.”

Hare 1963: 25.

Introduction

In the last few decades, philosophers of language and linguists have turned their attention to the evaluative and expressive dimensions of language. ‘Evaluative’ and ‘expressive’ are quite inclusive labels, ranging from expressive intensifiers like ‘damn’ to slurs like ‘wop’ and interjections like ‘shit’, from thick terms like ‘lewd’ to thin ones like ‘good’. In this dissertation I almost exclusively discuss two classes of terms for which I put forward a uniform account: slurs and thick terms.

The study of slurs and thick terms has been *mainly* conducted in two different – although related – fields: philosophy of language and linguistics for slurs and ethics and metaethics for thick terms. Despite the obvious relation between these disciplines, only few scholars adopted an interdisciplinary stance. On the one hand, the literature on thick terms focuses on issues related to philosophically central matters, such as the cognitivism/non-cognitivism dispute (i.e. whether evaluative judgements express cognitive or non-cognitive contents. This point affects the further issue of whether evaluative judgements are truth-apt), the fact/value distinction and so on. On the other hand, the debate on slurs mainly focuses on the linguistic properties of epithets such as the projection power and it investigates the linguistic relation between descriptive and pejorative content.

Väyrynen (2009, 2011, 2012, 2013) is among the first scholars to systematically apply the tools of linguistics and philosophy of language to the study of thick terms, in order to investigate the relation between description and evaluation. Väyrynen’s series of papers together with his book play a crucial role in the discipline as an example of method. Not only do they make the study of thick terms accessible to linguists and philosophers of language, but they provide the basis for the *possibility* to ask whether slurs and thick terms rely on the same linguistic mechanisms or not.

In this thesis, I develop a uniform account of slurs and thick terms, classes of terms that I shall call ‘hybrid evaluatives’. My work is a contribution in filling the gap between the research on thick terms conducted in metaethics on the one hand and the research on slurs conducted in linguistics and philosophy of language on the other. The aim of this work is to show that the mechanism that underlies slurs and thick terms is one and the same and that the phenomenal differences that we can observe depend on the peculiarities of their descriptive content. My investigation leaves aside the debate

on thick and pejorative *concepts* and focuses on thick and pejorative *terms* instead: If one wants to infer claims about thick and pejorative concepts from my work on thick terms and slurs, one can do so only under the assumption that the meaning of thick terms and slurs amounts to what the corresponding concept expresses. I do not want to commit myself to this simplified view according to which there is a one-to-one relation between language and concepts: I will just stay neutral with respect to the relation between hybrid terms on the one hand and pejorative and thick concepts on the other hand and focus on language.

The dissertation consists of three parts. Part I presents a novel presuppositional analysis of hybrid evaluatives, a category including slurs and thick terms. After some preliminaries about how to define hybrid evaluatives and how to distinguish them from related notions (chapter 1), I focus on their presuppositional behavior with respect to projection and rejection and discuss an alternative explanation of the data (chapter 2). In chapter 3, I sketch a semantics for hybrid evaluatives in which I analyze their evaluative and their descriptive content respectively. I assess the issues of reference and extension and spell out what is exactly presupposed by hybrid evaluatives. I discuss how to interpret the presuppositional failure in the case of evaluative presuppositions and show how these topics relate to highly debated questions about values in ethics. In chapter 4, I discuss the effects of the *use* of hybrid evaluatives in conversation, while investigating their function: I show that the two questions as to how slurs and thick terms encode evaluation and the question as to what effects slurs and thick terms produce are deeply related. I discuss how the complex semantics of hybrid evaluatives brings about interesting issues concerning negotiations and disagreement. Chapter 5 is dedicated to the main objections that have been raised against presuppositional accounts of slurs and of thick terms respectively. I show how these objections do not constitute an obstacle to the presuppositional approach. In chapter 6 I discuss the residual differences between slurs and thick terms. I compare certain aspects of the semantics of thick terms and slurs respectively and their projective behavior and show how the features of the semantics of slurs and thick terms explain the phenomenal divergences. In a nutshell, my analysis of hybrid evaluatives, as I presented it in Part I, is the following: Both slurs and thick terms pick out a certain descriptive property and at the same time they trigger an evaluation on

that content. For example, ‘wop’ means ‘Italian’ but at the same time it triggers the presupposition that Italians are bad because of being Italians; ‘lewd’ means ‘sexually explicit beyond conventional boundaries’ and triggers the presupposition that things or individuals that are sexually explicit beyond conventional boundaries are bad because of being so and so on. Slurs have a stronger projective power than thick terms, fact that can be explained by appealing to some features of their descriptive content and to some extra-linguistic factors.

Part II is a critical review of some alternative theories that have been put forward to account for slurs and thick terms respectively. Because slurs and thick terms have been mainly studied separately, the theories I consider often deal with only one class and neglect the other: The linguistic analysis of slurs is generally more detailed and developed than the one of thick terms. The overview I provide in this part is far from covering *all* the theories on evaluatives. Rather, I select three kinds of accounts for their relevance to specific aspects of my argument. In particular, I leave aside expressivist accounts of evaluatives and I do not assess the complicated debate between cognitivists and non-cognitivists. The first four chapters are dedicated to the “single-source” approaches, as Jeshion (2016a) calls them: the truth-conditional approach (chapters 7-8) and the deflationary approach (chapters 9-10). In chapter 7 I discuss the truth-conditional analysis of slurs, by presenting the proposal of Hom (2008, 2010, 2012) and Hom and May (2013, 2014, forthcoming); in chapter 8, I discuss the truth-conditional analysis of thick terms formulated by Kyle (2013). The main problem with the truth-conditional approaches is to account for the projection of the evaluative content. I show that the strategies employed to explain projection without appealing to non-truth-conditional meaning fail to account for the phenomena. In particular, I show (i) that Hom and May’s data in support of a non-pejorative readings of embedded slurs actually stem from metalinguistic readings and (ii) that pragmatic mechanisms (both Hom and May’s offense and Kyle’s conversational implicatures) cannot explain projection data. Therefore, the truth-conditional analysis of hybrid evaluatives fails to explain the basic data and should be abandoned. However, I endorse the following claim of the truth-conditional accounts of hybrid evaluatives: The evaluative content is linguistically encoded in the meanings of such terms and it does not arise only on the basis of pragmatic mechanisms. Chapters 9-10

are dedicated to the deflationary account of slurs and thick terms, that analyze the evaluative content associated with hybrid evaluatives in terms of pragmatic implications. In chapter 9, dedicated to slurs, I present and discuss the so-called deflationary accounts of slurs (following the label that Anderson and Lepore 2013a gave to their own approach), according to which slurs have a fairly unexceptional semantics. I include in this category all those accounts according to which the derogatory content associated with slurs is not part of their encoded meaning (again, broadly intended: not just truth conditions, but also conventional implicatures, presuppositions and the like). I present three main approaches developed by Anderson and Lepore (2013a, 2013b), Bolinger (2015) and Nunberg (forthcoming), that focus on three notions respectively: taboo, contrastive preference and manner implicatures. I will discuss the specific problems met by each account, but I show that in general the deflationary accounts of slurs meet some difficulties in explaining the genesis of slurs, the complicity phenomenon (i.e. the fact that the pejorative content associated with slurs does not seem to be ascribed to the speaker *only*) and the contrast between slurs and other ‘affiliatory’ terms. Chapter 10 is dedicated to the deflationary account of thick terms, developed by Pekka Väyrynen, according to which the evaluation associated with thick terms (*T-evaluation* in Väyrynen’s terms) conversationally arises as a pragmatic implication. These implications are not part of the asserted content, nor are they the main point of the utterance in literal uses of thick terms in normal contexts: In other words, they are typically backgrounded. The pragmatic account has the advantage of being more parsimonious than any other account of thick terms, but it fails to recognize some instances of projection of the evaluative content (the alleged *defeasibility* data). In sum, in the first four chapters of Part II, I show that the single-source approaches fail to properly account for the behavior of hybrid evaluatives and I argue that the difficulties met by these proposals speak in favor of a hybrid approach. In chapter 11 I discuss an alternative hybrid approach, the conventional implicature view: It agrees with truth-conditional accounts that the derogatory content of epithets is linguistically encoded, but it holds with pragmatic theories that the pejorative content does not belong to the truth-conditional dimension of meaning. Such an approach, that analyzes the derogatory content of slurs in terms of conventional implicatures, was proposed by Christopher Potts (Potts 2005, 2007) and developed by Eric McCready

(McCready 2010) and Daniel Gutzmann (Gutzmann 2011). It was also supported *i.a.* by Williamson (2009) and Whiting (2013). While the previous approaches I presented have been independently developed for both slurs and thick terms, the conventional implicatures view was proposed only for slurs:¹ The alleged defeasibility of the evaluative content of thick terms made scholars discard a conventional implicature approach. After presenting the account and showing how similar it is to the presuppositional approach (presupposition and conventional implicature being, in Abbott 2006’s words “close neighbor[s] in the linguistic literature”), I point out the aspects for which the presuppositional view should be preferred.

In Part III I propose a uniform account of two phenomena that have been treated as independent so far: Appropriation of slurs and Variability of thick terms. My proposal, designed to account for Appropriation and Variability within my presuppositional theory of hybrid evaluatives, is actually compatible with possibly any account of slurs and thick terms. In chapter 12 I discuss Appropriation, a particular use of slurs, and Variability, a particular use of thick terms. In chapter 13, I propose a uniform account of Appropriation and Variability, relying on the notion of “echoic use” of language, introduced by Relevance Theory (Sperber and Wilson, 1986, Carston, 1996, 2002, Wilson, 2006a, 2006b, 2007 and Wilson and Sperber, 2012) and employed by Bianchi (2014) to analyze the appropriated uses of slurs. I will argue that such an echoic approach does not only explain the Appropriation of slurs, as Bianchi claims, but it also accounts for the Variability of thick terms. I show how the relevance-theoretic distinction between ‘merely attributive’ and ‘echoic’ uses of language sheds light on the crucial distinction between the two often conflated phenomena of ‘suspension’ and ‘reversal’ of the evaluative content of slurs and thick terms. In chapter 14, I discuss the theoretical outcomes of my proposal for the theory of thick terms: If Variability cases are in fact derivative or parasitic uses of language, they do not count anymore as a counterexample to the claim that thick terms convey or express evaluation in a

¹ Copp (2001, 2009) did propose an analysis of moral terms based on conventional implicatures. However, he characterizes conventional implicatures in such a way that they turn out to be very different from Grice’s ones and from Pott’s ones. Copp himself decided to distinguish the inferences he had in mind from conventional implicatures by introducing the label of ‘conventional *simplicature*’: “The conventional implicature view was my proposed way to explicate realist-expressivism. It turns out that it needs to be replaced with the conventional *simplicature* view” (Copp 2009: 178).

systematic way. This suggests once more to consider slurs and thick terms as similar phenomena. Finally, I conclude by suggesting that the relevance-theoretic notions of “attributive” and “echoic” uses of language seem to be related to similar notions such as “perspectival shift” (Harris and Potts 2009a, 2009b and Harris 2012) and “polyphony” (Ducrot 1984 and Recanati 2006). The phenomenon at stake – namely the possibility to attribute a certain content to someone else, while possibly expressing one’s attitude towards it – is a very general feature of human language that goes beyond evaluative terms and deserves deeper investigation.

The picture that emerges from this work is that hybrid evaluatives represent a device through which language implicitly conveys linguistically encoded evaluations. In my account, hybrid evaluatives rely on presuppositions, which are – in Chilton (2004)’s words – “at least one micro-mechanism in language use which contributes to the building of a consensual reality”. By employing these terms, we implicitly take for granted a certain moral perspective, a certain set of beliefs concerning what is good and what is bad (an “ethos”, as Gibbard 2003b calls it). We implicitly apply a certain lens to the world and expect everyone else to do the same. Because the presupposed content is presented as not open to discussion, if it is not objected, it has the potential to shape contexts. In this sense, using hybrid evaluatives is a powerful tool through which language not only encodes evaluation, but is also able to impose it. Talking about the stereotypes evoked by slurs, Nunberg (forthcoming) appeals to the notion of “shortcut”: “Stereotypes, negative and positive, are among the cognitive shortcuts we rely on to make sense of the world and to guide our responses to it”. I argue that this is true not just for stereotypes, but for hybrid evaluatives in general, as they are devices through which language can convey evaluations in a way that is both linguistically encoded and implicit (see Sbisà 2007).

In this work, I am going to mention (though not use) several English slurs and a few other bad words, as a study on pejoratives cannot be conducted without constructing examples featuring these terms. I hope this will not offend the sensibilities of the readers. Since the readers might not be familiar with the slurs that I mention, the dissertation also includes an Appendix, a list of the slurs that feature in this work.

Part I. The Presuppositional Account of Hybrid Evaluatives²

The aim of this part is to develop a unified account of two classes of terms that have caught the attention of scholars in philosophy of language, linguistics and metaethics: slurs and thick terms. This interest is due to the hybrid nature of these terms, as they are believed to encode a descriptive and an evaluative component at the same time. The label ‘hybrid evaluative’, introduced in Cepollaro and Stojanovic (2016) works as a general term that covers both slurs and thick terms, leaving it open which other sorts of expression (if any) may count as hybrid evaluative. The central idea of the unified account of slurs and thick terms is the following: Both classes of expressions trigger a certain kind of evaluative presupposition. Both types of hybrid evaluatives are descriptive at the level of truth conditions, but presuppose an evaluative content. Thus, for example, the adjective ‘lewd’ has roughly the same truth-conditional content as ‘sexually explicit beyond conventional boundaries’, but at the same time it triggers the presupposition that things that are sexually explicit beyond conventional boundaries are bad because of being so. Along similar lines, I analyze a slur like ‘wop’ as having the same truth conditions as ‘Italian’ but triggering at the same time the presupposition that Italians are bad because of being Italian.

According to this view, slurs appear to be very similar to a particular kind of thick terms, the so-called ‘objectionable’ thick terms. Many authors (Eklund 2011, 2013, Harcourt and Thomas 2013, Kyle 2013, Väyrynen 2013) distinguish between objectionable and non-objectionable thick terms: The former are terms that convey an evaluation that speakers are *not* willing to accept as appropriate or warranted; the latter are terms “that embody values that ought to be rejected” (Kyle 2013: 13). Note that this feature – being objectionable or not – is completely independent from the evaluation being negative or positive. For example, a speaker may well share the negative evaluation conveyed by ‘brutal’, but reject the one conveyed by ‘lewd’ and, on the other hand, share the positive evaluation conveyed by ‘generous’, but reject the one conveyed by ‘chaste’. These two features, namely *polarity* – the conveyed evaluation being positive (P) or negative (N) – and *objectability* – the conveyed evaluation being warranted (W) or unwarranted (U) – give rise to 4 combinations.

² This part is largely based on Cepollaro (2015) and Cepollaro and Stojanovic (2016).

While thick terms seem to admit all 4 combinations (PW ‘generous’, PU ‘chaste’, NW ‘brutal’, NU ‘lewd’), slurs tend to instantiate NU only, i.e. they typically convey negative and objectionable evaluations (we shall come back to this point in section 1.1). Unlike the polarity of the evaluation, the objectionable/non-objectionable distinction is not lexically encoded, but depends on the sets of values endorsed by individual speakers.

The first chapter of this part is dedicated to settling some preliminary issues such as how to establish whether a term can count as ‘hybrid evaluative’. In the second chapter, I present and discuss the presuppositional behavior of hybrid evaluatives. In chapter 3 I focus on the semantics of hybrid evaluatives. Chapter 4 is dedicated to the discussion of the way in which these terms work in conversation: I will show how the presuppositional theory can satisfactorily account for certain peculiar features of slurs and thick terms. In chapter 5 I present the objections raised against the presuppositional account of slurs and thick terms and I show that they do not constitute an obstacle for my account. In chapter 6 I discuss the residual differences between slurs and thick terms and show how such phenomenal divergences are after all superficial.

1. What counts as a hybrid evaluative

‘Evaluative’, ‘expressive’, ‘normative’, ‘pejorative’, ‘slur’, ‘thick’, ‘thin’ are widespread labels in metaethics, philosophy of language and linguistics and yet there is not much consensus about how to understand these notions. In this chapter I put forward a criterion to distinguish hybrid evaluatives from descriptive terms that can be used evaluatively. Moreover, I elaborate on the controversial notion of ‘group’ involved in the definition of slurs (section 1.1). I discuss the problematic thin-thick distinction by relying on the distinction between at-issue vs. not-at-issue content and I underline some difficulties of obtaining a clear-cut divide (section 1.2). Finally, I briefly discuss some proposals that takes into consideration the idea that slurs and thick terms can be analyzed along the same lines (section 1.3).

In the literature about thick terms, the paradigmatic examples are expressions such as ‘lewd’, ‘generous’, ‘courageous’, ‘chaste’, ‘cruel’. The notion of a thick term has been first introduced in metaethics, and it is in the domain of metaethics that it has

been mainly discussed (Hare 1963, Williams 1985, Blackburn 1992, Gibbard 1992, among others). The defining characteristic of thick terms is that they combine descriptive meaning with evaluative meaning. The descriptive meaning can be typically paraphrased: The descriptive content of ‘lewd’ is something like ‘sexually explicit beyond conventional boundaries’; ‘generous’ descriptively corresponds to something like ‘willing to give without expectation of compensation’, ‘chaste’ to ‘abstaining from sexual intercourse’ and so on. In addition to this descriptive content, thick terms also convey (or express, depending on what theory one favors) a value-judgment about such an object or individual. Thus in describing a person as ‘lewd’ speakers usually imply something negative about the person, while in describing an act as ‘generous’, they typically convey a positive evaluation of the act at stake. Scholars investigate how the use of a thick term is associated with evaluative contents, but they disagree about what kind of semantic and pragmatic mechanisms make it possible to express or convey evaluations.

In recent years, philosophers and linguists have focused their attention on another class of terms that arguably mix description with evaluation, namely slurs. Slurs are derogatory terms targeting individuals and groups on the basis of their belonging to a certain category: They target on the basis of sexual orientation (terms such as ‘faggot’, targeting homosexual men; or ‘dyke’, targeting homosexual women, etc.), nationality (‘wop’, targeting Italians; ‘boche’ targeting Germans, etc.); ethnicity (‘nigger’ targeting black people, ‘chink’ targeting Asian people, especially Chinese etc.) and so on and so forth. Again, the questions that these terms raise concern whether and how they encode evaluation, together with related issues such as: Whether these expressions allow non-offensive uses (Brontsema 2004; Bianchi 2014; Croom 2014), whether they can be mentioned without provoking offense (Hornsby 2001; Anderson and Lepore 2013a, 2013b), whether they constitute a legitimate and uniform lexical category (Croom 2011, Nunberg forthcoming).

It is not always clear how to determine if a certain lexical item should count as a slur (or as a thick term) or not. For instance, there is some discussion whether terms such as ‘athletic’ should count as thick terms (see e.g. Eklund 2011: 37) and whether

terms such as ‘gypsy’ should count as slurs.³ The criterion we put forward in Cepollaro and Stojanovic (2016) is the following: In addition to descriptive content, hybrid evaluatives systematically convey some evaluative content that scopes out when they are embedded under negation, conditionals, modals and questions. In chapter 2 we will see in greater detail that thick terms such as ‘lewd’ and slurs such as ‘wop’ meet these requirements. This helps settle the issue whether a term like ‘athletic’ should count as evaluative: If the term has literal uses that do not convey any evaluation at all, then the term should not count as a hybrid evaluative. And indeed, when one looks into the Corpus of Contemporary American English (a.k.a. COCA; Davies 2008-), one finds occurrences on which ‘athletic’ simply means ‘related to sports’ as well as occurrences on which ‘athletic’ also appears to imply that being related to sports is a good thing:

1. In the Roman context, the heavy athletic disciplines were wrestling, boxing, and the pankration, a brutal blend of wrestling and boxing.
2. I wasn’t a particularly athletic or popular child.
3. Everybody in cirque is athletic and handsome or beautiful.

The same goes for slurs: There are terms that have derogatory and non-derogatory occurrences, i.e. terms that work *as* slurs or not, depending on the context. Here are some examples of the term ‘gypsy’ from COCA:

4. On the street, a violinist plays one of Bartok’s Gypsy melodies. The vibrato from his violin feels as if it is penetrating my heart.
5. Calling that green-eyed minx a slut is too good for her. They say she’s got Gypsy blood!

³ For example, there was a debate about ‘gypsy’ within the BBC, as The Guardian reports: <http://www.theguardian.com/media/organgrinder/2009/sep/24/gypsy-racist-ben-miller>.

In (2)-(3) ‘athletic’ conveys a positive evaluation toward sport-related things or people, in the same way in which a thick term would do; similarly, ‘gypsy’ appears to work like a slur in (5); however, this is not a systematic and encoded feature of these terms. As (1) and (4) show, the evaluative content that is sometimes associated with expressions such as ‘athletic’ or ‘gypsy’ is not part of its conventional meaning, which shows, in turn, that these terms, *qua* expressions, are not to be counted among hybrid evaluatives.

Note that the criterion that I have just provided does not *prima facie* include the so-called multivalent thick terms (see Dancy 1995), i.e. those terms that seem to convey positive and negative content at the same time: ‘eccentric’, ‘extravagant’, ‘kinky’, ‘unorthodox’ and the like. Similarly, my account would not apply to the terms that do not lexically encode evaluation, but can be used in an evaluative way both positively and negatively, for instance ‘intense’ (see Stojanovic 2016a, 2016b about “valence-underspecification”).

1.1 What slurs are about

Slurs are taken to denigrate groups or individuals on the basis of their belonging to a certain class. Scholars do not usually specify what counts as a target class, instead they present some examples involving ethnic or geographical origin, religion, gender, sexual orientation. Nunberg (forthcoming) calls these categories “the deep fatalities that have historically been the focus of discrimination or social antagonisms that we see as rents in the fabric of civil society”. As a result, there is no consensus whether certain terms, especially those that are not about the listed typical categories, can count as slurs. Take for example ‘foureyes’ for people wearing glasses, or ‘stinkpotter’, for owners of a motorized boat – “Sailing enthusiasts deprecate the owners of motor craft as ‘stinkpotters’ but we probably wouldn’t call the word a slur” – Nunberg (forthcoming). Nunberg claims that it was a sign of methodological incuriosity on the side of scholars to focus on a restricted set of examples exclusively including the worst slurs like ‘faggot’ and ‘nigger’. Such a restriction put at the center of the debate certain features that do not seem to necessarily characterize slurs as a whole.

“Someone says “That building is full of flacks [publicists]” or “Mes collègues sont tous fachos”. (...) I don’t think many people would want to argue that those utterances aren’t truth-evaluable, or that they’re purely expressive, or that they’re useless to us, all claims that people have made about words like nigger. Not that the claims about the use of that word are wrong—that’s another question—but they seem to apply only to words that convey unfounded or indefensible contempt for the members of a racial or ethnic group, which make for a poor candidate for a universal linguistic type.” (Nunberg forthcoming)

I agree with Nunberg and my proposal predicts that pejorative terms like ‘flack’ for ‘publicist’ and ‘*facho*’ for ‘fascist’, just like ‘foureyes’ and ‘stinkpotter’, can in fact count as slurs, as long as they fulfill the criterion I have proposed in the previous section, even if the categories they target are not prototypical. This brings up a very important issue concerning the descriptive content of slurs, namely what categories can count as a target group. Scholars have attempted to define what could count as a target group. By observing existing slurs in English, we can already discard the hypothesis that for a property to determine a target group, there must be some kind of self-identification among the people instantiating such a property. It is probably true for properties like ‘being Italian’ or ‘being Jewish’ targeted by the slurs ‘wop’ and ‘kike’, but epithets like ‘gook’ do not involve a property with which people can self-identify: ‘gook’ was used to call *all* natives of the regions occupied by the US army (Nicaraguans, Costa Ricans, Filipinos), who do not self-identify with each other (see Jeshion 2016b). Hom and May (forthcoming) have formulated three options to establish what count as a group G, the target of a pejorative:

“1) There are no restrictions on G; it can be instantiated by any group whatsoever. This is in effect the claim that group membership is not something that is morally evaluable.

2) There is a restriction on G supplied by a theory of natural groups. This theory would isolate racial, religious, gender, sexual orientation, etc. as natural groups, and hence as targets of pejoration.

3) There is a restriction on G provided by ideologies that are active in sociocultural contexts. A group could be a value of G only insofar as there is a discriminatory cultural norm that supports it.” (Hom and May forthcoming).

The authors embrace option (3), while keeping the idea expressed in option (1) that group membership is not a morally evaluable feature:

“What are the criteria for choices of G such that there will be a pejorative term with the meaning (sense) PEJ(G)? – (...) it is reserved for groups that for whatever odious reasons have associated with them an unjust, hateful or discriminatory ideology that is culturally ingrained within society.” (*Ibidem*)

I agree with the authors about rejecting option (2) as a non-starter, but I do not embrace their solution. In my account *any* property can individuate a target group, as long as some speakers find it interesting or relevant to pick out such a property and convey an evaluation about the objects instantiating it. Note that this criterion does not put any requirement on whether the evaluation is appropriate or not. My criterion to define a slur differs from Hom and May’s (and from most accounts of slurs) in that (i) slurs do not have to be wrong *by definition*, in the sense that it is not part of how I define a slur that the triggered evaluation is wrong and “unjust”,⁴ (ii) a group of speakers coining a term to express a systematic evaluation on people instantiating a targeted property is enough to have a slur, it is not necessary to have a culturally ingrained discriminatory ideology within society and (iii) slurs do not have to convey a necessarily negative “hateful” evaluation; it is conceivable that a slur might convey a positive evaluation. An example of what a slur with a positive polarity might look like is ‘Aryan’: It was used by Nazis to talk about Indo-Europeans, while conveying some positive evaluation, supporting the idea that being Indo-European is good in itself.

Note that once we acknowledge that the target class of a slur could be any (it does

⁴ For a discussion of the idea that the evaluation of slurs might be well-focused, that is, that in some cases the target class might deserve the evaluation conveyed (prototypical examples are ‘Nazi’, or the French ‘*facho*’), see Predelli (2010) and Mišćević (2011).

not necessarily have to do with nationality, ethnicity etc.) and in principle also the polarity of the evaluation could be any, it is easier to conceive a unified analysis of slurs and thick terms. For the time being, what needs to be stressed is that most slurs are both negative and unjust, as they typically convey a negative evaluation of their targets on the basis of characteristics that in no way ground or justify derogation, exclusion etc. However, I shall underline that the basic linguistic mechanism through which slurs encode evaluative content does not require that the evaluation conveyed by the slur is negative, nor that it is unjust. We shall discuss the effects of slurs and their power to not only reflect, but also promote discrimination in section 4.1.

1.2 Distinguishing thick terms from thin terms

In the previous section I have proposed a criterion that allows to distinguish hybrid evaluatives from ordinary descriptive terms that, in a suitable context, can be *used evaluatively*. A further question is how to distinguish between different kinds of evaluatives, for instance how to distinguish *hybrid* evaluatives from *pure* evaluatives. In other words, how do we draw the distinction between thick terms, such as ‘courageous’ or ‘lewd’, and thin terms, such as ‘good’ and ‘bad’? The thin vs. thick distinction, formulated in metaethics (Williams 1985), is far from being uncontroversial, even though it is widely accepted. Eklund (2011: 40) claims that if we analyze the evaluative content of thick terms as conventional implicatures (or presuppositions, I would add), then we end up drawing a “dramatic difference” between thin and thick terms, since for the former the evaluation belongs to the Fregean sense, whereas for thick terms it would belong to the Fregean tone. Some major difficulties raised by such a distinction are discussed in Väyrynen (2013), who defends the view that the so-called thick terms, such as ‘courageous’ or ‘lewd’, do not encode evaluations. According to Väyrynen, the evaluative content associated with thick terms arises from conversational mechanisms that can vary on a case-by-case basis. As argued in Cepollaro and Stojanovic (2016), we defend on the contrary the view according to which those evaluations are recorded in the meaning of those terms, to the extent that the terms trigger evaluative presuppositions. Thus our view has to answer the question as to what distinguishes pure evaluatives from hybrid ones.

One way of distinguishing thin and thick terms would be to say that in the case of thin terms, the evaluative content is at-issue content, while in the case of thick terms and slurs, evaluative content is *presupposed*, hence not-at-issue, and what is at issue is the descriptive content. The at-issue vs. not-at-issue distinction, widely accepted in discourse analysis,⁵ amounts to the idea that utterances typically have something like a ‘main point’, i.e. a question under discussion (QUD); we call ‘at-issue’ those propositions associated with an utterance that are relevant to the QUD, in the sense that they entail “a partial or complete answer to the QUD” (Simons et. al 2010). Every proposition associated with an utterance that is *not* relevant to the QUD is labeled as ‘not-at-issue’. Presuppositions are typically *not* relevant to the QUD, as they amount to content that is taken for granted. Going back to the thin vs. thick terms distinction, the idea is that evaluation is not-at-issue (it is taken for granted) in the case of thick terms, whereas it is at-issue (it is the ‘main point’, relevant to the QUD) in the case of thin terms. Although the idea sounds plausible to a certain extent, the picture appears to be quite complicated. For one thing, the evaluative content can *become* precisely what is at issue when using a hybrid evaluative. Consider:

6.

A: Madonna’s show was lewd.

B: I disagree. It was sexually explicit beyond conventional boundaries, indeed, but it was not bad in any way.

However, this is not surprising after all: As we will see in greater detail in chapter 4, presupposed content can become at-issue under the appropriate circumstances. We shall see how the cases of negotiation involving metalinguistic disagreement – like (6) – are particularly interesting for evaluatives. For the time being, I will consider other examples that seem to challenge the idea of a sharp divide between thick and thin terms.

⁵ See for instance Roberts (1996), Potts (2005, 2007b, 2007c, 2012), Simons et al. (2010), Tonhauser (2012).

The first problems are raised by those terms that are considered on the edge between thin and thick, such as ‘just’. ‘Just’ seems to have *some* descriptive content, but very ‘slim’. This impression of being on the edge between the thick and the thin can be explained in terms of how specific the descriptive content is. For example, the descriptive at-issue content of ‘just’ is not very specific nor rich, it amounts to something like ‘in accordance of standards and requirements’, which is quite vague. On the contrary, most thick terms (and the same goes for slurs) are specific in their at-issue content: For example ‘lewd’, meaning ‘sexually explicit beyond conventional boundaries’, ‘wop’ meaning ‘Italian’. However, one could simply acknowledge that some thick terms have a quite vague descriptive content, without taking this to be evidence that the linguistic mechanism through which thin and thick terms convey evaluation is the same.

In Cepollaro and Stojanovic (2016), we argued that further difficulties for postulating a sharp boundary between the thin and the thick stem from those terms that are very close to the all-purpose evaluatives ‘good’ and ‘bad’ in that they (i) are typically used to convey evaluative content and (ii) do not have a rich descriptive content in addition to the evaluative one, but seem however ‘thicker’ than ‘good’ and ‘bad’. We considered two cases. The first one concerned highly positive and highly negative expressions, such as ‘awesome’, ‘excellent’, ‘fantastic’, ‘magnificent’ on the one hand and terms like ‘awful’, ‘horrible’, ‘terrible’ on the other hand. The above terms look like intensified thin terms (something like ‘good (or bad) to a very high degree’). However, there seems to be some additional lexical information that distinguishes those expressions among each other, that is, *some* descriptive meaning: ‘Excellent’ is for what is superior to the rest in a certain area, ‘terrible’ is for what provokes terror and so on.

The second case we considered concerns those terms that look like thin terms, but relative to a specific domain. That is the case of ‘beautiful’, ‘ugly’, or ‘evil’. As a matter of fact, when discussed within the realm of aesthetics, ‘beautiful’ and ‘ugly’ are taken to be *thin* terms. Similarly, ‘evil’, would likely count as thin in ethics, as it means something along the lines of ‘morally bad’. At the same time, if instead of thinking about evaluatives from a specific domain, we consider them from a broader perspective, ‘beautiful’ and ‘ugly’ are narrower than the all-purpose ‘good’ and ‘bad’,

as they constrain the evaluation to *aesthetic* evaluation and ‘evil’ is narrower than ‘bad’, as the negative evaluation that it encodes is linguistically constrained to *moral* evaluation.

We drew the conclusion that the most reasonable thing to say is to say that the terms like ‘excellent’ and ‘horrible’, ‘beautiful’ and ‘ugly’ are *thicker* than ‘good’ and ‘bad’, but *thinner* than ‘courageous’ or ‘cruel’.

However, the cases we presented do not inevitably knock down the idea of a distinction between thin and thick terms. Let me sketch a tentative alternative explanation. One could analyze ‘beautiful’ and ‘ugly’ as thin terms *restricted* to a certain domain, corresponding to ‘aesthetically good’ and ‘aesthetically bad’. Similarly, as we said, one could analyze ‘excellent’ and ‘awful’ as intensified evaluatives, corresponding to ‘very good’ and ‘very bad’. In these terms, ‘beautiful’, ‘ugly’, ‘excellent’ and ‘awful’ appear to be *more complex* than ‘good’ and ‘bad’, but not necessarily *thicker*. Compare them to comparatives and superlatives, for example; we would not say that ‘better’ is thicker than ‘good’, only that it is more complex, as it is a modification of ‘good’ that roughly means ‘*more* good’. The same goes for ‘*very* good’ (‘excellent’): It is a modification (intensification, in this case) of ‘good’. Presumably we could say the same also for ‘*aesthetically* good’ (‘beautiful’) and ‘*morally* bad’ (‘evil’). These terms (‘better’, ‘beautiful’, ‘evil’, etc.) do not mix description and evaluation. They remain thin terms, whose content has been somehow modified, either by a comparative or an intensifier (as in the case of ‘better’ and ‘excellent’), or by a domain restrictor (as in the case of ‘beautiful’ and ‘evil’), without becoming thick. They can be analyzed as (modified) thin terms, that are still distinct from thick terms in that the latter mix descriptive and evaluative content. This is just a sketchy suggestion of one way in which one can try to resist the skeptic attitude with respect to the thin/thick distinction discussed in Cepollaro and Stojanovic (2016).

Note that such a tentative explanation does not take into account certain uses of predicates like ‘excellent’: ‘Excellent’ is a thin term only when it is used as intensifier of ‘good’; if ‘excellent’ is used to call something that is superior to the rest in some respects (i.e. etymologically), then it would not count as a thin term. One could provide a story like the following: Terms like ‘excellent’ and ‘terrible’ were

originally descriptive for ‘superior to the rest in some respects’ and ‘provoking terror’; because it is typically very good to be superior to the rest in some respects and it is typically very bad to provoke terror, those terms were also employed as intensified thin terms, which is presumably their most common use today. This approach has the advantage of maintaining the distinction between the thick and the thin, but it has to postulate some kind of polysemy for terms like ‘excellent’ and ‘terrible’, which is on the one hand theoretically costly, but on the other it is justified by the existence of two different uses of the terms: the etymological one and the intensifier one.

In chapter 3, I shall come back to the semantics of thin and thick terms, in order to investigate similarities and differences in more detail.

1.3 Previous suggestions for a uniform account

The idea that slurs and thick terms can be analyzed along the same lines was hinted at in some previous works. Even though there is no systematic and fully-fledged account, different authors over the years have suggested a comparison between slurs and thick terms, in different ways.⁶ In this section, I will briefly mention these works, starting from scholars working on thick terms and moving to scholars working on slurs.

Hare (1963) talks about terms and concepts that mix description and evaluation. In discussing whether it is necessary to endorse a certain moral perspective in order to understand this kind of terms (see section 3.3), he mentions two such expressions: ‘courageous’ and ‘nigger’ (Hare 1963: 187). He endorses the idea that both expressions convey at the same time evaluative and descriptive content and that, when speakers do not share such evaluations, they have to abandon these terms and substitute them with neutral counterparts.⁷ He also claims that speakers are more accurate in perceiving the evaluative content that they do not endorse, therefore suggesting that for most philosophers it is easier to perceive the evaluative content of the slur ‘nigger’ than the evaluative content of ‘generous’ (see section 6.2).

⁶ I will not discuss Hay (2013), as his proposal does not concern thick terms, but mainly *thin* ones and he analyzes them along the lines of general non-slurring insults (‘jerk’, ‘asshole’, etc.).

⁷ Hornsby (2001) discusses some aspects of Hare’s theory of evaluatives that resemble her notion of ‘uselessness’ of slurs.

While supporting a different view from Hare's one, Blackburn (1992) treats slurs and thick terms along similar lines. He endorses a deflationary view according to which there is no such thing as a term encoding description and evaluation at the same time. According to Blackburn, the attitudes associated with slurs and thick terms are typically communicated via intonations, prosody and the like:

“The picture I want to paint, then, is of a multiplicity of attitudes and feelings, and a flexible and changing repertoire of linguistic expression, with feeling more naturally signaled by intonation, and only unreliably read back from vocabulary except in very few cases” (Blackburn 1992: 296)

However, he seems to acknowledge that some negative content might be in fact lexicalized for slurs, whereas this is not the case for most thick terms: “The dictionary puts no ‘positive’ indicator of attitude by any of Hume’s terms [i.e. positive thick terms], in the way that it puts ‘derog.’ or ‘usually contempt.’ by certain epithets of abuse.” (Blackburn 1992: 286).

Gibbard (2003a) suggests that slurs might match the definition of ‘thick term’ provided by Williams, as they mix classification (i.e. description) and attitude (i.e. evaluation). However, slurs (understood as thick terms *à la* Williams) would work differently from how prototypical thick terms (like ‘chaste’ and ‘cruel’) work according to Gibbard. For Gibbard, when one does not endorse the attitude conveyed by a thick term like ‘chaste’ and ‘cruel’, one cannot regard an utterance featuring such a term as true or false. On the other hand, according to Gibbard, one can (more or less) evaluate the truth value of a slurring utterance:

“Those who lack the attitudes that would make them users [of thick terms, as defined by Williams], I am saying, cannot join users in regarding those thick claims as true (...). Racial epithets may sometimes work this way: Where the local population stems from different far parts of the world, classification by ancestry can be factual and descriptive, but, alas, the terms people use for this are often denigrating. Nonracists can recognize things people say as truths objectionably couched.” (Gibbard 2003a: 300)

It is worth noting that Gibbard (1992, 2003a, 2003b) interestingly analyzes thick terms as (i) having both descriptive and evaluative content and (ii) triggering a presupposition. Nevertheless, his presuppositional account – deeply inspiring as it is – is not fully developed and detailed. Gibbard appeals to the notion of presupposition, but he is not interested in deploying the set of tools provided by the theory of presupposition (in terms of projective behavior, common ground, accommodation, failure, rejection, complicity, etc.). In a sense, the way in which Gibbard means ‘presupposition’ is very broad: “The ethos of a community offers presuppositions on which thick meanings depend” (Gibbard 2003b: 177). Gibbard’s main goal is not to provide a linguistic analysis of thick terms: “In this sense, the term ‘zowy’ [an invented thick term] carries (...) a presupposition. Making this loose sense precise, though, would be a tricky matter, to say the least, and I won’t attempt to do so” (Gibbard 2003b: 146). Rather, he investigates broader questions, such as whether there are normative facts, he draws crucial distinctions between properties, concepts, states of affairs, natural and non-natural facts; whether normative facts, beliefs and truths are pseudo- or genuine and so on. These issues are beyond the scope of the present work. On the other hand, my aim is to develop a more refined analysis of the semantics and the pragmatics of slurs and thick terms, detailed enough to generate predictions and to allow a systematic comparison between these two classes of expressions. My interest, as Väyrynen (2011) puts it, concerns the ‘location of evaluation’, i.e. the question as to how to characterize the relation between description and evaluation, rather than issues such as the relation between moral terms (thin and thick) and reasons to action, the debate between internalism and externalism on reasons and the like (see Gibbard 2003a).

More recently, Elstein and Hurka (2011) have analyzed thick terms and slurs along the same lines. They call an epithet like ‘kraut’, “fully thick”, as opposed to terms such as ‘selfish’, that they analyze as in a middle position between thin terms and slurs (Elstein and Hurka 2011: 524). In this respect, their account resembles Gibbard’s one, as they analyze slurs as the items that most precisely fulfill the definition of ‘thick term’ provided by Williams (1985).

Finally, Väyrynen (2013: 150-156) discusses the idea of treating pejoratives as objectionable thick terms. Väyrynen (2016a) leaves room for the possibility that his deflationary account of thick terms could be applied to slurs. In particular, he underlines some analogies between his account of thick terms and Bolinger (2015)'s deflationary account of slurs. We shall come back to this point in chapter 10.

When the debate on slurs and pejoratives started flourishing, scholars working on derogatory epithets have started considering the comparison between slurs and thick terms. Richard (2008) underlines that slurs and thick terms raise the same kind of problems: Questions as to how the evaluation is related to a certain lexical item, whether speakers change the meaning of an evaluative term when they succeed in using it *without* its typical evaluative content and the like. Williamson (2009) is agnostic on whether his analysis of slurs in terms of conventional implicature could apply to evaluative moral terms: "The bearing on such terms of the present paper is a large question". Hom (2010) defines his own account of slurs as "thick", because he postulates a semantics of pejoratives that includes descriptive and evaluative content (see chapter 7); nevertheless, he does not seem to suggest a uniform (truth-conditional, in this case) analysis for slurs and thick terms. Mišćević (2011), in commenting on Richard (2008), endorses the idea that slurs and thick terms could get a uniform treatment and argues *contra* Richard in favor of the truth-aptness of utterances involving these evaluative terms (we will come back to this point in section 3.3).

All these suggestions constitute crucial and inspiring hints, but none of them was satisfactorily developed: A fully-fledged account of how to analyze thick terms and slurs along the same lines is needed.

2. The projective behavior of slurs and thick terms

In this chapter I will focus on the presuppositional behavior of hybrid evaluatives. In section 2.1 I discuss projection and present an alternative classification of projective content put forward by Tonhauser et al. (2013). In section 2.2 I discuss rejection, that is, the strategies employed to reject the descriptive and the evaluative content respectively. I will point out that the strategies of rejection reveal some interesting differences between descriptive and evaluative presuppositions. Finally, in section

2.3 I will discuss an alternative way to explain the projection of the evaluative content that does not rely on presupposition.

2.1 Projection

A well-known feature of both thick terms and slurs is projection (Väyrynen 2009, 2013, Eklund 2011, 2013 for thick terms, Croom 2011, Camp 2013, Jeshion 2013a, 2013b for slurs). Consider the following examples (Väyrynen 2013: 64, 70, 78):

7. Madonna's show is lewd.
8. Madonna's show isn't lewd.
9. Is Madonna's show lewd?
10. If Madonna show is lewd, tabloid press will go nuts.
11. Madonna's show might be lewd.
- π_7 . Things that are sexually explicit beyond conventional boundaries are bad because of being so.

All the above occurrences of the term 'lewd' seem to involve a negative evaluative content along the lines of π_7 , even when the predicate is embedded under negation, questions, conditionals or modals. The same happens with slurs:

12. Madonna is a wop.
13. Madonna is not a wop.
14. Is Madonna a wop?
15. If Madonna is a wop, I don't want to go to her concerts.

16. Madonna might be a wop.

π_{12} . Italians are bad because of being Italian.

For both thick terms and slurs, we observe that the presupposed content triggered by hybrid evaluatives scopes out of semantic embeddings. One reason to test these intuitions by employing an objectionable evaluative is that it is hard to perceive whether a certain utterance presupposes a certain content, when such a content is something that speakers and addressees already take for granted (see Hare 1963). However, because in principle *any* thick term could be seen as objectionable, projection data concern thick terms in general (see Väyrynen 2013: 56).

Projection is typically taken to be the distinctive feature of presuppositions. However, recent work in linguistics emphasized the heterogeneity of projective content. In particular, Tonhauser et al. (2013), while leaving aside notions such as ‘presupposition’, ‘conventional implicature’ etc., individuate four classes of expressions and constructions that give rise to projective contents. The four classes of triggers are distinguished by two properties, (i) “strong contextual felicity” and (ii) “obligatory local effect”.

	Projective content	Strong contextual felicity	Obligatory Local effect
Class A	yes	yes	yes
Class B	yes	no	no
Class C	yes	no	yes
Class D	yes	yes	no

The strong contextual felicity requirement demands that the trigger at stake imposes a strong constraint on the context, i.e. that the occurrence of the trigger requires that a certain content is entailed by the context in order for the utterance to make sense. Take a context where Lucy is on the bus eating a cucumber salad and a woman that she does not know starts talking to her and says:

17. Your salad looks nice. I never make cucumber salad at home cause my daughter hates cucumbers.

Even though Lucy did not know that the woman had a daughter, the utterance is acceptable: The projective content triggered by ‘my daughter’ does *not* have a strong contextual felicity requirement. Now imagine that the woman goes on and says:

18. Our bus driver is eating empadanas, too. (Tonhauser et al. 2013: 78)

The utterance is not felicitous unless a salient person in the context – different from the bus driver – is eating empanadas, which – as far as Lucy knows – is not the case. We conclude that the projective content triggered by ‘too’ has a strong contextual felicity requirement.

Slurs and thick terms pattern like the possessive ‘my daughter’ and unlike ‘too’ with respect to the strong contextual felicity requirement. As a matter of fact, speakers who do *not* endorse the presupposed evaluation of certain expressions can nevertheless make sense of what the interlocutor is saying. Suppose that Lucy and the woman on the bus are talking about a movie. The woman on the bus utters something like (19) and (20):

19. That movie is lewd!

20. The director is a wop.

Lucy would have no problem in making sense of (19) and (20), even if she does not share at all what the woman is presupposing about Italian people and sexual explicitness. In sections 2.2 and 3.3 we will come back to this point and discuss what happens when speakers do not share the presuppositions that the interlocutor is taking for granted.

The second parameter with respect to which Tonhauser et al. (2013) classify projective content is obligatory local effect, that is, they test whether the projecting

content needs to be part of the local context of a belief-predicate (i.e. the attitude holder's epistemic state) or not. For instance, suppose that Mary says:

21. Jane believes that Bill has stopped smoking.

The projective content – namely that Bill used to smoke – *has to* be ascribed to Jane (and it *can* be ascribed to Mary, but that is not obligatory). We conclude that the projective content triggered by 'stop' shows obligatory local effect. Now suppose that Mary says:

22. Jane believes that Bill, who is Sue's cousin, is Sue's brother.
(Tonhauser et al. 2013: 92)

The projective content of the appositive – namely that Bill is Sue's cousin – is not ascribed to Jane, but only to Mary: Appositives do *not* show obligatory local effect. Slurs and thick terms do not show obligatory local effect, either. Suppose that Mary says:

23. Jane told me that if the movie was lewd, it would have been more popular.⁸

24. Jane told me that the director is a wop.

In (23) and (24), the projective contents – namely that Italian people are bad and that things that are sexually explicit beyond conventional boundaries are bad – do not have to be ascribed to Jane. Mary's utterances are *compatible* with Jane lacking any negative attitude towards Italians and sexually explicit things (attitudes that would be ascribed only to Mary).

In conclusion, we observe that the projective content of slurs and thick terms does not have a strong contextual felicity requirement, nor obligatory local effect. In Tonhauser et al. (2013)'s classification, both slurs and thick terms would belong to

⁸ I owe this example to Isidora Stojanovic (p.c.)

‘class B’, the class of those expressions that (i) give rise to projective content; (ii) do *not* have a strong contextual felicity requirement; (iii) do *not* have an obligatory local effect. It is noteworthy that this category includes items that have been analyzed in the literature in terms of conventional implicatures (for instance, appositives) and expressions that have been analyzed in terms of presuppositions (for instance, personal pronouns and possessive noun phrases).

2.2 Rejection

In addition to the projective behavior described above, speakers cannot reject the evaluative content of hybrid evaluatives through mere denial (Väyrynen 2009, 2013 for thick terms, McCready 2010, Croom 2011, Camp 2013, Jeshion 2013a, 2013b for slurs). Observe:

25.

A: Madonna’s show was lewd.

B1: No, it wasn’t lewd.

B2: Well, the show was indeed sexually explicit beyond conventional boundaries, but there’s nothing bad in this.

B3: Madonna’s show wasn’t lewd, because there is no such thing.

In the above exchange, the disagreement concerns the application of the predicate ‘lewd’ to a certain object, namely, Madonna’s show: B1 acknowledges that there are things that are sexually explicit beyond conventional boundaries and bad because of that; however, according to B1, Madonna’s show is not one of those things. B2 and B3, on the other hand, reject the idea itself that being sexually explicit beyond conventional boundaries can be bad in itself: B2 retrieves, articulates and rejects the presupposition triggered by A; B3 uses a metalinguistic negation. A similar case of metalinguistic rejection along the same lines of B3 comes from the trials of Oscar Wilde (Foldy 1997: 8):

26.

Carson: Do you consider that⁹ blasphemous?

Oscar Wilde: “(...) the word ‘blasphemous’ is not a word of mine”.

The same happens for slurs:

27.

A: Madonna is a wop.

B: No, she isn’t.

B’s answer is not apt to target the evaluation and it is at least ambiguous between denying the descriptive content and metalinguistically rejecting the evaluative content.¹⁰ The issue of how to reject slurs in conversation is tricky and, unlike the case of thick terms, many social issues are at stake at the same time. Observe some examples of rejection of slurs from the COCA:

28.

A. I understand Black culture. I grew up around black people all my life. If the truth be told I probably know niggers better than you, Monsieur Delacroix.

B. I would prefer you not use that word in my presence.

29.

A. You think he’s a fag.

B. You think he’s gay.

⁹ ‘That’ refers to *The Priest and the Acolyt*, a story from John Francis Bloxam.

¹⁰ Note that the metalinguistic reading is easier to access with some kind of follow up or with a marked intonation.

30.

A. I heard he was a fag (...)

B. Historians don't generally put it quite that way, but yes, da Vinci was a homosexual.

In (28)-(30) speakers forbid or correct the use of a slurring term, but they do not articulate the presupposition properly. But we can find explicit articulations of the derogatory meaning of slurs in pedagogical contexts, where experts teach how to respond to slurs (homophobic, in the case at stake):

31. What you just said was really inappropriate because you are implying that there is something wrong with being gay or lesbian when there isn't.¹¹

We observe that speakers tend not to articulate the presupposed content and we might wonder why. A first point to outline is that there is a crucial difference between rejecting descriptive and evaluative presuppositions. Consider (32) and (34) and their presuppositions π_{32} and π_{34} . From now on, we will use '?' before an utterance to signal that the utterance is odd or infelicitous; '??' signals that the utterance is strongly infelicitous; '*' signals that the utterance is redundant.

32. Jane stopped smoking.

π_{32} . Jane used to smoke.

33. Hey wait a minute, I didn't know Jane used to smoke!

34. You finally realize you were dating a loser.

π_{34} . You were dating a loser.

¹¹ <http://www.healthiersf.org/LGBTQ/InTheClassroom/docs/Responding%20to%20Homophobia.pdf>.

35. ? Hey wait a minute, I didn't know that I was dating a loser!

36. Hey you shouldn't talk like that about my ex.

Compare (32) and (34) and the presuppositions that they trigger: The presupposition conveyed by (32) only involves a state of affairs (someone having used to smoke), whereas the one conveyed by (34) involves an evaluation about someone (someone being a 'loser'). Now suppose that for each utterance, there is an interlocutor who does not take the presupposed content for granted and who is not willing to accommodate it immediately (see Lewis 1979, Stalnaker 2002, von Fintel 2008). In response to (32), a reply along the lines of 'Hey wait a minute, I didn't know that...',¹² as in (33), will be appropriate. Nevertheless, (35) shows that such answers are not appropriate anymore when an evaluation is at stake. If the speaker presupposes that someone is 'a loser' and the interlocutor does not take that for granted, she cannot reply with something along the lines of 'Hey, wait a minute, I didn't know that...'.¹³ Something like (36) would be more appropriate.

The general point that I want to make is that when an interlocutor wants to reject a presupposed evaluation, the issue is not whether the interlocutor knew already the presupposed content or not, but rather whether she is willing to accept it or not. In general, when the presupposed content involves an evaluation, articulating the presupposition on the model of 'Hey wait a minute, I didn't know that...' is not the best way to reject the presupposed content. As the above data suggest, the alternative strategies employed can include metalinguistic moves ((26), (28)) or corrections ((29)-(30)). The proper articulation and rejection of the evaluative presupposition is more likely to be found in pedagogical contexts, such as (31) and it does not involve 'Hey wait a minute; I didn't know that...' kind of answers.

In addition to this first consideration, there are at least three other factors to take

¹² A standard test for presuppositions introduced by von Fintel (2004) relies on the fact that speakers can felicitously assess the presupposed content by saying something like 'Hey, wait a minute. I had no idea that...'.
¹³ Unless what is at stake is the interlocutor's knowledge regarding the speaker's belief, in which case answers such as 'I didn't know you thought my ex is a loser' or 'I didn't know you had such a bad opinion about my ex' sound appropriate.

into account to explain why speakers tend not to articulate evaluative presuppositions:

(i) In general, to retrieve, explicitly articulate and reject a presupposition in a conversation is both cognitively and socially costly (see Chilton 2004, 64), so that simply rejecting the use of a term is easier for speakers rather than articulating the presupposition.

(ii) Speakers articulate the presupposed content only when they deem it necessary; for instance, they do it when the question under discussion is precisely why it is wrong to use a slur, as in pedagogical contexts.

(iii) Slurs are prohibited words, whose use can also be prosecuted.¹⁴ This might explain why, in the case of slurs, people prefer a metalinguistic or silentistic solution, in order not to get involved with those terms: They tend to forbid the use of the term (“I would prefer you not use that word in my presence”), rather than explicitly articulating the presupposed content; nevertheless, as (23) shows, speakers are able to provide the relevant articulation if required.

2.3 An alternative explanation of projection

In the previous sections of this chapter, I discussed the projective behavior of slurs and thick terms. The proposal that I put forward accounts for the projective behavior of slurs and thick terms by relying on well-known projective features of presuppositions. This means that the explanation of the projection of the evaluative content does not rely on the *nature* of the evaluative content itself, but on the *way* in which it is encoded, i.e. presuppositionally. In other words, the derogatory content of ‘wop’ and the negative evaluation of ‘lewd’ do not project because they are evaluative; they project because they are presupposed. This is a crucial difference between my account and the expressivist account of slurs developed by Jeshion (2013a, 2013b, 2014, 2016a, 2016b, forthcoming). According to the expressivist strategy, the derogatory content of slurs projects *because* it is expressive rather than

¹⁴ <http://fox2now.com/2015/10/06/missouri-company-faces-federal-lawsuit-over-racial-slur-allegations/>.

descriptive. Jeshion analyzes slurs as truth-conditionally equivalent to their neutral counterparts. In her framework, ‘wop’ has the same truth conditions of ‘Italians’: Utterances like ‘Leonardo da Vinci is a wop’ are strictly speaking true, utterances like ‘Leonardo da Vinci is a boche’ are strictly speaking false, etc.¹⁵ Note that Jeshion’s account is not committed to the existence of a neutral counterpart for every slur: All she needs to postulate is the existence of a group-defining property. As we shall see in chapter 6, this is a crucial point in my account and it becomes especially important if one wants to analyze slurs and thick terms along the same lines. For Jeshion, what distinguishes non-loaded expressions (neutral counterparts or more complex non-loaded paraphrases) from epithets is that the semantics of slurs also includes – at a *non*-truth-conditional level – an additional element, i.e. an expressive component.¹⁶ In particular, slurs encode *contempt*. What ‘wop’ means is something like ‘Italian’, pronounced with a contemptuous intonation, or accompanied by pejorative modifiers (such as ‘dirty’, ‘rotten’, etc). Contempt is to be understood as the attitude that one holds towards those that one regards as inferior *as persons*. In Mason’s words, contempt should be understood as “presenting its object as low in the sense of ranking low in worth as a person in virtue of falling short of some legitimate interpersonal ideal of the person, one the contemner endorses if not one that she herself succeeds in meeting” (Mason 2003: 241). In Jeshion’s analysis, contempt is “a highly structured affectively- and normatively-guided moral attitude” (Jeshion 2013a: 242). In this sense, contempt does not only express the speaker’s feelings – like other attitudes do (fear, anger, etc.) – but it also represents its target as *worthy* of contempt. Moreover, contempt does not necessarily give rise to clearly recognizable behavior. Jeshion (forthcoming) argues that empirical data support the claim “that when contempt is felt, the confrontational behaviors are less prevalent than the distancing and excluding behaviors”. This would mean that in order to express contempt, slurs-users do not need to display their attitude towards the target in particularly patent ways other than using a slur.

In addition to these two components (truth-conditional and expressive), the semantics of slurs also has a third component, namely the identifying one: A slur-

¹⁵ For Richard (2008, 2011), the fact that the meaning of slurs is expressive implies that slurring utterances are never truth-apt. For a ‘pure expressivist’ view on slurs, see Hedger (2012).

¹⁶ See also Tenchini and Frigerio (2016).

user aims at identifying the target by pinpointing “what the target is”. The identifying component is needed to explain how contempt works when encoded in slurs: In a way, the slur-user would not express contempt towards the target *qua* person, if the properties selected by the truth-conditional component were not taken to be fundamental (that is different from *essential*) to the identity of the target.

As we saw, Jeshion can explain the phenomenon of projection by relying on the features of expressive content. However, if the reason why the derogatory content projects is that – being expressive – it does not interact with semantic operators, this expressivist strategy implies that *any* occurrence of a slur must involve the expression of contempt. If that is the case, we rule out the possibility of suspending (or attributing to someone different from the speaker) the derogatory content of slurs in certain constructs and in certain circumstances (see Part III).

3. The semantics of hybrid evaluatives

In the previous chapters, I proposed a unified account of slurs and thick terms in terms of presupposition and I focused on the presuppositional patterns of these terms with respect to projection and rejection. The core of my proposal is that slurs and thick terms encode a certain descriptive content and at the same time trigger an evaluative presupposition. For example, ‘wop’ means ‘Italian’ but at the same time triggers the presupposition that Italians are bad because of being Italians; ‘lewd’ means roughly ‘sexually explicit beyond conventional boundaries’ and triggers the presupposition that things or individuals that are sexually explicit beyond conventional boundaries are bad because of being so; ‘generous’ means ‘disposed to give without expectation of receiving anything in return’ and triggers the presupposition that those who are disposed to give without expectation of compensation are good because of that. In the present chapter, I characterize in greater detail the evaluative and the descriptive content of hybrid evaluatives and I will try to explain the behavior of these expressions in terms of the features of their evaluative and descriptive content. I shall outline a semantics for the basic evaluative terms, i.e. ‘good’ and ‘bad’ that occur in the evaluative content (3.1). We will see that such terms are highly context-sensitive and multidimensional. This under-determination gets inherited by the evaluative content of hybrid evaluatives and we

shall say something about the factors that contribute to the determination of the evaluation in the context; moreover, we shall observe the descriptive content of hybrid evaluatives and note that most adjectival thick terms, including ‘lewd’, ‘generous’, etc., are multidimensional; on the other hand, most slurs are nouns rather than adjectives and their descriptive content does not display the same amount of context sensitivity (3.2). In section 3.3 I will discuss the issues of reference and presuppositional failure and we shall see how such issues are related to controversial philosophical questions.

3.1 The evaluative content

When scholars discuss the attitude that slurs convey, they talk about ‘contempt’, ‘despise’, ‘negative moral evaluations’, etc. However, each definition is unsatisfactory on its own, as slurs do not involve a single attitude. Some uses of slurs are typically associated with contempt, but others seem to be associated with different attitudes, such as condescendence, disgust, dismissiveness, fear, hate, mockery, rage (Camp 2013: 338-339, Jeshion 2013a: 233). As Hom and May (2013: 297) notice, the same slur can be associated with a different attitude on different uses:

“This morally negative evaluation encompasses a wide range of negative potential attitudes from hate and derision on one extreme to mere dislike or resentment on the other. The ‘skinhead anti-semitic’ and the ‘country club anti-semitic’ express the same semantic content with their uses of ‘kike’, though their surrounding dispositions and psychologies are likely to vary tremendously”.

Because of this variability, Hom and May (forthcoming) avoid a precise characterization of the attitude conveyed by slurs:

“We have purposely left loose the notion of negative moral evaluation. If making it more precise is helpful, we are not adverse to thinking of it as contempt, in the sense explicated by Jeshion (2014), following Mason

(2003). On this view, contempt is an affective, person directed and focused moral notion based on perceived character defining traits. In the current context, however, nothing hinges on this, so we stick with our more generic characterization.”

In my proposal, the content presupposed by both slurs and thick terms is underdetermined. The evaluative content of hybrid evaluatives crucially features thin terms such as ‘good’ and ‘bad’ whose well-known context sensitivity can account for the under-determination and context dependence that are at stake here. ‘Good’ and ‘bad’ are all-purpose evaluatives: We apply them to all sorts of things, from good/bad weather, cars, meals, results, to good/bad books, movies, actions, souls etc. and each time there are important differences among the evaluations that are conveyed in those different cases. This means that, for one thing, ‘good’ and ‘bad’ need to be disambiguated into ‘good’/ ‘bad’ *as* something (good as a book, bad as an action, etc.).¹⁷ Moreover, ‘good’ and ‘bad’ are adjectives that are *gradable* and *multidimensional* (for a characterization of evaluative adjectives, see McNally and Stojanovic 2017).

3.1.1 Gradability

Gradable adjectives (like ‘tall’, ‘expensive’, ‘heavy’ etc.) refer to properties that come in degrees: Something can be *very* tall, or *more* expensive than something else, or *slightly* heavy. Consider an utterance like (37):

37. That book is expensive.

In order to assign a truth value to an utterance like (37), one has to know what counts as ‘expensive’ in that context and what is *for a book* to be expensive.

¹⁷ Note that such a disambiguation is crucial, as underlined by Kyle (2013). If ‘good’ was not disambiguated as ‘good *as* something’, ‘good’ could truthfully apply to anything. Take the utterance ‘Gandhi’s murder is good’. Take a speaker S who would not apply the predicate ‘good’ to Gandhi’s murder and judges the utterance false. The interpretation of ‘good’ that the speaker actually has in mind when judging the utterance false is something like ‘good *as* an action’ or ‘good *as* an event’. If the context did not provide such a disambiguation, ‘Gandhi’s murder is good’ would always be true under at least *some* interpretation of ‘good’, such as ‘good *as* an example of unjustifiable violence’. As Kyle (2013) points out, it is crucial that context disambiguates the interpretation of ‘good’ and ‘bad’.

Gradable adjectives are functions that map their arguments onto abstract representations of measurement (i.e. degrees), which are points or intervals ordered along a dimension (e.g. height for ‘tall’, cost for ‘expensive’, weight for ‘heavy’ and so forth). In other words, gradable adjectives locate objects on scales that are relative to a *dimension*: ‘Tall’ locates things on the scale of tallness, ‘good’ on the scale of goodness, etc. (Kennedy and McNally 2005, Kennedy 2007).¹⁸ Moreover, there need to be a threshold on the scale. Consider ‘tall’. In order to define something ‘tall’, it is not enough that the object has *any* height, or ‘tall’ would be trivially true of any (at-least-) two-dimensional object. The threshold varies from context to context. Suppose that someone runs the 100 meters in 12 seconds. 12 seconds is a good time in an informal race among amateurs, but it is not a good time at the Olympiads: The threshold for ‘good *as* a time in the 100 meters’ varies between the amateur race and the Olympiads. Consider (38):

38. It’s good to run the 100 meters in 12 seconds.

In order to evaluate the truth value of a sentence such as (38), the context needs to supply a mapping from objects to degrees, a scale *S* and a threshold *t* on that scale. (38) is true iff the degree on the scale to which the function maps the object *o* (the running time) goes beyond the threshold. In a context *c* and world *w*, (38) is true iff $d(o) \geq t$, where $d(o)$ is the degree to which *o* is mapped on the scale of goodness. We can imagine that the degree to which running the 100 meters in 12 seconds is mapped goes beyond the threshold fixed in the case of the amateur race, but it does not go beyond the threshold fixed at the Olympiads.

Gradable adjectives, unlike non-gradable adjectives like ‘geological’, ‘married’ and ‘wooden’, take degree modifiers such as ‘very’ and ‘somewhat’:

39. Nancy is very tall.

¹⁸ Among the theories of gradable adjectives, one should distinguish between two kinds of approach: One according to which the context determines a comparison class with respect to which the threshold is fixed and one according to which the context directly fixes the threshold. I will stick to the latter for simplicity without exploring the differences between the two approaches, nor the arguments in favour of one over the other. For a discussion of the two approaches, see Klein (1980) and Glanzberg (2007).

40. The book is somewhat good.

41. ? The table is very wooden.

42. ? Luca is somewhat married.

Note that the parameters that need to be contextually fixed can vary from speaker to speaker. For this reason, gradable adjectives typically give rise to disagreement (*i.a.* see Glanzberg 2007, McNally and Stojanovic 2017).

3.1.2 Multidimensionality

Adjectives such as ‘tall’, ‘bald’, ‘salty’ are called ‘unidimensional’, because their interpretation depends on values along a singular scalar dimension (e.g., height for ‘tall’ and quantity of hair for bald, density of salt for ‘salty’, etc.). Some adjectives require more than a scale and are called ‘multidimensional’ (Bierwisch 1989, Sassoon 2013). Take ‘good’ and ‘bad’: A theatrical performance can be ‘good’ or ‘bad’ in terms of plot, scenic design, acting, dialogues, etc. and these criteria can differ from context to context (see Thomson 1992, 2008). Unlike unidimensional adjectives, multidimensional ones such as ‘identical’, ‘normal’, ‘happy’, ‘healthy’, etc. can felicitously take constructions such as ‘in every/some/most respect(s)’ (Sassoon 2013: 336):

43. The theatrical performance is good in most respects.

44. The acting was good in every respect.

45. ? The actor is bald in some respects.

The truth or falsity of (43)-(44) requires the determination of several parameters. While unidimensional adjectives only require the determination of a threshold on a scale, in order for a multidimensional adjective to determine such a threshold, it has

to be contextually determined also “the relative weight of each of the dimensions that contribute to the property in question” (McNally and Stojanovic 2017). How exactly these dimensions combine together is still largely under-explored, or, in Glanzberg (2007)’s terms, it is a “messy subject, about which we understand relatively little”. However, it constitutes a very interesting point, as the selection and weighted combination of the relevant dimensions is an additional source of context dependence and therefore an additional source of disagreement. Speakers could not only take different dimensions to be relevant, but also weigh them in different ways, thus producing diverging judgements about the same state of affairs (McNally and Stojanovic 2017 and Stojanovic 2016b).

So far we said that when one utters something like ‘Madonna is a wop’, one is asserting that Madonna is Italian and, in addition, one is presupposing something like: ‘Italians are bad for being Italian’. Now we can be more precise and say that in view of the characteristics of the all-purpose gradable and multidimensional adjectives ‘bad’ and ‘good’, the presupposed content of ‘Madonna is a wop’ uttered in a context *c* actually amounts to the following: ‘Italians are bad *in a way W* for being Italian’, where *W* is a way (or a combination of ways) in which something can be bad and, as we shall see in the next section, *W* is determined by the context and by the descriptive content of hybrid evaluatives.¹⁹

In a nutshell, the gradability and multidimensionality of thin terms result in the under-determination of the evaluative meaning of slurs and thick terms. Therefore, it is part of the predictions of my account that the evaluation is, in Hom and May’s words, left “loose”. One way to analyze ‘good’ and ‘bad’ is to claim that they are respectively associated with a pro-attitude and an anti-attitude (see section 3.3.4). Under the umbrella of ‘pro attitude’ we find admiration, approval, desire, liking, praise, pursuit, whereas ‘anti-attitude’ covers contempt, disgust, dismissiveness, fear, hate, mockery, rage, etc. This loose nature of the notions of ‘pro’ and ‘anti’ attitudes is consistent with the under-specification and multidimensionality of thin terms.

In these sections, we have discussed the under-determination and context sensitivity of the evaluative content of hybrid evaluatives stemming from the

¹⁹ Note that for the sake of brevity and simplicity, I will most of the times use just ‘good’/‘bad’, while meaning ‘good/bad in a way *W*’.

occurrence of expressions like ‘good’ and ‘bad’. Note however that the fact that the presupposition triggered by hybrid evaluatives features gradable and multidimensional adjectives like ‘good’ and ‘bad’, does not make *every* hybrid evaluative gradable and multidimensional. As a matter of fact, the evaluative content of such terms is *presupposed*. In order for hybrid evaluatives to display the typical behavior of gradable and multidimensional adjectives, the descriptive at-issue content (that is, their truth-conditional meaning) has to display these features. In section 3.2 we shall see how the descriptive content of most thick terms consists of gradable multidimensional adjectives, whereas slurs are typically nouns and, even when adjectives, their truth-conditional meaning is not characterized by gradability and multidimensionality. This can explain why, despite the gradability and multidimensionality of ‘bad’, slurs do not typically show the same context sensitivity of thick terms (and they probably give rise to less complex cases of disagreement and negotiation than thick terms, as we shall see in chapter 4).

3.1.3 How to interpret the evaluative content

The under-determination of the evaluative content of hybrid evaluatives raises interesting issues, such as *how* the evaluative content can be each time determined, how the relevant dimensions are picked in each context, and so on. Glanzberg (2007) points out that while containing contextual parameters, the semantics of gradable (and multidimensional) terms does not say *how* the parameters get contextually specified: This is rather a matter of *metasemantics*. We can name some factors that can contribute to the determination of the evaluative content. For one thing, the interpretation of the evaluation can be driven by the descriptive content. As argued in Cepollaro and Stojanovic (2016: 474), the truth-conditional meaning of thick terms can constrain the interpretation of the triggered evaluation. For example, the descriptive content of ‘generous’, namely ‘disposed to give without expectation of receiving anything in return’ calls into play a *moral* dimension. Moreover, the kind of constraint provided by the truth-conditional meaning can be even more specific than this, as the moral dimension is quite broad: The descriptive contents of terms like ‘lewd’, ‘chaste’, ‘unchaste’, ‘obscene’ call upon the moral dimension in relation to sexuality; on the other hand, terms like ‘generous’ and ‘selfish’ have to do with the

moral dimension with respect to the management and dispensation of goods and so on. As we said, the occurrence of ‘good’ and ‘bad’ in the evaluative presuppositions should be read as ‘good/bad *in a way W*’, where *W* is determined by the context and the descriptive content of the hybrid evaluative.

Consider the following utterance:

46. She is courageous, but she is selfish.

(46) is felicitous: This means that it is not contradictory to deem ‘selfish’ (and therefore bad) someone who is at the same time taken to be ‘generous’ (and therefore good). This supports the claim that the evaluations conveyed by ‘generous’ and ‘selfish’ do not involve something like ‘morally good/bad’ but something like ‘good/bad *in a way W*’.

In the case of slurs, it is less clear how the descriptive content could narrow down the available interpretation of the presupposed evaluation. What kind of dimension would ‘being Italian’ or ‘being Jewish’ call upon? The descriptive content of most slurs refer to monadic properties: no scales, respects, dimensions, degrees, etc. This kind of descriptive content does not narrow down very much the meaning of the all-purpose evaluative ‘bad’ that features in the associated presupposition. A plausible hypothesis put forward in Cepollaro and Stojanovic (2016) would be to take the evaluative content of slurs to be fairly unconstrained, something like ‘bad in all (or most) respects because of being (...)’. However, I do not take the evaluation conveyed by slurs to be a *tout court* evaluation, that is, something that diminish people under *every possible respect*. This (quite radical) hypothesis is discarded by the felicity of the following utterances:

47. That wop is a good philosopher.

48. He is a beautiful wop.

I suggest that in the case of slurs it is the stereotype associated with the target class that plays a crucial role in determining the evaluative content. Suppose that a speaker is exposed to an environment where the stereotype associated with German

people involves something about inelegance and lack of taste; the dimension that will be relevant when the speaker is confronted with a germano-phobic slur will be aesthetic; on the other hand, if a speaker is exposed to a stereotype according to which German people are cruel, the dimension that will be salient is going to be moral. One might wonder whether the term ‘boche’ used by these two hypothetical speaker is in fact one and the same. Or, to make the case clearer, suppose that the Italian slur for ‘German’, namely ‘*crucco*’, is associated with a stereotype involving inelegance, whereas the English slur for ‘German’, ‘boche’, is associated with a stereotype involving cruelty (see Dummett 1973). One might claim that the two slurs are not in fact synonymous. My theory, on the contrary, is committed to the claim that two such words are synonymous as long as they share the same descriptive and evaluative content (which is, as we just saw, quite loose). In the framework that I propose the stereotype is *not* linguistically encoded (see Richard 2011 and Jeshion 2013b). The contrast between (49) and (50) speaks in favor of the hypothesis that while the descriptive content is linguistically encoded, the stereotype is not: (49), where a certain stereotype associated with ‘wop’ is negated, is felicitous, whereas (50), where the descriptive content of ‘wop’ is negated, is infelicitous.

49. That wop doesn’t know how to cook.

50. ? That wop is not Italian.

3.1.4 Inter-variation and intra-variation

One might wonder whether the under-determination just described can be responsible for the *inter-variation* of slurs (see Hom 2008, Jeshion 2016a and Popa-Wyatt 2016), i.e. the fact that the degree to which slurs are perceived as offensive varies from item to item: For instance, ‘kike’ is typically taken to be more offensive than ‘kraut’. My answer is no: The fact that some slurs are perceived as deeply offensive (like ‘nigger’) and others as barely offensive (like ‘tree-hugger’) does not mean that they must *encode* a different evaluation. I take the inter-epithets variation to depend on extra-linguistic features, such as the historical fact that certain people and not others were victims of violence and discrimination. In a counter-factual

situation where English people were subjected to discrimination and violence for centuries, ‘limey’ would probably sound like a very derogatory term.²⁰ I do not take the difference in perceived offensiveness as evaluative encoded information. As a matter of fact the same point can be made of non-slurring terms. It is more difficult to use ‘man’ as an insult than it is to use ‘woman’, but this does not tell us anything about what the two predicates *mean*. The fact that one is perceived as a worse insult than the other entirely depends on extra-linguistic matters.²¹

Another question is whether specific lexical items can get typically associated with a specific attitude. Several factors can contribute to this regular association, such as the agents who typically use the term and the attitude that they typically convey. For example, in Italy some people (especially elder ones) use ‘*cioccolato*’ (literally: chocolate praline) to call black people, typically children, with an attitude of condescendence. The use of ‘*cioccolato*’ (independently of *who* uses it) is not typically associated with rage, hatred or threat, but with condescendence. A plausible explanation of why the term is associated with condescendence rather than with hatred is that it is usually used by speakers who accompany it with such an attitude (see Bolinger’s and Nunberg’s account of affiliation and co-occurrence expectations in chapter 9).²² Another factor is the description that the term evokes: Slurs against Jewish people like ‘firewood’ or ‘lamp shade’ evoke the Nazi genocide and are probably associated with threat rather than fear or condescendence and might be perceived as more offensive than a non-evoking term like ‘kike’; similarly, ‘gorilla’ used as a term targeting black people is associated with mockery in virtue of the description that it provides. Different descriptions are connected to different stereotypical properties that have different degrees of offensiveness.

In a nutshell, we observe that despite the fact that the evaluative content of slurs is typically under-determined, sometimes particular attitudes can get attached to a

²⁰ Hom (2008: 426, 433) makes a similar point and talks about “derogatory variation”. See also Cepollaro (2015) and Čupković (2015).

²¹ About the perception of derogation, see O’Dea C.J. et al. (2014) and Saucier (2014).

²² Another similar case again from Italian is the term ‘*negretto*’ for black children, conveying an attitude of condescendence. The term is morphologically complex: It takes ‘*negro*’, the Italian equivalent of ‘nigger’, and combines it with the suffix that typically signals affection, namely ‘-etto’. The use of ‘*negretto*’ (independently of who uses it) is not typically associated with rage, hatred or threat, but with condescendence.

specific lexical item in virtues of different mechanisms (the attitudes usually involved when the term is used, the specific description provided by the term, etc.). Note that this regularity can explain the so-called intra-variability (see Jeshion 2016a), that is, the fact that different slurs targeting the same group can convey a different attitude and different degrees of offensiveness (for example, ‘coon’, ‘negro’, ‘nigger’, ‘spade’ are usually taken to convey different attitudes in different degrees).

In conclusion, the presupposed evaluative content of hybrid evaluatives is context-sensitive because of the occurrence of the thin terms ‘good’ and ‘bad’, which are gradable and multidimensional. The evaluative content is therefore under-determined and can vary from context to context. However, the presupposed content being gradable and multidimensional does not result in the term itself displaying gradable and multidimensional behavior. We shall see in the next section that the descriptive content plays a crucial role in making the hybrid evaluative gradable and multidimensional (or not).

3.2 The descriptive content

The descriptive content of slurs displays different features from that of thick terms. As we said, most slurs are nouns rather than adjectives and slurs are typically about people rather than things, actions and behavior. Moreover, while the descriptive content of thick terms is typically gradable and multidimensional, the descriptive content of slurs is not. In chapter 6, I shall argue that these differences are only superficial contrasts and do not challenge the claim that slurs and thick terms rely on the same linguistic mechanism of encoded evaluation.

Let’s now consider the descriptive content of hybrid evaluatives in greater detail. As we said, the descriptive content of most thick terms consists of gradable expressions, making them gradable adjectives, too. When thick terms can felicitously take degree modifiers, also their descriptive counterparts can:

51. Nancy is very generous.

52. Nancy is very disposed to give without expectation of receiving anything in return.

53. Nancy is somewhat lewd.

54. Nancy is somewhat sexually explicit beyond conventional boundaries.

Moreover, when thick terms locate objects on a scale, their non-loaded counterparts do, too:²³

55. Their last movie is lewder than the previous one.

56. Their last movie is more sexually explicit beyond conventional boundaries than the previous one.

57. Nancy is more generous than Lucy.

58. Nancy is more disposed to give without expectation of receiving anything in return than Lucy.

While it is no wonder that slurs – *qua* nouns – are not gradable, we observe that slurring adjectives, in their literal uses, do not seem to be gradable either. They do not normally take degree modifiers such as ‘very’ and ‘somewhat’ or admit comparative constructs, nor do their neutral counterparts:

²³ This is in a way reminiscent of Kaplan (1999)’s contrast between the intuitively valid argument A and the intuitively invalid argument B:

A. Premise A: That damn Kaplan was promoted. Conclusion A: Kaplan was promoted.

B. Premise B: Kaplan was promoted. Conclusion B: That damn Kaplan was promoted.

Kaplan’s point is that one can safely infer from an evaluatively charged statement (like Premise A) a descriptive statement (like Conclusion A), but not the other way around. The reason is that in the evaluative utterance “there is more semantic information” than in the descriptive one.

For a development of Kaplan’s proposal in terms of “use-conditional meaning”, see Gutzmann (2015). For an analysis of slurs in terms of inferential role, see Tirrell (1999).

59. ? John is somewhat queer.²⁴

60. ? John is somewhat homosexual.²⁵

61. ? John is very queer.

62. ? John is very homosexual.

63. ? John is queerer than Luca.

64. ? John is more homosexual than Luca.

65. ? Bianca is more of a wop than Luca.

66. ? Bianca is more of an Italian than Luca.

Moreover, most thick terms are multidimensional, as their descriptive content is multidimensional. As a matter of fact, when thick terms can be uttered felicitously with ‘in every/some/most respect(s)’ constructions, also their descriptive counterparts can:

67. The movie is lewd in some respects.

68. The movie is sexually explicit beyond conventional boundaries in some respects.

²⁴ I can access two readings of (59) under which the utterance is felicitous. They would be (i) a reading where ‘queer’ is a synonym for ‘weird’, ‘strange’, not a homophobic slur and (ii) a reading where ‘queer’ is a slur, but it is used in a sort of figurative way, used to talk about the stereotypical properties and applicable to a subject that does not belong to the target class. As a matter of fact, one can force a scale on non-gradable predicates in a figurative way. Take for instance utterances such as ‘My bag is more hand-made than the one you bought’ or ‘Jane is more of a New Yorker than Peter’. These figurative readings, however, do not do any harm to the argument. Thanks to Pier Marco Bertinetto for drawing my attention to these cases.

²⁵ Again, there is reading where ‘homosexual’ is used in a sort of figurative way to talk about the stereotypical properties. This does not show that ‘homosexual’ is gradable. As a matter of fact we can force non-gradable adjectives into figurative uses: ‘this is more hand-made than that’, ‘Sue is more an engineer than Peter’, etc.

69. Nancy is generous in every respect.

70. Nancy is disposed to give without expectation of receiving anything in return in every respect.²⁶

Unlike most thick terms, the descriptive content of slurs is typically *unidimensional* and slurs inherit this unidimensionality, as we can see from the infelicity of the following constructions, that include both nominal and adjectival slurs:

71. ? Bianca is a wop in some respects.

72. ? Bianca is Italian in some respects.

73. ? Bianca is queer in most respects.

74. ? Bianca is homosexual in most respects.

I acknowledge that the description that I sketched of how thick terms work is quite complex. However, this is not surprising: It is required the amount of contextual contribution involved for *any* multidimensional adjective, plus the determination of ‘good’ or ‘bad’ featuring in the evaluative content. The case of most slurs looks simpler, since, as we saw, the truth-conditional content of slurs is typically non-gradable.

In a nutshell, the evaluation triggered by both slurs and thick terms carries a certain amount of context sensitivity stemming from the occurrence of terms like ‘good’ and ‘bad’. Moreover, in the case of most thick terms, the descriptive content is a further source of context sensitivity, as it features expressions which are gradable

²⁶ Note that when the descriptive content is very complex, it can be a bit unclear what is the scope of expressions such as ‘in every respect’. A better formulation might be: ‘Nancy is in every respect disposed to give without expectation of receiving anything in return’.

and multidimensional. This additional source of context sensitivity constitutes a further source of disagreement and negotiation, as we shall see in chapter 4.

3.3 When things go wrong: presuppositional failure

In this section, I will address the issues of reference and extension (3.3.1). I will spell out what is exactly presupposed by hybrid evaluatives (3.3.2) and discuss how to interpret the presuppositional failure in the case of evaluative presuppositions (3.3.3). We shall see that these topics constitute a very complicate matter in the analysis of hybrid evaluatives, as they relate to highly debated issues about values in ethics (3.3.4). My aim is not to establish which theory of values is the best one, but to clarify the relations between the linguistic analysis and the issues at stake in ethics.

3.3.1 Reference and extension

One of the most debated questions for slurs and thick terms concerns reference (Richard 2008, Hom and May 2013 for slurs; Eklund 2011 for thick terms): What do they refer to, if they refer *at all*? Consider slurs. Take the following utterance:

75. Bianca is a wop.

Authors like Christopher Hom and Robert May (Hom and May 2013, 2014, forthcoming; see chapter 7) take slurs to be terms with a necessarily empty extension, as they roughly mean something along the lines of ‘worthy of negative moral evaluation for being G’, where ‘G’ refers to a discriminated group (we have already seen in section 1.1 that for Hom and May to belong to a group cannot under any circumstance ground the negative evaluation conveyed by slurs). So both (75) and (76) are false:

76. Bianca is a boche.

A problematic prediction of this approach is that ‘Bianca is not a wop’ turns out to be necessarily true, whether the subject is Italian or not.

Other scholars (*i.a.* Potts 2005, 2007c, Whiting 2013, Jeshion 2013a, 2013b) claimed that slurs refer to the same objects as their neutral counterparts. According to this approach, if the subject is Italian, (75) is true and (76) is false. This view has the controversial consequence that (75) turns out to be true, as well as utterances like ‘Italians are wops’. However, the defenders of this approach can advocate for a distinction between a *technical* sense of ‘true’, that only involves truth conditions, and a *folk* sense of ‘true’ that takes into consideration pragmatic factors of various kind. In this spirit, ‘true’ should not be confused with ‘assertable’.

The presuppositional analysis offers a fairly direct answer to questions about reference: An utterance in which a hybrid evaluative occurs (like (75) or (76)) has the same truth conditions as those of the utterance in which the hybrid evaluative is substituted by its neutral counterpart, but in addition it also has failure conditions. Because they feature presupposition triggers, they are appropriate²⁷ only if the presupposed evaluation is correct. So (75) and (76) have the same truth conditions as (77) and (78) if the evaluation triggered by ‘wop’ and ‘boche’ respectively is correct:

77. Bianca is Italian.

78. Bianca is German.

Note moreover that the gradability of thick terms makes the issue of reference more complicated, as in order to establish whether the evaluation conveyed is correct, scales and thresholds need to be established first. The reference of thick terms is in a way a complicate issue, as we shall see in section 4.2.

3.3.2 What is exactly presupposed

Note that in order to address the issues of reference and extension, it is crucial to establish what is exactly presupposed by hybrid evaluatives. In the literature, two options were put forward, that we can call the ‘objective’ and ‘subjective’ option (see

²⁷ Sometimes, I will talk about truth conditions and appropriateness when speaking about a hybrid evaluative term. Obviously, it is not the predicate that is appropriate or that has truth conditions: It is a shortcut to talk about the sentence in which the hybrid evaluative occurs and to which it contributes.

Predelli 2010 and Cepollaro 2015). For ‘wop’ and ‘lewd’, the objective and subjective options would be respectively:

W_o. Italians are bad for being Italian.

W_s. The speaker believes that Italians are bad for being Italian.

L_o. People and things that are sexually explicit beyond conventional boundaries are bad for being so.

L_s. The speaker believes that people and things that are sexually explicit beyond conventional boundaries are bad for being so.

The ‘objective’ option is about the target class (‘Italians are bad because of being Italian’) whereas the ‘subjective’ option is about the speaker (‘*the speaker believes that* Italians are bad because of being Italian’). The OO-SO opposition resembles the debate about presupposition triggers under the scope of *attitude* verbs (Beaver and Geurts 2014):

79. Mary wants to phone her brother.

B_o. Mary has a brother.

B_s. Mary believes she has a brother.

Some scholars hold that (79) presupposes (B_o), some that it presupposes (B_s), some others that it presupposes both (see *i.a.* Karttunen 1974, Heim 1992, Beaver and Geurts 2014). The dilemma here is quite similar: Either the presupposition concerns a state of affairs (Mary having a brother, certain people or things being such and such), either it regards the speaker’s beliefs (Mary’s belief that she has a brother, the speaker’s belief about certain people or things). Beaver and Geurts (2014) suggest that it might be possible to infer the ‘subjective’ option from the ‘objective’

one and vice versa: In the former case, the addressee would infer from the speaker’s utterance about a state of affairs that the speaker *believes* it to be correct; in the latter, the addressee would infer from the speaker’s expression her beliefs, something about a state of affairs through a process that has some commonalities with testimony (see for instance Fricker 1987). Such a possibility of mutually infer one option from the other would be why it is hard to tell one option from the other; they propose a visual recap in a table that, applied to the slur ‘wop’, would be the following:

	The speaker believes that Italians are bad for being Italian	Italians are bad for being Italian
Objective Option	Inference	Presupposition
Subjective Option	Presupposition	Inference

However, choosing one option over the other has important consequences concerning a few issues, including the account of failure. In section 4.1 we shall see some reasons to favor the objective option.

3.3.3 Failure

In the previous sections, I claimed that when a hybrid evaluative occurs in an utterance, the utterance has the same truth conditions that it would have had if it featured a corresponding non-loaded term instead of the hybrid evaluative, but, *qua* presupposing utterance, it is appropriate only if the presupposed evaluation is correct. Let us now consider what happens when the evaluation is *not* in fact correct, that is, let us consider failure. According to which option one favors, subjective or objective, different accounts of failure follow.

Consider the subjective option proposed by Schenker (2007)²⁸:

“with respect to a context (c) and a world (w). [...]

[[honky]] (c)(w) $\neq \#$ iff the agent of c believes in the world of c that white people are despicable. If $\neq \#$, [[honky]] (c)(w) = [[white]] (c)(w)”
(Schlenker 2007: 237-238)

The subjective presupposition holds whenever the speaker truly endorses the relevant attitudes (that is, whenever the speaker is sincere in despising the target class). According to Schlenker’s account, whether ‘Bianca is a wop’ is a case of presuppositional failure depends on whether the speaker sincerely believes that Italians are despicable.

On the other hand, the objective presupposition holds whenever the presupposed evaluation is correct. It is easy to see what a delicate question the objective option brings up: What does establish whether the evaluation that hybrid evaluatives convey is correct? That is, how do we establish whether ‘Italians are bad for being Italians’ is true or correct in the context of utterance? The answer to these questions crucially depends on what theory of values one favors. I will have to leave the ‘hot potato’, that is the ultimate core of the problem, i.e. how to establish when an evaluation is correct, to philosophers working in ethics, but in what follows I will try however to unravel part of the knots at stake.

²⁸ Macià (2002: 505-506) seems to hold something like the objective option, when he says:

“Por otro lado, trataremos ‘el imbécil de Alfonso’ asignándole una función f que a cada contexto C le asigna la persona Alfonso, si Alfonso es un imbécil según C, y f no asigna ningún valor a ‘el imbécil de Alfonso’ si Alfonso no es ningún imbécil según C”.

“On the other hand, we will analyze ‘the idiot Alfonso’ as a function f that, for every context C, assigns Alfonso (the person) if Alfonso is an idiot in C, and it does not assign any value to ‘the idiot Alfonso’ if Alfonso is not an idiot in C ” (translation mine).

Then he seems to have changed his mind in Macià (2006, 2011):

“what ‘That bastard John came into the room’ presupposes is that we are willing to treat him in a contemptuous way” (Macià 2006)

“[[Frog]] (C,w) (a) = *True* if in C speakers are disposed to treat French people with a certain kind of contempt because they are French, and a is French
False if in C speakers are disposed to treat French people with a certain kind of contempt because they are French, and a is French;

Nothing, otherwise” (Macià 2011, italics mine)

In the framework I presented, presuppositional failure is brought about by certain utterances featuring slurs and objectionable thick terms. They are characterized by the fact that speakers refuse to apply the predicate and do not accept to judge an utterance where the hybrid evaluative is used as true or false. Let us remind ourselves of the examples proposed in section 2.2, that show how speakers refuse to apply a certain predicate:

80.

A. I heard he was a fag [...]

B. Historians don't generally put it quite that way, but yes, Da Vinci was a homosexual.

81.

Carson: Do you consider that blasphemous?

Oscar Wilde: (...) the word 'blasphemous' is not a word of mine.

One might doubt that the notion of 'failure' is convincing here: After all, it is not a catastrophe for the meaning of the utterance if the presupposed evaluation fails, as competent speakers typically understand utterances featuring objectionable hybrid evaluatives.

An utterance of (82) crucially diverges from (83):

82. The current king of France is bald.

83. Hollande is a frog.

When exposed to (82) in 2016, the addressee has a hard time making sense of the utterance and she might wonder what the speaker wants to say with it. On the contrary, it is quite clear what is happening when (83) is uttered: The speaker is gallophobic. Let me make two remarks. The first one concerns the idea that there are

different kinds of failure. Yablo (2006) introduces the notion of catastrophic vs non-catastrophic failure: “Non-catastrophic presupposition failure then becomes the phenomenon of a sentence still making an evaluable claim despite presupposing a falsehood” (Yablo 2006: 164). The idea is that competent speakers can understand utterances featuring hybrid evaluatives even if the presupposition fails and, crucially, they are able to distinguish those that ‘sound true’ from those that ‘sound false’. Compare:

84. Leonardo da Vinci is a wop.

85. Leonardo da Vinci is a kraut.

A competent speaker is not only able to understand (84)-(85) even if she does not share either of the presupposed evaluations about Italian and German people respectively; she can also see that the ways in which the two utterances are wrong is different. By ignoring the evaluations triggered, she can interpret (84)-(85) as if they meant ‘Leonardo da Vinci is Italian’ and ‘Leonardo da Vinci is German’; that is why we have the intuition that despite the failure, since Leonardo da Vinci was Italian and not German, (84) sounds like something ‘true but with wrong assumptions’ (or true but “objectionably couched”, as Gibbard 2003a puts it) and (85) seems something ‘false and with wrong assumptions’. The second consideration concerns the evaluative dimension. Note that the notion of non-catastrophic presupposition is completely independent from the presupposition being descriptive or evaluative. However, it is reasonable to expect some divergences. Recall that in section 2.2, I discussed the possible reasons for which speakers tend not to articulate the presupposed content of hybrid evaluatives when they want to reject it. Among the possible reasons, I claimed that speakers articulate the presupposed content only when they deem it necessary; when the presupposed content is evaluative (e.g. that French people are bad for being French), speakers might be less prone to intervene than when the speaker is presupposing a false fact (that France has a king in 2016). Simply informing the speaker that she is presupposing something false might change

her beliefs right away, whereas things are more complicated with evaluations. Consider the following dialogue between Lucy and Mary:

86.

Lucy: So your wife is Italian.

Mary: Yes, she's Italian. She's my fiancée, though, we are not married.

This can be an example of non-catastrophic failure. Unlike cases *à la* (82) – the king of France case –, in the exchange above Mary can easily understand to whom Lucy is referring, even if the presupposition that Mary has a wife is strictly speaking false. In this short exchange Lucy is informed that her piece of information about Mary's marital status was wrong: It suffices to Mary to correct Lucy. It is different with evaluative content:

87.

Lucy: So your wife is a wop.

Mary: My wife is Italian and there is nothing bad in that.

Whatever answer Mary could give Lucy, it is not going to immediately change Lucy's evaluation of Italian people: If that was possible, all we would need to do to end racism would be to *explain* to racists that they are wrong about a couple of things. *Pace* Socrates, this is not how things work. While it is relatively effective to intervene to correct some piece of information,²⁹ it is harder to obtain similar results with moral judgements.

This contrast between the descriptive and the evaluative dimension comes in handy in the discussion of failure. When a speaker is presupposing some wrong evaluative content, the participants to the conversation can still make sense of the

²⁹ However, as noticed by Isidora Stojanovic (p.c.), this feature should *not* be taken as the distinctive feature of the evaluative, as it might be just as hard to correct descriptive information involving assumptions on non-existing entities (for instance, it would be hard to intervene when someone utters something like "Mother Teresa is watching us from Paradise").

utterance because what determines the failure of the evaluative presupposition (thus the inappropriateness of the utterance) is an attitude that speakers can detach from the rest of the utterance if they do not endorse it.³⁰ In sum, the contrast between (82) – ‘The current king of France is bald’ – and (83) – ‘Hollande is a frog’ – can thus be explained in terms of catastrophic vs non-catastrophic failure, but also in terms of the *reasons* for which the speaker is taking for granted wrong assumptions: false information on the one hand, wrong moral judgement on the other. In the first case, conversation participants are left wondering what the speaker is talking about when mentioning ‘the king of France’; in the second case, they typically know what the speaker has in mind when talking about ‘frogs’, even without sharing the gallophobic attitude.

3.3.4 A theory of value

The crucial question to which I will *not* provide a conclusive answer concerns, as I mentioned earlier, the following issue: What does it mean for the presupposed evaluation to hold or to be correct in the context of utterance? Or, in other words, what is for a certain property (being Italian, being sexually explicit, etc.) to be bad or good? It depends on what theory of value one favors. I will not choose a specific account, nor I will argue in favor of one. However, I do find interesting insights in the fitting attitude theories (FA). The firsts to develop this kind of account were Henry Sidgwick and Franz Brentano. Sidgwick, from at least the 1884 third edition of *The Methods of Ethics*, defines “the ultimately good or desirable” as “that of which we should desire the existence if our desire were in harmony with reason” (Sidgwick 1884: 108). Brentano (1969 [1889]: 18) similarly claims that “the good is that which is worthy of love, that which can be loved with a love that is correct”. The formulation of the FA theories that is more canonical is provided by Ewing, that defines “good” as “fitting object of a pro attitude” (Ewing 1947: 152), where “pro attitude” covers “any favorable attitude to something” (*Ibidem*: 149), such as “choice, desire, liking, pursuit, approval, admiration”. Similarly, “(...) what is evil

³⁰ See Sreenivasan (2001) for the view according to which one can perfectly understand and assign the correct truth conditions to utterances featuring “alien thick terms”, i.e. thick terms involving a perspective (or “ethical outlook”) that a subject is not willing to take. Interestingly, Sreenivasan endorses a presuppositional analysis of thick terms.

[is] a suitable object of anti-attitudes” (Ewing 1939: 9). The FA approach has had a contemporary revival (see *i.a.* Lewis 1989, Scanlon 1998, Rabinowicz and Rønnow-Rasmussen 2004). I will try to sketch a very general characterization of the proposal, without going through all the possible options (for an overview of the alternative variants of FA, see Jacobson 2010). Jacobson 2010 pinpoints the two main claims that a FA account typically holds. The first claim of FA accounts is the reduction of the evaluative to the deontic: The positive evaluation expressed by ‘good’ is analyzed in terms of what one *has reasons to* feel. The second claim is that the notion of ‘value’ has a response-dependent nature, in the sense that values are partly determined by human responses and attitudes.³¹ In the words of David Wiggins (1987: 206):

“x is good if and only if x is the sort of thing that calls forth or makes appropriate a certain sentiment of approbation given the range of propensities we have to respond in this or that way.”

Note that this account – unlike an entirely dispositional view – leaves room for error. An individual can be mistaken in her evaluation, when the object of her evaluation is unfitting the relevant attitude. This way of characterizing values might shed light on how an account of hybrid evaluatives can make sense of the notion of failure. Utterances featuring hybrid evaluatives are cases of presuppositional failure when the object the descriptive content is about does not fit the expressed attitudes. In other words, if one has reasons to feel some kind of anti-attitudes toward Italians *qua* Italians, then the presupposition expressed by ‘wop’ is correct and an utterance like ‘Bianca is a wop’ does not fail. On the contrary, if that is not the case, if the object in question does not fit the anti-attitude, then it is incorrect to take for granted a certain evaluation and it does not really make sense to utter such a thing.

Note again that I am not defending fitting-attitude accounts³² *over* alternative theories of value, I just observe interesting interplays between this kind of

³¹ For a response-dependent account of evaluatives, see Mišćević (2006) and Thommen (2014).

³² For criticism of this view, see Väyrynen (2006).

approach and a presuppositional analysis of hybrid evaluatives. However, the crucial question as to *how* to established whether certain sentiments of approbation (or reprobation) are warranted remains unanswered and it is somehow beyond the scope of the present work.

Before concluding, note that the issue of the correctness (or fitness) of evaluations does not concern ethicists only, as it has important consequences for the theory of presuppositions in general. We say that for an utterance to be appropriate, its presuppositions must be true. Nevertheless, when the presupposition is evaluative rather than descriptive, it is not clear how this requirement should be interpreted. Either we add a clause that says ‘*when the presupposition is descriptive*’, or we change the truth requirement on the presupposition into something like ‘correctness’, which brings us back to the issue of correctness determination.³³

In a nutshell: With respect to reference and failure, the contrast between objectionable and non-objectionable hybrid evaluatives, i.e. between the evaluatives that convey correct evaluations and those that do not (namely, slurs and objectionable thick terms), is stronger than the contrast between slurs and thick terms. The criterion to establish whether a presupposition is correct (thus, whether an utterance featuring a hybrid evaluative is appropriate or not), depends *inter alia* on what kind of theory of value one endorses, which is still an open question. However, I hope I clarified the relations between the linguistic analysis and the issues at stake in ethics.

4. The Dynamics of hybrid evaluatives: complicity, propaganda, negotiation³⁴

In the previous sections, I presented a presuppositional account of hybrid evaluatives. In the present chapter, I will show how such a linguistic analysis can account for some crucial features of slurs and thick terms that philosophers have been pointing out. I will start from slurs (4.1), by considering features such as the complicity phenomenon and the propagandistic power of slurs (see, among others, Croom 2011, Camp 2013, Stanley 2015, Jeshion 2016b) and then I will move to thick terms (4.2), by considering negotiation, the role of perspectives and conceptual ethics.

³³ I thank Uriah Kriegel for helpful and insightful discussion about this topic.

³⁴ The present section is largely based on Cepollaro (forthcoming).

4.1 Slurs in conversation

As we said, according to the presuppositional account, the derogatory content of epithets is a non-truth-conditional component of meaning. This feature is not just a linguistic technicality. The information that utterances *presuppose* rather than *assert* has a particular status, as in the former case the content is presented as taken for granted, as not open to discussion. A presupposition is a content that is taken for granted in a conversation, that is, something we take to be already part of the conversation background, or to be very compatible with it. We shall see how this feature of presuppositions is crucial to understand how derogatory epithets not only *reflect* discriminatory beliefs and attitudes, but they in fact *promote* them. In particular, I will consider what happens when slurs are used in conversation, by focusing on people's reactions. I will distinguish three scenarios: (a) endorsement, (b) complicity and propaganda and (c) rejection. Before considering the three scenarios, note that the notion of 'common ground' at stake here does not only include descriptive 'informational' content *à la* Stalnaker, but also attitudes and evaluative content. As García-Carpintero points out:

“[contexts] cannot be just propositions, or more in general representational contents, but rather these contents together with commitments toward them in different modes by speakers” (García-Carpintero 2015: 11)

“The main reason is that several phenomena (presuppositional treatments of pejoratives and predicates of taste, forces other than assertion) require that the common ground includes non-doxastic attitudes such as appraisals, emotions, etc.” (*Ibidem*: 11)

(a) Endorsement

The first scenario is the racist context, in which the participants to the conversation endorse the derogatory presupposition. Imagine for example that Jane says ‘Peter’s boss is a chink, isn’t she?’ while she talks to her colleagues. Jane is taking for granted

that Asian people are bad for being Asian, which is in fact common ground in the context of that conversation. Presupposing this content is a way to reinforce the shared discriminatory beliefs as well as to reinforce the racist identity of the group. Geoffrey Nunberg points out that to attack targets is not the only function of slurs (and according to Nunberg, it is not even the primary one); slurs are used by slurs-users to identify themselves, to create a bond based on discriminating someone else (see also Mišćević 2011: 165).

“It’s misleading to characterize all of the latter [the literal uses of slurs] in terms of their use as ‘weapons’ or for that matter to say that they’re in their nature tools for insulting or offending others. Those are unquestionably among the most troubling effects of the uses of certain slurs (...) the focus on the offensiveness of slurs tends to obscure what is usually their primary *raison d’être*. (...) [The] vast majority of the uses of these words occur among the members of the group they belong to, out of earshot of the people they denote. Indeed, a community may have a slur for a group of people that its members have no expectation of ever encountering” (Nunberg forthcoming)

(b) Complicity and propaganda

The second scenario is the one in which the derogatory presupposition triggered by the slur is *not* common ground among the participants. It might sound strange to say that speakers presuppose something that is not common ground, as presuppositions are what we take for granted in a conversations. However, presuppositions can also get accommodated: We will rely on the notion of “informative presupposition” (see, among others, Stalnaker 2002, Simons 2005, Gauker 2008). Speakers can presuppose information that is not already part of the common ground, when they expect that the audience will just accept it without objections. Imagine that a speaker says to her colleagues ‘Peter’s boss is a chink, isn’t she?’, even if she does not know whether everyone agrees that Asian people are despicable for being Asian. If no one says anything, the introduction of derogatory and discriminatory contents transforms and shapes the common ground and it legitimates further uses of slurs as well as

corresponding discriminatory practices. The non-challenged racist utterance turns the context into an explicitly racist one (see Sbisà 1999 for a discussion about how implicit communication relates to ideology). The presuppositional account of slurs provides a reasonable explanation of why slurs are taken to have a propagandistic power: They trigger presuppositions that – if not objected – slip into the common ground and change what the participants to a conversation take everyone else to assume. Here we see the danger of slurs: They do not only reflect the discriminatory practices and beliefs of a certain community, but they also spread them, by imposing such beliefs in contexts where they were not explicitly accepted by everyone. The propagandistic power of slurs is closely related to the phenomenon of ‘complicity’ often noticed by scholars:

“If slurs are devices for introducing presuppositions in the way that descriptions and clefts are, they are devices which through our silence make us complicit in the bigotry of others” (Richard 2008: 19)

“just hearing (...) [the slur ‘nigger’] can leave one feeling as if they have been made complicit in a morally atrocious act” (Croom 2011: 343)

“slurs make recalcitrant hearers feel complicit in the speaker’s way of thinking” (Camp 2013: 333)

“In some cases, merely overhearing a slur is sufficient for making a non-prejudiced listener feel complicit in a speaker’s slurring performance” (DiFranco 2014)

As I argued in Cepollaro (2016), the presuppositional analysis of slurs can naturally account for complicity, as in the absence of objections, speakers are taken to endorse any content that is accommodated in the common ground (Lewis 1979, Stalnaker 2002), no matter what they think or what their intentions are. In other words, when a slur is used literally in a context, if speakers do not object they are *responsible* for letting the derogatory content in the common ground. The way in

which standard presuppositions work in conversation (the fact that they are typically backgrounded and, when they are not, they can be accommodated) can explain complicity and – more broadly – the fact that when hybrid evaluatives are used, a certain moral perspective is taken. This goes along the lines of Elizabeth Camp’s proposal in terms of “slurring perspectives”. As Camp (2013: 330-331) argues, slurs “present contents from a certain perspective, which is difficult to dislodge despite the fact that it is precisely what a nonbigoted hearer most wants to resist” and they “conventionally signal a speaker’s allegiance to a derogating perspective on the group identified by the slur’s extension-determining core”.³⁵ For Camp, “perspectives” are “ways of thinking, feeling” and when speakers employ them, they also invite the addressee to do the same, which explains complicity. Camp claims that the relation between the meaning of perspectival expressions and perspectives is not pragmatic, but semantic. However, she leaves open the question as to how this component of meaning should be spelt out and what properties it is supposed to display. In a way, the presuppositional account aims to explain *what* such a perspective is (a set of moral beliefs and evaluations) and *how* it comes in place (it is presupposed: Either it is backgrounded or it gets accommodated). In other words, while agreeing with Camp on the idea that slurs and thick terms involve a perspective, I also believe that it is crucial to establish *how* these perspectives are encoded in language.

Going back to the phenomenon of complicity, note that the version of the presuppositional account of hybrid evaluatives that explains complicity more naturally is the ‘objective’ option (OO), according to which the pejorative content is analyzed as being about the target class (‘Italian people are bad for being Italian’): The content that slips into the common ground and that is therefore taken to be accepted by every participant is that the target class is despicable. On the other hand, the ‘subjective’ option, proposed in Schlenker 2007 for slurs, has no direct way to explain the complicity phenomenon: If the presupposed content is not about targets (Italians), but it is about the beliefs of the speaker (‘*the speaker believes that* Italian people are bad for being Italian’), it is slightly more difficult to explain why the

³⁵ Note that for Camp, thick terms (together with formal and informal ways of addressing like the French ‘*tu*’/‘*vous*’ or the Italian ‘*tu*’/‘*lei*’) are perspectival expressions in that they signal the speaker’s relation to the semantic content.

participants to a conversation have any responsibility to say something when they do not endorse certain evaluations (see section 3.3.2 for a discussion of the ‘objective’ and ‘subjective’ option). An objection along the same lines has been raised by Camp about Jeshion’s expressivist account. Camp’s point is that if the derogatory content of slurs amounts to the expression of one’s feelings, then the only person who is responsible for the derogation conveyed by the slur would be the speaker (it would be “just the speaker’s problem”). Jeshion’s answer hinges on the idea that contempt has different features from purely subjective feelings such as fear, pain and the like. It has a normative character: The expression of contempt strongly invites the addressee to “sign onto” the speaker’s attitude. If this is correct, expressivists can explain the sense of squeamishness that non-bigots get when they are exposed to slurs as a result of the pressure that contempt provokes. However, they do not automatically explain how the audience can be taken to be *responsible* for the derogatory content conveyed by slurs, or at least not as directly as the presuppositional account of slurs can.

Similarly, I can think of two strategies for the ‘subjective’ version of the presuppositional account to explain complicity, none of which I find completely convincing. The first one is to claim that people are in certain ways responsible for the beliefs endorsed by the people they deal with; in this sense, avoiding challenging the use of a slur would be endorsing the attitude associated with the slur, even if all the slur was doing was to inform about the mental state and dispositions of the slur-user. The second strategy would be to rely on a general notion of so-called presupposition of commonality, i.e., the idea – introduced by Lopez de Sa (2008: 304) for predicates of taste – that “the relevant expression (...) triggers a presupposition of commonality to the effect that the participants of the conversation are all alike in the relevant respects”. If we admit an extension from the taste domain to the moral dimension, we would have the following account: A slur like ‘wop’ activates the presupposition that the speaker despises Italians *qua* Italians and in addition it also triggers the presupposition that the conversation participants feel the same way about Italians (are all alike with respect to their feelings towards Italians).

As we said, most scholars agree on considering complicity as a feature of slurs. However, according to Jason Stanley, the power to shape the context and impose a

certain derogatory content on the audience does not only belong to slurs and in fact does not primarily belong to slurs.³⁶

“In the United States at least, the focus philosophers have placed on explicit slurs is misplaced. (...) The problem with the literature on slurs is that it suggests that there is a clear dividing line between expressions with the properties of slurs and expressions that are not slurs. This assumption is false. Many and perhaps most expressions have the properties that only slurs are supposed to have, not-at-issue content that cannot be denied and is directly added to the common ground. Most words carry with them, in all of their occurrences, not- at-issue meanings that cannot be easily expunged, if at all. Politics involves a constant search for words that do not appear to be slurs, but that carry a not- at- issue content that prejudices political debate. (Stanley 2015: 154-155)

Stanley’s point is that the use of slurs is too obviously a violation of human dignity to be actually able to do any propaganda: The real danger comes from less explicit tools of dehumanization. While agreeing with Stanley on the fact that the propagandistic power of slurs does not primarily concern the *public* political debate, I think that Stanley’s objection misses two crucial points. For one thing, the propagandistic power of slurs concerns *all* the contexts where slurs are used. When scholars talk about the propagandistic power of slurs, they refer to the private settings where slurs are in fact employed, not to public debates where they are typically banned: These private settings are just as effective as public debates in shaping people’s political views. As Stanley himself acknowledges ‘Citizens gather to speak about politics in all sorts of informal settings (...) [that] guide us in our political choices’ (88). Moreover, Stanley claims that many non-slurring expressions carry not-at-issue content with a pejorative content, for example ‘illegal alien’.³⁷ If the association between the pejorative content and these terms is stable, if they carry such a content in all their occurrences, then in my account these terms just count as

³⁶ Cf. Anderson et al. (2012), Langton (2012) and Lenehan (2015).

³⁷ For a discussion of apparently innocuous terms that carry socially problematic content, see Haslanger (2013).

slurs. On the other hand, if they only *sometimes* carry pejorative content, then the association is not systematic and it is not the case that the associated pejorative content cannot be expunged. An example from Stanley (2015) is ‘welfare’, that refers to social assistance and in the States has a negative and heavily racialized connotation. However, such an association is not hard to expunge: It is easy to find non-loaded uses of ‘welfare’, as these examples from the COCA (Davies 2008-) show:³⁸

88. Eisenhower took his turn with the interstate highway system and the creation of the department of health, education and welfare.

89. When we sat down, he was eager to discuss welfare. The Islamic State may have medieval-style punishments for moral crimes (...), but its social-welfare program is, at least in some aspects, progressive to a degree that would please an MSNBC pundit. Health care, he said, is free.

90. The hukou registration system that ties people to their place of ordinary residence deprives migrant workers of rights to welfare, education or property.

In conclusion, Stanley’s worry is that slurs are not the terms that primarily carry and impose certain evaluations, since other non-slurring terms do the same. I argue that when certain expressions behave like slurs (see for example the case of ‘illegal alien’), then they just *are* slurs: *If* certain expressions systematically convey a certain evaluation as not-at-issue content in all their literal occurrences, then such expressions count as hybrid evaluatives. When they do not, that is, when they have non-loaded literal unembedded uses, then there is an important contrast with prototypical slurs, where slurs are in fact strikingly more powerful in imposing a certain perspective on the audience. We shall come back to the non-slurring expressions that can implicitly convey evaluative content in chapter 9.

³⁸ And of course ‘welfare’ has no negative connotation when it means ‘well-being’, such as in “It is very clear that carbon pollution is a danger to public health and welfare”, again from COCA (Davies 2008-).

(c) Rejection

The third and last scenario I consider is the one where, just like in the second one, the derogatory presupposition is *not* already part of the common ground, but unlike the second one, the audience is *not* willing to accommodate the derogatory presupposition. It is the case of rejection, presented in section 2.2. As we saw, to reject a presupposition, it is not sufficient to state ‘it’s not true’. Take the verb ‘to stop’, which is a well-known presupposition trigger. (91) presupposes (π_{91}). As a matter of fact, (π_{91}) survives semantic embedding in (92) and (93).

91. Jane stopped smoking.

π_{91} . Jane used to smoke.

92. Jane didn’t stop smoking.

93. Did Jane stop smoking?

Consider now the following exchange:

94.

A: Jane stopped smoking.

B: No, she didn’t.

B’s denial in (94) presupposes that Jane used to smoke, even if it denies that now she does not: In general, denial targets the at-issue content, while it does not interact with the not-at-issue one. It follows that for the presuppositional account the derogatory content of slurs cannot be rejected through denial, as presuppositions are not part of the at-issue content. Observe:

95.

A: Peter's boss is a chink, isn't she?

B: No, she isn't.

B's denial in (95) is not apt to reject the pejorative content of A's utterance; it only gets to deny the descriptive truth-conditional content, namely that Peter's boss is Asian (see also Jeshion 2013b: 320). We can imagine for B's utterance a continuation like 'No, she isn't. She's a spic'.

In order to reject the presupposed content some other strategy needs to be found. As we saw in section 2.2, one strategy with presupposition is to explicitly articulate the presupposition and explicitly reject it, or to deny it metalinguistically.

96. Jane stopped smoking.

97. Hey wait a minute, Jane didn't used to smoke.

98. I wouldn't say that she "stopped": She never smoke in the first place.

For slurs, we can imagine similar patterns: One can explicitly articulate the derogatory presupposition and reject it or simply refuse the use of the term metalinguistically.

99. Peter's boss is a chink, isn't she?

100. Peter's boss is Asian and there is nothing bad in that.

101. I don't allow you to use racist words in my presence.

The discussion about the three scenarios, together with the way in which the derogatory content of slurs slips into the common ground when it is not challenged, should have revealed so far why slurs are considered dangerous tools. In general, presuppositions are so effective in entering the common ground because rejection has a cognitive as well as a social cost. As noticed in section 2.2, the implicit nature of the derogatory content (that is – in less evocative terms – the fact that epithets encode the derogatory content at the level of presuppositions rather than at the level of truth conditions) makes it hard for speakers to retrieve and reject it, as the best way to reject presupposition is to stop the flow of the conversation, which means to stop being cooperative. So either one accommodates the presupposition or gives up cooperation. In this sense, slurs, intended as means to produce discrimination, are crucially different – and in way worse – from explicit manifestos endorsing and fostering discrimination. If a speaker is confronted with a homophobic manifesto, there can be room for discussion. Take for example this statement from the Congregation for the Doctrine of the Faith:

“[Homosexual] tendency comes from a false education, from a lack of normal sexual development, from habit, from bad example, or from other similar causes, and is transitory or at least not incurable”.³⁹

As unreasonable as it is, the homophobic content is explicitly articulated and thus it can be rejected. It is something to which one can answer ‘That is false’. Also statements that *imply* homophobic beliefs can be relatively easily challenged. Consider these statements⁴⁰ of Antonin Gregory Scalia, an Associate Justice of the Supreme Court of the United States, known for his homophobic ideology:

102. “If we cannot have moral feelings against homosexuality, can we have it against murder? Can we have it against other things?”

³⁹ http://www.vatican.va/roman_curia/congregations/cfaith/documents/rc_con_cfaith_doc_19751229_persona-humana_en.html. Persona Humana is a declaration of the Congregation for the Doctrine of the Faith released in 1975 regarding sexual ethics.

⁴⁰ <http://www.towleroad.com/2016/02/302540/>.

103. “[The Texas anti-sodomy law] undoubtedly imposes constraints on liberty. ... So do laws prohibiting prostitution, recreational use of heroin, and, for that matter, working more than 60 hours per week in a bakery.”

104. “The death penalty? Give me a break. It’s easy. Abortion? Absolutely easy. Nobody ever thought the Constitution prevented restrictions on abortion. Homosexual sodomy? Come on. For 200 years, it was criminal in every state.”

For how wrong and offensive Scalia’s opinions can be, it is something people can disagree about. They can argue that Scalia is wrong by presenting arguments (for example, with respect to (104), one can say that the fact that something was illegal for 200 years tells nothing about the legal status it should have today). On the contrary, when slurs are used, it is much harder to argue, as the derogatory content they trigger is presuppositional and thus taken for granted: It is *imposed on* the audience.

4.1.1 The shortcut of discrimination

In the previous section, I have dwelt on the way in which slurs work in conversation. The question about the effects elicited by slurs brings up the question of the *function* that slurs fulfill. I shall argue that the function of slurs is not just to insult and diminish people, as Jeshion points out, “slurs function to derogate or dehumanize, by which I mean, that they function to signal that their targets are unworthy of equal standing or full respect as persons, that they are inferior as persons” (Jeshion 2013a: 232). As a matter of fact, there might well be other non-slurring expressives that are apt to dehumanize, derogate and label the addressee as inferior (probably expressions like ‘son of a bitch’, ‘piece of shit’, ‘shit-eater’ and the like). The function that most existing slurs – as opposed to other expressives – specifically fulfill is to discriminate: To demean the targets by labelling them as inferior on the basis of characteristics that in *no way* ground or justify derogation, exclusion etc.

Talking about the relation between slurs and the discriminatory stereotype they are associated with, the psychologist Leon Rappoport (Rappoport 2005, quoted in Nunberg forthcoming) talks about “shorthands”:

“Ethnic slurs serve as a kind of shorthand way of referring to the negative qualities associated with any particular group. They are quite specific. Hispanics might be called ‘spics’ and Jews ‘kikes’; each term would stand for a specific cluster of traits assumed to be typical of Hispanics and Jews...” (Rappoport 2005: 46).

Nunberg revises the notion of shorthand with the related notion of “shortcut”:

“Stereotypes, negative and positive, are among the cognitive shortcuts we rely on to make sense of the world and to guide our responses to it. The utterance of a slur very often evokes or foregrounds a negative stereotype of its target, which is one reason why people use these words: “What do you expect from a ____?” Those stereotypes in turn can serve to legitimate various responses to the group, sometimes by dehumanizing or marginalizing its members, and other times merely allowing us to discount them (when we call a publicist a flack it’s by way of questioning her journalistic integrity, not her basic humanity)” (Nunberg forthcoming)

Nunberg’s idea is that stereotypes count as cognitive shortcuts because they allow people to form judgements that do not have to be based on experience and acquaintance, but just on the assumption that belonging to a certain category warrants certain features. If you assume that, for example, being Italian goes together with being a latecomer, talking loud, being obsessed with food and the like, then you are justified in inferring from the fact that someone is Italian that she instantiates most of the above properties; and you are justified, for instance, in avoiding hiring an Italian person for a job that requires punctuality. Nunberg’s description of what a stereotype is is convincing. Nevertheless, I shall argue that this is not the end of the story when it comes to slurs and the reason has to do with certain features of presuppositions. According to the analysis I just presented, slurs are presupposition triggers. As we saw, presuppositions are what speakers take for granted in conversation: When speakers presuppose a certain content, they present it as

accepted and uncontroversial (see Sbisà 2007). Because of this feature, those scholars that analyze discourse in Court dedicate special attention to how presuppositions are used to impose a content that was not actually shared (see *i.a.* Coulthard and Johnson 2010 and Ehrlich and Sidnell 2006): “Presupposition is at least one micro-mechanism in language use which contributes to the building of a consensual reality” (Chilton 2004: 64). My point being that slurs are not just the *expression* of a stereotype, which would already be, in Nunberg’s terms, a shortcut, but they also impose the stereotype, as we saw, by presenting it as shared, which is an additional shortcut. As said before, slurs do not just reflect discriminatory beliefs, but they actually promote discrimination without arguing for it, just presenting it as given.

Wrapping up, presupposing in general is a useful tool to give for granted what everybody accepts already and it is also a means to establish commonality, possibly imposing it. This can concern not just the knowledge speakers share, but also the values they share. If the set of beliefs underlying discrimination was object of discussion, we might have far less discriminatory practices than we do. The reason why slurs are considered dangerous tools is that they impose discriminatory contents without leaving room for disagreement. Partly because of such reasons, both scholars (Hornsby 2001, Anderson and Lepore, 2013a) and authorities promoted ‘silentism’ as policy: In order to fight the discrimination perpetrated by slurs, their use (and according to some, their mention, too) must be prohibited. Such terms need to be removed from circulation until their offensive potential fades away. However, it is not clear how this would happen. Bianchi (2014) proposes an alternative analysis, where she rejects silentism as policy and she suggests that the process through which slurs can lose their derogatory potential is actually Appropriation (see Part III). ‘Appropriation’ refers to those cases where the members of a group can use among themselves the slur targeting their own group, in such a way that the slur is not offensive anymore in those contexts; on the contrary, appropriated slurs are used to express solidarity and underline intimacy. This is for example what happened to the term ‘gay’, that used to be derogatory and is neutral nowadays, after a process of appropriation (see Brontsema 2004). According to Bianchi, silentism is not apt to diminish the derogatory potential of slurs, while Appropriation is able to subvert it

and, in time, even delete it. The analysis of epithets that I just proposed does not take a clear stance with respect to silentism. What it does suggest is that the main source of danger of slurs is their conveying a derogatory content in an implicit way. So a good practice would be to train people to properly respond to slurs when they come across them. Just as in the case of any presupposition, a proper response would consist in articulating the presupposition, i.e. making explicit the derogatory content of the slur and reject it. Some LGBTQ rights organizations do promote this kind of training, by teaching people how to react when they come across an utterance of an homophobic slur. This is one of their suggestions mentioned in section 2.2, which closely resembles what it would mean to retrieve and reject a presupposition in the case of slurs.

105. What you just said was really inappropriate because you are implying that there is something wrong with being gay or lesbian when there isn't.⁴¹

4.2 Thick terms: negotiation and conceptual ethics

In section 2.2, we discussed how people can react to the use of hybrid evaluatives. We have observed that in order to reject the presupposed evaluative content of an utterance, denial is typically not enough, as it targets the at-issue content. Looking at the semantics of hybrid evaluatives, we have discussed different ways in which hybrid evaluatives are context-sensitive. The first one is brought about by the evaluative content, where terms like 'good' and 'bad' occur. The second one concerns thick terms and it has to do with their descriptive content, that present a higher degree of context sensitivity than slurs (3.2). We shall now see that the occurrence of terms like 'bad' and 'good' and the context sensitivity of certain descriptive contents give rise to interesting observations concerning disagreement and negotiation. In general, as Väyrynen (2013: 34) points out:

“ (...) an expression that is evaluative in meaning needn't fully specify the content of the standard as a matter of meaning. What things count as morally good, for instance, isn't determined purely by conceptual structure.

⁴¹ <http://www.healthiersf.org/LGBTQ/InTheClassroom/docs/Responding%20to%20Homophobia.pdf>.

A shared concept of moral goodness allows for substantive disagreement over which things are morally good.”

In addition to the issue of whether certain descriptive properties (e.g. ‘being Italian’) correctly deserve a certain evaluation (e.g. ‘being bad’), in the case of thick terms there is room for further negotiation about the parameters that need to be established, that is, the threshold after which the evaluation is triggered.

We shall distinguish between at least three types of disagreement involving thick terms. Imagine someone says ‘That movie is lewd’. One can disagree at least in three ways:⁴²

(i) *Descriptive disagreement*

One can disagree on whether the object in question can count *at all* as sexually explicit beyond conventional boundaries.

E.g.: ‘No, that movie is not lewd at all. You are confusing it with *La vie d’Adele*’.

(ii) *Metasemantic disagreement*

One can disagree on whether the object in question passes the relevant threshold for ‘lewd’, that is, one can disagree about where to put the threshold.

E.g.: ‘No, that movie is not lewd. The sexual tension between the characters is just alluded to’.

(iii) *Evaluative disagreement*

One can disagree on a more radical ground, that is, disagree with the very idea that being sexually explicit beyond conventional boundaries can in *any* circumstance deserve negative moral evaluation.⁴³

E.g.: ‘No, that movie is not lewd: it is sexually explicit, I grant that, but there is really nothing bad in that’.

⁴² See *i.a.* Hare (1963), Blackburn (1992) and Gibbard (1992, 2003a, 200b).

⁴³ Gibbard (2003b) calls this kind of radical disagreement “demurring”.

In (i) we imagine that the speakers share the very same concept of ‘lewd’, they use the term in the same way; the disagreement only concerns the descriptive level. In (ii), on the other hand, we imagine that both speakers are ‘lewd’-users, but they associate the term to different standards, to the extent that for one a certain object can be called ‘lewd’ while for the other it cannot. In (iii), speakers disagree on whether it makes sense at all to employ the term ‘lewd’, that is, whether there is something like being bad on the basis of being sexually explicit beyond conventional boundaries. Note that evaluative disagreement does not only concern objectionable thick terms. Another such case would be the following. Suppose that for a given descriptive property, two distinct terms convey opposite evaluations. For example, suppose that the terms ‘brave’ and ‘reckless’ have exactly the same descriptive properties.⁴⁴ Two speakers might disagree on whether a certain action is brave or reckless, while agreeing on the descriptive properties of the action at stake. This case counts as ‘evaluative disagreement’ because speakers disagree on whether certain properties instantiated by an action deserve a negative or positive evaluation.

Let’s see whether we can have the same kind of disagreement for slurs. Imagine someone says ‘Bianca is a wop’.

(i) *Descriptive disagreement*

One can disagree on whether the object in question instantiates the descriptive property or not.

E.g.: ‘No, she is not a wop. You are mistaking her for Maria. Bianca is German’.

(ii) *Metasemantic disagreement*

⁴⁴ This might be an extreme abstraction, made for the sake of the argument. As a matter of fact, the descriptive content of ‘brave’ does not simply concern the amount of risk a person is willing to take, but also the *reasons* for which she is taking it. The same action could be judged ‘brave’ or ‘reckless’ in light of what it is done for. For instance: Take the action of entering a building on fire and risk your life. A enters the building and risks her life to get her favorite coffee cup. B enters the building and risks her life to save her neighbor. It is plausible to say that if a speaker applies the predicate ‘brave’ to B and ‘reckless’ to A, she is conveying a positive and a negative evaluation about different descriptive properties. And given that both A and B risk their lives by entering a building on fire, the difference must be in *why* they do so.

? One can disagree on whether the object in question passes the relevant threshold of ‘Italianness’ for ‘wop’.

E.g.: ‘She is not a wop. She only has Italian parents, but she was raised in Germany’.

(iii) *Evaluative disagreement*

One can disagree on a more radical ground, that is, disagree with the very idea that being Italian can in any circumstance deserve negative moral evaluation.

E.g.: ‘She is not a wop, there is no such thing. She is Italian and there is nothing bad in this’.

As (i) and (iii) show, it is easy to imagine descriptive and evaluative disagreement with slurs, whereas the metasemantic type is harder to get. We can build up cases like (ii) that resembles the metasemantic disagreement, but they are definitely atypical uses of slurs. My guess is that there is some kind of figurative talking in (ii), that employs a certain term, ‘wop’, to talk about the stereotypical properties associated with ‘Italian’. In general, we might expect that when people disagree about slurs, it is usually a case of descriptive or evaluative disagreement, whereas we expect metasemantic disagreement to typically concern the use of gradable thick terms.

Let’s now consider (ii) and (iii)-kinds of examples. One might want to discard these uses of hybrid evaluatives as metalinguistic uses. As a matter of fact, in chapter 7 I will provide three linguistic tests to show that the uses of negated evaluatives *à la* (iii) are metalinguistic uses of language. However, there is an interesting point to make about this kind of disagreement: The fact that disagreement is metalinguistic does not make it any less substantial or interesting. Plunkett and Sundell (2013) talk about *metalinguistic negotiation*, cases of metalinguistic disagreement that are tacit negotiations over how a certain term should be used. It involves two notions: ‘metalinguistic usage’, from Barker (2002) and ‘conceptual ethics’, from Burgess and Plunkett (2013a, 2013b). The former label, ‘metalinguistic usage’ (as opposed to ‘mention’) refers to the cases where a certain term is used (not mentioned) in order to

convey information concerning how that term should be employed in context. The latter label, ‘conceptual ethics’ comes from Burgess and Plunkett (2013a, 2013b):

“claims about how one ought (or would do well) to think and talk are nearly as ubiquitous in philosophy as their descriptive counterparts, not to mention their prevalence in ordinary discourse. (...) we might call such normative and evaluative issues about representation ‘conceptual ethics’” (Burgess and Plunkett 2013a: 1091)

“That is, the ethics of concepts or of what concepts to use – rather than ethics done in an especially conceptual way (whatever that might look like)” (*Ibidem*: 1098)

I endorse Plunkett and Sundell (2013)’s claim that the disagreement being metalinguistic does not mean that disagreement is ‘merely verbal’: It is worth engaging in such a dispute because, as Plunkett and Sundell claimed, how we use words *matters*. Moreover, I believe that the case of disagreement and negotiation involving thick terms satisfactorily fits in the field of conceptual ethics. As a matter of fact, Burgess and Plunkett themselves include, among the various instances of conceptual ethics, the discussion about slurs and epithets, while Plunkett and Sundell (2013: 13) consider linguistic expressions that are context-sensitive in virtue of being gradable (or “relative gradable adjectives”, after Kennedy 2007). The way in which Plunkett and Sundell (2013:3) characterize metalinguistic negotiations is very interesting with respect of hybrid evaluatives:

“(...) we argue that many disagreements about conceptual ethics are not expressed explicitly. Metalinguistic negotiations are the most important instances of this implicit, or tacit, kind. Such disputes may not at first glance appear — either to the speakers themselves or to the theorist — to reflect disagreements about concept choice. But in fact they do reflect disagreements about concept choice. Many disputes that theorists have thought must be analyzed in terms of the shared literal content are thus best

analyzed as speakers using their words in different ways, advocating (metalinguistically) for their preferred usage.”

To characterize metalinguistic negotiation as typically implicit accounts for the observations discussed in section 2.2 about the fact that when hybrid evaluatives are objects of a dispute, the contested issue is rarely articulated and more often left implicit. Moreover, in my account hybrid evaluatives convey their evaluative content in an implicit manner, that is, *via* presupposition. One of the main points of Plunkett and Sundell is to make room for non-canonical genuine disagreement, that is, genuine disagreement that does not necessarily focus on the truth or correctness of the content *explicitly* expressed by the speakers.

Note, moreover, that in my analysis of the disagreement involving slurs and thick terms, the contestability that gives rise to disagreement does not especially belong to the evaluative dimension, but it characterizes both descriptive and evaluative content. To be systematically contestable is not what is distinctive of the evaluative. In this respect, I agree with Väyrynen’s claim that “essential contestability offers no deep illumination of the evaluative in particular (...)” (Väyrynen 2014: 472), as essential contestability has mainly to do with context dependence and multidimensionality, rather than with evaluation as opposed to description.⁴⁵

In conclusion, the complex semantics of hybrid evaluatives – and thick terms in particular – brings about interesting issues concerning negotiations and disagreement. I propose to distinguish at least three types of disagreement, that I call descriptive, metasemantic and evaluative. By appealing to the works on Burgess and Plunkett (2013a, 2013b) and Plunkett and Sundell (2013), I suggest that at least two kinds of disagreement involving hybrid evaluatives and thick terms in particular are very relevant within the domain of so-called conceptual ethics and support Plunkett and Sundell’s claim that genuine disagreement can well be non-canonical.

⁴⁵ Väyrynen’s target is the idea – often implicit – that the essential contestability can be a sufficient condition for the evaluative dimension (Gallie 1956) or that it can at least help demarcating the evaluative (see, *i.a.* Roberts 2013: 89: “A concept being evaluative is sufficient for it being essentially contestable, though it may not be necessary: there may be non-evaluative concepts that are essentially contestable”).

5. Defending the presuppositional account

In this chapter I will assess the main objections that have been raised against presuppositional accounts of slurs and of thick terms respectively, that were discussed in Cepollaro and Stojanovic (2016). As far as slurs are concerned, the main objection involves the so-called *cancellability*, namely the possibility to prevent an embedded presupposition from projecting. Moreover, I will discuss *redundancy* and *expressivity*; in the case of thick terms, there are more general perplexities, among which I will consider *triggering*, *appropriateness* and *negation-infelicity*. I will show how these objections do not constitute an obstacle to the presuppositional approach.

5.1 Slurs and presuppositions

The main objection that has been raised against a presuppositional analysis of slurs concerns the so-called ‘cancellability’ problem, namely the fact that presuppositions usually do not survive under certain embeddings. The prototypical case is the case of ‘binding’, in which the antecedent of a conditional entails the content presupposed by the consequent; in such environments, the presupposition does not project, as we can see in (106). The case of slurs seems to be different under this respect, as we see in (109). Compare:

106. Jane’s sister must be smart.

π_{106} . Jane has a sister.

107. If Jane has a sister, Jane’s sister must be smart.

108. Jane must be a dyke.

π_{108} . Homosexual women are bad because of being homosexual.

109. ? If homosexual women are bad because of being homosexual, then Jane must be a dyke.

(109), in addition to being odd, does not seem to prevent (π_{108}) from projecting,

whereas (107) succeeds in preventing (π_{106}) from projecting. Whether one should consider this contrast between (107) and (109) as a knock-down argument against the presuppositional account of slurs depends on whether one should consider cancellability as the crucial test for presupposition. In Cepollaro (2015) and Cepollaro and Stojanovic (2016), we argue that independently from slurs, that is not the case. Our claim is that the ease of cancellation comes in degree, rather than being the distinctive feature of presupposition, and that many other factors come into play.

Ever since cancellability was proposed as a diagnostic test for presupposition (Karttunen 1971, Chierchia and McConnell-Ginet 1990, *i.a.*), many authors have come up with problematic examples showing that the presuppositions of some triggers are very hard to neutralize.

The gender features of pronouns are a case at point; according to some authors like Heim and Kratzer (1998), ‘she’ presupposes that the referent of the pronoun is female.⁴⁶ Embedding the pronoun in a conditional does not neutralize such a presupposition:

110. ?? If John were female, she would be popular among the boys. (Sudo 2012)

Likewise, most presuppositions are neutralized (they do not project), in disjunctions of the form ‘not π or φ_π ’, where φ_π is a presupposing utterance and π is the content presupposed by φ_π . For example, take (111), that presupposes (π_{111}). If we embed (111) in the disjunction, we observe that the disjunction (112) does not presuppose that Jane used to smoke:

111. Jane stopped smoking.

π_{111} . Jane used to smoke.

112. Either Jane never smoked, or she stopped.

⁴⁶ “On this account of gender features, it is not strictly speaking impossible to use a feminine pronoun to refer to a man. But if one does so, one thereby expresses the presupposition that this man is female” (Heim and Kratzer 1998: 245).

Nevertheless, this neutralizing effect does not work for gender features. (113) presupposes (π_{113}), but such a presupposition is *not* felicitously suspended in (114).

113. She hangs out with boys.

π_{113} . Jesse [the referent of ‘she’] is a girl.

114. ?? Either Jesse is not a girl, or she hangs out with boys. (From Sudo 2012)

Similarly, presuppositions triggered by slurs are *not* neutralized in disjunction of the form ‘not π or φ_π ’.

115. ?? Either homosexual women are not despicable for being homosexual, or Jane is a dyke.

Again, the evaluative content of slurs is not neutralized in constructs of the form ‘If π , then φ_π ’, where φ_π is a presupposing utterance and π is the content presupposed by φ_π ; the antecedent of the conditional contains the presupposition triggered by the consequent; the conditional is preceded by a clause where the speaker expresses agnosticism with respect to the presupposed content in order to make the utterance easier to understand. In this respect, slurs pattern like gender features of free pronouns and particular triggers like ‘too’ (see Chemla and Schlenker 2012) and unlike standard presupposition triggers like ‘stop’:

116. ?? I am not sure whether Italians are contemptible. But if Leonardo is a wop, I’m sure he will understand some French.

117. ?? I don’t know if Jesse is a boy or a girl. But if she is grumpy, I will take her to the zoo. (Sudo 2012: 30)

118. ?? I have no idea whether John read that proposal. But if Bill read it

too, let's ask them to confer and simply give us a yes no response. (Abusch 2010: 40)

119. I do not know whether John used to smoke. But if he stopped smoking, it must have been hard.

The fact that the evaluation conveyed by slurs always survives embedding does not discard the presuppositional analysis: As we have seen, there are independent reasons to doubt that cancellability and suspendability are distinctive features of presuppositions in general. On the contrary, if such a presuppositional analysis of hybrid evaluatives is correct, it will constitute further evidence that challenges the cancellability requirement.⁴⁷

The hard neutralization of the evaluative component of slurs requires nevertheless an explanation: Why is it so hard or nearly impossible to prevent the derogatory content of slurs from projecting in embedded contexts? One notion that could be helpful here is the soft/hard triggers distinction. According to some scholars (see Abusch 2002, Abbott 2006, *i.a.*), we can distinguish soft and hard triggers, on the basis of how hard it is to neutralize the presupposition that they trigger.⁴⁸ According to Abusch (2002), who introduced this distinction, *soft* triggers introduce a pragmatic presupposition and *hard* triggers a semantic one. Leaving aside the whole semantic/pragmatic debate on presupposition, we can take the distinction to amount to the following: Some presupposition triggers are more sensitive to the context and the extra-linguistic knowledge than others; this soft/hard parameter results in neutralization being easier or harder. In the case of slurs, the presupposition always tends to project and survive semantic embedding.

Nevertheless, as Abrusán (2016) points out, things are way more complicated than that and many factors are called into play when it comes to neutralization. Moreover, we should not neglect the fact that on top of all this, slurs bring about a taboo-effect that should not be underestimated. The fact that slurs are prohibited words could

⁴⁷ About the heterogeneity of projective content, recall the proposal of Tonhauser et al. (2013) discussed in section 2.1.

⁴⁸ The distinction between hard and soft triggers is not uncontroversial, though. For a recent criticism, see Abrusán (2016).

seriously affect our linguistic judgments in detecting projection. People often find it offensive to merely mention taboo terms (think about the need for euphemisms like ‘the F-word’, ‘the N-word’ etc.). The meaning (*lato sensu*) is not enough to explain why it sounds offensive to mention a term without using it. On the other hand, we expect such a taboo effect to be less salient, if not entirely absent, in the case of thick terms.

The second objection voiced against the presuppositional account of slurs concerns *redundancy* and it was raised by Nunberg (forthcoming), Bianchi (p.c.) and Jeshion (p.c.).⁴⁹ The idea is the following: If certain contents are lexically encoded in the meaning of a term, then the explicit assertion of those contents would strike us as redundant.⁵⁰ We will use the symbol ‘*’ to indicate redundancy (while ‘?’ and ‘??’ indicate infelicity and strong infelicity, respectively). Note however that redundancy is *per se* a slippery issue, as one can formulate examples and contexts such that it makes sense to emphasize a certain content without redundancy. Assuming that ‘bachelor’ means ‘unmarried man’, we expect (120) and (121) to be redundant:

120. * Bachelors are unmarried.

121. * Bachelors are men.

I will focus on three important factors to take into account when discussing redundancy and hybrid evaluatives: (i) order, (ii) reinforceability of non-descriptive content and (iii) precise individuation of descriptive content.

The first thing to notice when discussing redundancy is the order of the clauses. For standard presupposition we expect redundancy to occur for utterances of the form ‘ φ_π and π ’, where φ_π is a presupposing utterance and π is the content presupposed by φ_π and *not* for utterances of the form π and φ_π . We can observe this pattern for the factive ‘stop’:

⁴⁹ Note that Kyle (2013) has an argument based on redundancy against the conventional implicature analysis of thick terms; however, for how he characterizes the two approaches, my presuppositional proposal resembles the conventional implicature one more than it resembles the account that Kyle calls ‘presuppositional’. We shall come back to his objections in chapter 8.

⁵⁰ Except for metalinguistic uses aimed at defining the term at stake.

122. * Bianca stopped smoking and she used to smoke.

123. Bianca used to smoke and she stopped.

We see that slurs pattern like standard presuppositions in this respect:

124. Italians are despicable for being Italian and Bianca is a wop.

125. *Bianca is a wop and Italians are despicable for being Italian.

In the case of thick terms, the contrast is less strong (and we shall formulate some hypotheses of why this is the case in chapter 10):

126. It is good to be liberal in one's giving and sharing and Bianca is generous.

127.* Bianca is generous and it is good to be liberal in one's giving and sharing.

The second factor to take into account is that while the repetition of descriptive content typically leads to redundancy, non-descriptive content can just get strengthened when repeated (Potts 2005). Nunberg's perplexity neglects this aspect:

“As Sadock (Sadock 1978) notes: ‘Since conversational implicatures are not part of the conventional import of utterances, it should be possible to make them explicit without being guilty of redundancy.’ In this regard prejudicials [slurs] contrast with the hybrid words with which they're often lumped, where the evaluation is genuinely part of the word's meaning and hence can't be nonredundantly predicated of it. Utterances like ‘Toadies are obsequious,’ ‘Fleeing someone is unfair’ and ‘Shrill sounds are unpleasant’ are likely to elicit the reaction ‘So what else is new?’”

I disagree on two points: (i) I do not find a contrast in redundancy between slurs and thick terms, which for Nunberg should – unlike slurs – lexically encode evaluative content (see the examples above and below) and (ii) I do not find infelicitous the utterances where the *evaluative* content is strengthened, such as the following:

128. Wops are bad.

129. Nancy is generous and she is good in that respect.

On the other hand, I do find redundant the following utterances, where the *descriptive* content is repeated:

130. *Wops are Italian.

131. *Nancy is generous and she is liberal in her giving and sharing.

The third issue to consider is that in order to test speakers' judgements on redundancy, both the evaluative content and the descriptive content need to be spelled out precisely. Nunberg's remarks on redundancy stem, in my opinion, from an imprecise choice of the descriptive content of evaluatives. Take slurs with a complex descriptive content such as 'JAP', that applies to wealthy young Jewish women:

132. *Nancy is a JAP and she's Jewish.

133. *Nancy is a JAP and she's a girl.

134. Nancy is a JAP and she's bad in that respect.

If one takes the descriptive content of 'JAP' to be just 'Jewish', then one cannot explain the redundancy of (133). The same goes for 'fairy', a slur against

homosexual men that display features or behaviors that are taken to be stereotypically feminine. If one takes the descriptive content to be ‘homosexual’, then one cannot explain the redundancy of (135)-(136), that contrasts with the non-redundancy of (137).

135. *Fairies are guys.

136. *Fairies are effeminate.

137. Fairies are disgusting/bad/ridiculous.

In a nutshell, redundancy does not constitute an argument against the presuppositional view, if we take into account the three factors I have underlined.

The last objection that I discuss is the worry, voiced by Richard (2008) and Jeshion (2013a), that the presuppositional account is ‘too descriptive’, or ‘not expressive enough’. Let me consider two aspects: what slurs presuppose and what they can express. For one thing, what slurs trigger is an evaluative content: It would not be correct to characterize it as ‘descriptive’. What slurs presuppose is an evaluation that involves moral judgement,⁵¹ not a description of a state of affairs. In using a slur a bigot is assuming a certain negative evaluation of the target, as if such an evaluation was not the personal opinion of the speaker, but a well-known widely accepted *fact*. As we discussed in chapter 4, this explains how slurs are used among bigots, who reinforce their shared beliefs towards the targets. Richard’s and Jeshion’s worry could then concern the second-person uses of slurs, i.e. the cases where slurs are used as insults, as weapons to express contempt. As Jeshion claims, contempt can be expressed in different ways: prosody, intonation, facial expressions and the like. As a matter of fact, people can clearly express their contempt while using any kind of expressions (even though, depending on what the most common ideologies are, it is easier to express contempt by using certain expressions rather than others: For instance, in many countries it is easier to use ‘woman’ as an insult rather than ‘man’).

⁵¹ For a discussion of the different ways in which moral judgements can be analyzed as having propositional and conative content, see Campbell (2015).

It is plausible to say that when slurs are used as insults, people express their contempt toward the targets by employing various means and, above all, by using those terms that trigger a derogatory content. Moreover, as we saw, the under-determination of the moral judgement is compatible with the under-determination of the attitude conveyed: not just contempt, but also dismissiveness, derision, fear, patronage, superiority and so on. In this sense, rather than being ‘not expressive enough’, the presuppositional account can accommodate the plethora of different attitudes that bigots hold towards their targets.

5.2 Thick terms and presuppositions

In this section, I will consider four worries voiced against the presuppositional account of thick terms, the first two from Väyrynen (2013) and the second two from Kyle (2013).

The first issue, the *triggering* problem, is that thick terms do not resemble any of the standard presupposition triggers such as factive verbs, definite descriptions, cleft sentences and the like. According to Väyrynen, it would be *ad hoc* to simply add new items to the list of presupposition triggers, as argued in Cepollaro and Stojanovic (2016). We believe that this worry is ill-grounded. Presupposition triggers are not items that have been introduced as such to any natural language, in the way in which lists of items are stipulated e.g. for formal languages. For this reasons, it is only to be expected, from the way natural language works and evolves, that properties of expressions that have gone unnoticed are brought to light. The evaluative dimension of language is one that, at least in the realm of formal linguistics, from which the notion of presupposition comes, has been under-investigated. I thus fully embrace Väyrynen’s point that hybrid evaluatives are a brand new kind presupposition-triggers. However, Väyrynen’s point underlines that, unless we intend to add to the list of presupposition triggers each hybrid-evaluative term as a singular item (as it happens for triggers such as ‘too’, ‘again’, ‘even’), the presuppositional account is committed to the claim that hybrid evaluatives constitute a linguistic category (like factive verbs or aspectual predicates).

The second issue raised by Väyrynen, the *appropriateness* problem, has to do with the idea that presupposing is taking something for granted, assuming that the

audience either does so as well or is ready to do so: “A cooperative speaker shouldn’t presuppose something that is more properly put forward as a debatable assertion” (Väyrynen 2013: 113). Consider a conversation among ‘lewd’-users and ‘lewd’-objectors, namely people who respectively accept and refuse that things that are sexually explicit beyond conventional boundaries are bad because of that. Suppose that a ‘lewd’-user uses the term ‘lewd’. Since, by hypothesis, the evaluation is *not* already part of the background knowledge, given that it is rejected by the ‘lewd’-objectors, it would be uncooperative of the speaker, i.e. inappropriate, to presuppose such an evaluation. As Väyrynen points out, the relevant scenario is one where the speaker is fully aware that the audience (or part of it) does not share the presupposed evaluation. This corresponds to the well-known case of *informative presupposition* mentioned in section 4.1 (see *i.a.* Stalnaker 2002, Simons 2005 and Schlenker 2007), where speakers can introduce new information by means of presupposing it. In these cases, we find it plausible to imagine that speakers do act as if their evaluations were shared and also that sometimes they impose them on purpose (expecting or not some kind of protest). In fact, in those scenarios, using a thick term whose evaluation is rejected by the majority of bystanders *is* indeed inappropriate. In a way, slurs themselves provide a crystal clear example of how inappropriate it is to presuppose an evaluation that is *not* shared by the participants of a conversation. However, this case of inappropriateness does not lead to infelicity, as we have seen in section 2.1.

Let’s now consider Kyle (2013)’s objection about *negation-infelicity*, that is the fact that an utterance like (128) sounds contradictory.

138. ?? “Nancy is generous, and she’s not good in a way”. (Kyle 2013: 4)

Kyle takes (138)’s infelicity to prove that “the thick concept *generous* conceptually entails *good in a way*”. However, this is an argument against a strictly pragmatic account of thick terms: The presuppositional account that I endorse holds, as we saw, that the evaluation is part of the *encoded* linguistic meaning of hybrid evaluatives and predicts the patterns observed in (138). As a matter of fact, standard presuppositions such as those triggered by the factive ‘stop’ display the very same behavior:

139. ?? Nancy stopped smoking, and she never smoked.

We shall come back to this point in chapter 8 and make a few remarks on Kyle's use of the label 'semantic', that he takes to be synonymous with 'truth-conditional', but in fact should also include the notion of lexical presupposition. In general, the objections raised by Kyle allegedly against the presuppositional account are against a pragmatic theory that takes presuppositions to be conversationally driven, not against a theory of lexical presuppositions. In particular, he argues that it is possible for (138) to sound felicitous, if we take a context where the speaker does not take for granted nor background any positive evaluation of people who are liberal in giving and sharing. He offers the following example:

“Bob: «Nancy is highly controlled in her giving and sharing, and that's what makes her fiscally smart. She is not generous. But she's not selfish either. I admire her approach to finances».”

?? “Sue: «I disagree with you, Bob. Nancy is generous, and she's not good in any way»”. (Kyle 2013: 11)

According to Kyle the awkwardness of Sue's utterance cannot be explained by a presuppositional account of thick terms, because the speakers are by hypothesis not taking for granted the evaluative content that gets negated in the last conjunct. However, such an hypothesis does not really make sense in the first place, if we have in mind presuppositions that are activated by a certain lexical trigger. It would be like saying: Take a context where no one takes for granted or backgrounds that Nancy smoked in the past and observe the following dialogue:

140.

Bob: «Nancy does not smoke, nor has she stopped smoking»

Sue: ?? «I disagree with you, Bob. Nancy stopped smoking, and she

never smoked»

Sue's utterance is infelicitous because it does not matter whether the speakers are taking for granted that Nancy used to smoke in the past; by saying that she 'stopped' smoking, they are activating such a content and *this* is what determines the infelicity of Sue's utterance.

6. Conclusion: the residual differences between slurs and thick terms

In this first part, I presented a presuppositional analysis of hybrid evaluatives, a category including slurs and thick terms. After some preliminaries about how to define hybrid evaluatives and how to distinguish them from very close notions (chapter 1), I focused on their presuppositional behavior with respect to projection and rejection and I discussed an alternative explanation of the data (chapter 2). In chapter 3, I sketched a semantics for hybrid evaluatives analyzing their evaluative and their descriptive content respectively. I assessed the issues of reference and extension and spelt out what is exactly presupposed by hybrid evaluatives. I discussed how to interpret the presuppositional failure in the case of evaluative presuppositions and showed how these topics relate to highly debated questions about values in ethics. In chapter 4, I discussed the effects of the use of hybrid evaluatives in conversation, while investigating their function: I showed how the complex semantics of hybrid evaluatives brings about interesting issues concerning negotiations and disagreement. Chapter 5 was dedicated to assessing the main objections that have been raised against presuppositional accounts of slurs and of thick terms respectively and showing how these objections do not constitute an obstacle to the presuppositional approach. The picture that emerges from these first chapters is the following: Both slurs and thick terms pick out a certain descriptive property and, at the same time, they trigger an evaluation on that content. For example, 'wop' means 'Italian' but at the same time it triggers the presupposition that Italians are bad because of being Italians; 'lewd' means 'sexually explicit beyond conventional boundaries' and triggers the presupposition that things or individuals that are sexually explicit beyond conventional boundaries are bad because of being so and so on.

I will conclude this part by discussing the residual differences between slurs and thick terms. My aim is to observe the phenomenal differences and distinguish superficial differences from authentic ones. To this end, I will compare certain aspects of the semantics of thick terms and slurs respectively and their projective behavior. I will try to show how the features of the semantics of slurs and thick terms can explain the phenomenal divergences.

6.1 The semantics

We observed some differences between the *descriptive* content of thick terms and slurs. For one thing, the descriptive content of thick terms consists of gradable and multidimensional properties, whereas for slurs the descriptive content usually consists of properties that do not display gradability or multidimensionality. Moreover, the former are about people, behaviors and things, whereas the latter seem to specifically target individuals and groups. In the case of a slur like ‘wop’, it is sufficient that the subject is Italian for the slur to apply; in the case of thick terms like ‘lewd’, there must be a certain amount of sexual explicitness that is contextually determined for the term to apply. In addition to these differences in their extension, another divergence in their descriptive content concerns the features of the so-called neutral counterpart. As a matter of fact, slurs usually have a specific term as a non-loaded counterpart (‘Italian’ for ‘wop’, ‘gay’ for ‘faggot’, ‘Jew’ for ‘kike’, etc.), whereas thick terms tend to lack an exact counterpart and their descriptive content typically corresponds to a more complex paraphrase.⁵² I do not take the existence of a neutral counterpart consisting in a specific term to be an essential feature of hybrid evaluatives. Having a specific descriptive counterpart or making use of a paraphrase instead makes no relevant theoretical difference, nor should it be taken as a parameter that distinguishes slurs and thick terms: It is true that most slurs derogate people on the basis of characteristics for which languages have a specific term (for example nationality and sexual orientation), but this is not always the case. Think about slurs like ‘foureyes’ for people wearing glasses, or ‘JAP’ for wealthy American

⁵² See Elstein and Hurka (2009: 521) for a discussion about the descriptive content of slurs and thick terms.

Jewish women, or ‘fairy’ for homosexual men that display features or behaviors that are taken to be stereotypically feminine, ‘rice queen’ for gay men particularly attracted to East Asian men, etc. These examples show that slurs can have complex descriptive meanings for which languages lack an exact non-loaded counterpart, *contra* Camp (2013: 330), who talks about a “marked contrast” between slurs and thick terms in that the former have a neutral counterpart and the latter do not. Scholars have diverging opinions on this point. For example, consider the following remark from Nunberg (forthcoming):

“DiFranco (2015) notes that a certain class of idiomatic expressions, such as slanteyed, curry muncher and Jewish American Princess, don’t have neutral equivalents, but should be considered slurs all the same. These fall, I think, into several classes (...), but the absence of a default synonym is likely due to the fact that, like wog and gook, it denotes what most consider an illegitimate social category – i.e., those who ‘look Oriental’ – and hence lacks a coreferential default equivalent.”

I agree with DiFranco’s claim that terms like ‘curry-muncher’ are slurs despite their complex descriptive content and disagree with Nunberg about the illegitimacy of the social category being a crucial feature. The problem is not whether a certain social category is legitimate, rather whether a certain linguistic community has had or not the interest (legitimate or not, it does not matter) to coin a term to call a certain group instantiating a certain property.⁵³ This is, in a way, accidental and it does not constitute a crucial fact in the linguistic analysis of the underlying mechanism through which certain terms are associated with certain evaluative contents. In the same way, it might be the case that the only ones interested in naming a certain feature (for instance ‘being disposed to share one’s resources without expectations of receiving anything in return’) have been the ones who wanted to *praise* such a property. The fact that thick terms do not have an exact descriptive counterpart does not make them crucially different from slurs.

⁵³ About the idea that the mere existence of a term signals the interests of a certain linguistic community, see Saka (2007).

However, we shall see in the next section that this feature of having an exact term as a neutral counterpart rather than a more complex expression – feature that I discard as superficial – does play a role in giving rise to some phenomenal divergences between slurs and thick terms.

As far as the *evaluative* content is concerned, the main differences concern (i) polarity – the conveyed evaluation being positive (P) or negative (N) – and (ii) ‘objectability’ – the conveyed evaluation being warranted (W) or unwarranted (U) –. As we saw, these two parameters are orthogonal and give rise to 4 combinations. While thick terms seem to admit all four combinations (say: PW ‘generous’, PU ‘chaste’, NW ‘brutal’, NU ‘lewd’), slurs tend to instantiate NU only, i.e. they typically convey negative and morally wrong evaluative content (even though, as we saw in section 1.1, it is discussed whether slurs can convey positive or warranted evaluations). Even though it might be sociologically interesting that slurs only instantiate one of the four combinations, it does not constitute a strong reason to treat slurs and thick terms differently, because, as we discussed at a greater length in section 1.1, it is conceivable for a language to have a positive slur, or a slur that conveys a correct morally right evaluation.

6.2 The projective behavior

The second aspect under which slurs and thick terms differ is their projective behavior. In general, it is much easier to neutralize the evaluative content of thick terms rather than slurs by embedding them, as the evaluative content of slurs displays a stronger projective power.

Take the following constructs considered in chapter 5, of the form ‘If π , then φ_π ’, where φ_π is a presupposing utterance and π is the content presupposed by φ_π ; the antecedent of the conditional contains the presupposition triggered by the consequent; the conditional is preceded by a clause where the speaker expresses agnosticism with respect to the presupposed content:

141. I don’t know whether John used to smoke in the past. But if he did, then we can truthfully say that he stopped.

142. I don't know whether France has a king. But if it does, then we can truthfully say that the king of France tends not to appear on TV.

143. I don't know whether it is bad in any distinctive way to be sexually explicit beyond conventional boundaries. But if it is, we can truthfully describe John's show as being lewd.

144. ? I don't know whether Italian people are despicable for being Italian. But if they are, then we can truthfully say that your boss is a wop.

For standard presupposition triggers like 'stop' or 'the', the presupposition does not project ((141)-(142)); for thick terms, the presupposed evaluation does not project either, as we see in (143). The case of slurs seems to diverge from standard presuppositions and thick terms. (144) sounds odd and it does not entirely prevent the racist content triggered by 'wop' from projecting.

The same goes for disjunctions of the form 'not π or φ_π ', where φ_π is a presupposing utterance and π is the content presupposed by φ_π :

145. Either Jane never smoked, or she stopped.

146. Either France does not have a king, or the king of France tends not to appear on TV.

147. Either there is nothing bad in things that are sexually explicit beyond conventional boundaries, or this show is lewd.

148. ? Either Italians are not despicable because of being Italian, or your boss is a wop.

The presupposition triggered by 'stop', 'the' and 'lewd' does not project, while the presupposition triggered by 'wop' is hard to neutralize, as the oddity of (148) shows.

Let us now consider reported speech. *Verba dicendi* are plugs that can in general prevent a presupposition from projecting. Consider the following utterances featuring a standard presupposition trigger ('stop'), a thick term ('generous') and a slur ('wop'):

149. Lucy told me that Mary stopped smoking, but I don't think that Mary used to smoke in the first place.

150. Lucy told me that she found the donation generous, but I don't think there is anything good in giving without expectation of receiving anything in return.

151. ? Lucy told me that she saw a wop at the charity ball, but I disagree that being Italian is a reason for despise.

Even though (150) is slightly odd, it is still harder to embed slurs under speech report than how it is for standard presuppositions and thick terms.⁵⁴

Let us consider some factors that can explain such divergences. For one thing (and I am thinking of the examples with conditional and disjunctions) a possible explanation is that not to know whether certain people are despicable on the basis of their nationality is enough to be racist: Acknowledging the *possibility* that Italians *might* be despicable for being Italian already reveals a racist attitude. The fact that slurs negatively evaluate people on the basis of features such as nationality or sexual orientation makes it unacceptable not only to use them but also to admit that the negative evaluation that they trigger might be in fact correct. People that are sensitive to the harm inflicted by slurs will perceive utterances like (148) and (151) as racist. Another factor to take into account in explaining why slurs display a stronger projective behavior than thick terms concerns the social repercussions of slurs. Väyrynen (2013:151) underlines that thick terms do not provoke the same emotional reactions that slurs typically cause:

⁵⁴ Väyrynen (2013: 152) points out that slurs and thick terms do not pattern the same under belief reports.

“Whatever the correct semantics of these words may be, certain social and historical facts, standards and attitudes give them such power to cause emotional harm that for anyone who isn’t part of the targeted group to utter them in any context would show a failure fully to grasp their toxic power”.

However, involving or not social tension does not necessarily signal a linguistic difference: Predicates such as ‘pedophile’ or ‘thief’ might provoke particular social reactions that other predicates such as ‘teacher’ or ‘waiter’ lack, but this does not suggest that these predicates have different *kinds* of semantics. The social effects of slurs play nevertheless a role: Because of the socially toxic power, slurs are *prohibited* terms. Competent speakers typically know that slurs are taboo words and they avoid employing them unless they are willing to activate the evaluative content. On the contrary, there usually is no social prohibition that forbids the use of thick terms.

The context dependence of thick terms contributes to the contrast in projective behavior. As a matter of fact, a lot of contextual parameters need to be fixed to determine the meaning of a thick term. The value of several variables need to be negotiated. This does not happen for slurs: When a non-racist speaker is confronted with a slur, then she already knows that her interlocutor is making assumptions that she does not under any circumstances share, whereas with thick terms things are less clear. For example, a speaker might find most uses of ‘lewd’ objectionable: For this speaker, movies or shows cannot be bad for transgressing conventional boundaries in their being sexual explicit. But she might find some uses of ‘lewd’ possibly correct: For instance, when ‘lewd’ is used with respect to the behavior of a professor at school, the speaker can deem an action bad *for being* sexually explicit beyond conventional boundaries. This would in part explain the stronger projective behavior of slurs: The mere fact that a speaker is employing a slur alerts the participants to the conversation; we do not expect the same thing to happen with thick terms. Moreover, as far as reports featuring slurs are concerned, they are ambiguous between two readings: One where the content of the pejorative is ascribed to the speaker, one where it is ascribed to the reported person. Speakers who are sensitive to the harm of

slurs avoid the risk of misinterpretation by employing dissociative strategies: They use scare quotes or expressions like ‘literally’, ‘exact words’, or paraphrase ‘Lucy told me that she saw an Italian person at the charity ball, except she actually used a racist slur to call this guy’.

Another case where slurs allegedly display a stronger projective power than thick terms is the a-moralist kind of utterances. Some scholars (Brink 1989, Eklund 2013: 169, Väyrynen 2013) deem felicitous to apply a thick term while signaling agnosticism about the relevant evaluation. According to these authors – and I shall disagree – it is felicitous to utter (152):

152. Whether or not it is bad to be sexually explicit beyond conventional boundaries, John’s show can truthfully be described as being lewd.

On the other hand, the same kind of utterance involving standard presupposition triggers and slurs is definitely infelicitous.

153. ? Whether or not John used to smoke in the past, we can truthfully say that he stopped.

154. ? Whether or not Italians are bad because of being Italian, John can truthfully be described as being a wop.

I am inclined to discard this kind of examples as infelicitous, but I have to explain why to certain people, such as the authors I mentioned, (152)-like utterances appear felicitous. The answer, that will be discussed in chapter 10, involves the availability of the neutral counterpart. The lack of an available exact non-loaded term makes speakers more charitable in making sense of an utterance like (152). Note moreover, that the prototypical example of this kind of defeasibility does not involve an objectionable thick term like ‘lewd’, but a positive and (possibly) non-objectionable one such as ‘considerate’ (‘Whether or not this is good in any distinctive way, John can be truthfully and neutrally described as being considerate’ Bergström 2002: 5, quoted in Väyrynen 2013: 84). This can make speakers fail to perceive the infelicity

of (152)-like constructions. As a matter of fact, it is easier for speakers to detect the evaluation associated with objectionable thick terms rather than that associated with non-objectionable ones. The more they disagree with the conveyed evaluation, the more they recognize its projection. This would also give a hint of why it has been claimed that the evaluation conveyed by slurs is systematic, whereas there is less consensus for thick terms and why the projection of the evaluative content of slurs is perceived as stronger than that of the evaluative content of thick terms.

Hare (1963) underlines the same thing:

“ (...) a person who is not inclined to despise people just because they are negroes⁵⁵ will reply that (...) he prefers not to use the word ‘nigger’. (...) He substitutes, let us suppose, the neutral word ‘negro’. (...) The only reason why such a course with the word ‘courageous’ would seem to us strange is that we are, most of us, very firmly wedded to the attitude which the word incapsulates. If there were a person who was not in the least disposed to commend those who preserved the safety of others by disregarding their own, then he could say, as before, ‘I prefer not to use the word ‘courageous’, just because it incapsulates this attitude to which I do not subscribe. I prefer the longer, morally neutral expression, ‘disregarding one’s own safety in order to preserve that of others’”. (Hare 1963: 188-189)

In this section I observed the differences between slurs and thick terms by focusing on their semantics and on their projective behavior. I showed how such differences do not constitute reasons to provide a different linguistic analysis of slurs and thick terms. As suggested in section 3.3.4, more interesting differences can be found between objectionable and non-objectionable hybrid evaluatives, i.e. between the evaluatives that convey correct evaluations and those that do not (like slurs and objectionable thick terms).

⁵⁵ Hare, who writes at the beginning of the Sixties, takes the term ‘negroes’ to be a descriptive non-derogatory term.

Part II. Rival theories

The present part is a critical review of some alternative theories that have been put forward to account for slurs and thick terms respectively. The two separate histories of these terms make it the case that the theories I consider often deal with only one class and neglect the other; nevertheless, the objections I raise against one kind of approach to slurs, for instance, will often apply to a corresponding analysis of thick terms.

The overview I provide in this part is far from covering *all* the theories on evaluatives, nor it treats the *main* theories on evaluatives. Rather, I selected three kinds of accounts for their relevance to specific aspects of my argument. In particular, I leave aside expressivist accounts of evaluatives and I do not assess the complicated debate between cognitivists and non-cognitivists.

The first four chapters are dedicated to the “single-source” approaches, as Jeshion (2016a) calls them: the truth-conditional approach (chapters 7-8), according to which the evaluative content of hybrid evaluatives is reducible to their truth conditions and the deflationary approach (chapters 9-10), according to which the evaluation associated with hybrid evaluatives is not part of their encoded meaning, but it is the result of pragmatic effects. I will show that the single-source approaches fail to properly account for the behavior of hybrid evaluatives. The difficulties met by these proposals speak in favor of a hybrid approach.

In chapter 11 I will present a hybrid approach alternative to mine, the conventional implicature view, and I will point out the aspects for which the presuppositional view should be preferred.

7. Truth-conditional theories of slurs: It’s just a matter of semantics⁵⁶

In this chapter I will present and discuss the truth-conditional analysis of slurs. The truth-conditional approach accounts for the intuition that the meaning of slurs crucially differs from the meaning of their neutral counterparts. Slurs and their counterparts are not synonymous, as the former derogate their targets, while the latter

⁵⁶ This chapter is largely based on Cepollaro and Thommen (ms), which is currently under review.

do not. The main difficulty for the truth-conditional account is to explain the projection of the derogatory content of slurs out of semantic embedding. In what follows, I will present the theory developed in Hom (2008, 2010, 2012) and Hom and May (2013, 2014, forthcoming). I will present the strategies that the authors propose to solve the projection problem and I will show why they are inadequate.

I claim that the truth-conditional approaches fails to properly account for the data involving hybrid evaluatives and should thus be rejected. However, as I will argue in chapters 9-10, I take the initial intuition to be correct: The difference between hybrid evaluatives and their neutral counterparts is not just a matter of pragmatic effects.

The bulk of the truth-conditional theory of slurs is the following: Slurs are derogatory because they ascribe negative properties to their targets. The simplest version of the approach would be to analyze a slur like ‘wop’ as ‘Italian and bad because of that’. Hom (2008, 2010, 2012) argued for a more complex version of the truth-conditional theory, introducing what he called “combinatorial externalism (CE)”. In his view, a slur means:

“Ought be subject to $p_1 + \dots + p_n$, because of being $d_1 + \dots + d_n$, all because of being NPC*, where $p_1 + \dots + p_n$ are deontic prescriptions derived from the set of racist practices, $d_1 + \dots + d_n$ are the negative properties derived from the racist ideology, and NPC* is the semantic value of the appropriate non-pejorative correlate of the epithet. For example, the epithet ‘chink’ expresses a complex, socially constructed property like: Ought to be subject to higher college admissions standards, and ought to be subject to exclusion from advancement to managerial positions, and ..., because of being slanty-eyed, and devious, and good-at-laundering, and ..., all because of being Chinese.”
Hom (2008: 431).

The externalist move consists in the fact that the properties ascribed to the target group (both $p_1 + \dots + p_n$ and $d_1 + \dots + d_n$) are established by experts in the community; in the case of ‘chink’, for example, the experts are the racists; for ‘faggot’, the experts are the homophobes and so on. The externalist component of Hom’s analysis was somehow left aside in Hom and May (2013, 2014, forthcoming). According to Hom

and May (2013, 2014, forthcoming), slurs are complex predicates consisting of a second order marker of pejoration PEJ that applies to a first order predicate referring to a group. For example, ‘kike’= PEJ (Jew). Slurs can be paraphrased as follows: “Ought to be the target of negative moral evaluation because of being a member of G, where G is the first order group⁵⁷ concept term” (Hom and May forthcoming).

However, I believe that both versions of the truth-conditional account (*à la* Hom, with an externalist semantics or *à la* Hom and May, with the pejorative operator) meet the same kind of problems. I will therefore focus on the latest version of the account.

As we said, Hom and May’s *semantic* claim is that a slur like ‘kike’ means something along the lines of “ought to be the target of negative moral evaluation because of being Jewish”. In addition to the semantic claim, they make a *moral* claim, according to which, necessarily, no one ought to be the target of negative moral evaluation because of belonging to a group. Note that not all scholars agree on the moral claim.⁵⁸ For example, for Predelli (2010):

“Notwithstanding the confused and unpleasant attitude apparently conveyed by uses of slurs of that sort, these uses may hardly qualify as linguistically defective. For one thing, racist and xenophobic attitudes are empirically incorrect: There is no conceptual (and, more importantly, no meaning-encoded) difficulty in supposing that membership in an ethnic or national group provides satisfactory motivation for a hostile attitude”

And again:

“It is a lamentable historical accident that bigots and racists have taken the lead in the production of ‘lexically encoded’ slurs. Still, exceptions abound: ‘Pigs’ may well have been an eminently usable derogatory term for typical police officers in the sixties (or at least for individuals belonging to certain repressive institutions). Moreover, the statistical prominence of

⁵⁷ For a discussion of what counts as a group according to Hom and May, see section 1.1.

⁵⁸ For a discussion of whether and when contempt can be well-focused, see Mason (2003).

racial, sexist, and xenophobic simple English slurs is easily compensated by the ‘expressive compositionality’ partially discussed above: ‘Damn fascist’ and ‘fucking racist’ are as subjectively expressive and linguistically non-defective as xenophobic slurs, but presumably eminently usable by at least some readers of this essay.” (*Ibidem*)

As discussed in sections 1.1 and 3.3, I leave room for the possibility that certain hybrid evaluatives might trigger an appropriate evaluation, that is, evaluations that are warranted or correct; this is particularly convincing for certain non-objectionable thick terms such as ‘courageous’, but not for prototypical examples of slurs. So I shall in a sense endorse Hom and May’s moral claim about slurs, but I do not maintain that it holds for *every* hybrid evaluative: I agree with Predelli that the fact that certain hybrid evaluatives (namely slurs) target people on the basis of incorrect assumptions is a “lamentable historical accident”, not an essential feature of hybrid evaluatives in general.

Going back to Hom and May’s proposal, the conjunction of the semantic claim and the moral claim has the following consequence: Slurs turn out to necessarily have an empty extension. In this sense, the two authors analyze slurs as fictional terms. The analogy between slurs and fiction allows to distinguish two types of truth: material and fictional truth. Consider the following pair:

1. Leonardo da Vinci was a wop.
2. Leonardo da Vinci was a limey.

Speaking about fictional entities, we can say that ‘Unicorns are white’ is *fictionally* true, but *materially* not true. In the very same way, Hom and May take (1) to be *fictionally* true (in the racist fiction of Italianophobia) and *materially* false, whereas (2) is both *fictionally* and *materially* false. According to the authors, the intuition that ‘Leonardo da Vinci is a wop’ is strictly speaking true ‘embeds a mistake of fiction for fact’ (Hom and May forthcoming). To illustrate the parallel, Hom and May consider the following case: In the Middle Ages, people took some real objects

– the tusks of narwhals – to be unicorn horns and they created a myth by ascribing them some magical properties (to neutralize poisons, heal diseases, etc.) that the real objects obviously lack. The same goes for slurs: Anti-Semites mistakenly take real individuals – Jewish people – to be ‘kikes’, by ascribing them some properties that they necessarily lack (like being contemptible because of being Jewish). That is to say, for Hom and May, anti-Semites are fiction-believers.

As previously mentioned, the main problem for the truth-conditional account is to explain the projection data presented in sections 2.1-2.2. Hom (2008) and Hom and May (2013, forthcoming) have two strategies to account for the data. The first one is to claim that there is in fact no real projection when slurs are embedded, for instance under negation. In particular, they take examples like ‘There are no kikes’ to convey no pejorative content. However, they recognize that sometimes embedded slurs do convey a pejorative content and they explain it by invoking an additional mechanism that they call “offense”. I shall consider both strategies and argue that they fail to account for the projection data. In particular, I will focus on two points put forward in Cepollaro and Thommen (ms): In section 7.1 I will argue that the non-pejorative readings of embedded slurs are actually metalinguistic and propose three tests for metalinguistic readings. In section 7.2, I will present and discuss the notion of offense and I will show that it does not explain the projection data either, by providing an example from fiction.

7.1. Denying some projection data

Let us acknowledge that it is possible for utterances like ‘There are no kikes’ to have a non-pejorative reading. The question is whether such a reading can prove that the pejorative content of slurs is truth-conditional. I maintain that it does not and I claim that the non-pejorative reading of embedded slurs is actually brought about by metalinguistic effects.⁵⁹

Consider the following statement:

⁵⁹ Note that the aim of this section is to investigate whether literal embedded uses of slurs can be non-pejorative; a separate issue is whether some slurs can have non-pejorative uses in general. An interesting example is the case of Appropriation (see Part III).

3. Lin is not happy.

There are (at least) two readings of (3), depending on how negation is interpreted. A *propositional* interpretation of negation (P) gives rise to a reading under which Lin is not in a particularly good mood; whereas a *metalinguistic* interpretation of negation (M) brings about a reading under which the speaker finds that the predicate ‘happy’ is not accurate to describe Lin’s state. The two readings can be made explicit with the following continuations:

4. Lin is not happy, she is quite sad. P⁶⁰
5. Lin is not happy, she is ecstatic. M

The phenomenon of metalinguistic negation has been intensely debated in the literature, but there exists no consensual theory of the phenomenon.⁶¹ I will not develop a general account of metalinguistic negation here, I will only show that non-offensive uses of slurs belong to this class of uses whose very existence is not debated. Recall that for Hom (2008, 2010, 2012) and Hom and May (2013, 2014, forthcoming) the non-offensive readings of the utterances in (6) show that truth-conditional operators can successfully affect the pejorative content of slurs, which is therefore truth-conditional in nature:

- 6.
- “Yao Ming is Chinese, but he’s not a chink.
- There are lots of Chinese people at Cal, but no chinks.
- Chinese people are not chinks.
- (...)

⁶⁰ ‘P’ indicates that a *propositional* reading is accessed; ‘M’ indicates that a *metalinguistic* reading is accessed; ‘#’ indicates that no reading is successfully accessed because the utterance is infelicitous.

⁶¹ About the variety of metalinguistic uses, from clear cases of reported speech to less clear cases of mixed or perspectival uses of expressions, see *i.a.* Horn (1985,1989) and Recanati (2001, 2007).

There are no chinks; racists are wrong.” (Hom 2008: 429)

The claim I shall defend in this section is that, *contra* Hom and May, the non-offensive readings of the above data are actually metalinguistic readings and crucially *not* propositional readings. Horn (1985) describes metalinguistic negation as follows:

“A device for objecting to a previous utterance on any grounds whatsoever - including its conventional or conversational implicata, its morphology, its style or register, or its phonetic realization” (Horn 1985: 121)

In other words, metalinguistic negation is not a device that would operate at the propositional level by reversing the truth value of a proposition (or selecting the complement of a predicate), but it would operate at the level of discourse representations. It is no surprise to find in Horn (1985) an example of metalinguistic negation involving embedded slurs:

“It is relevant that metalinguistic negation can be employed by a speaker who wishes to reject the bigoted or chauvinistic point of view embodied in an earlier statement within the discourse context: (c) I beg your pardon: Lee isn’t an ‘uppity {nigger/broad/kike/wop/...}’ - (s)he’s a strong, vibrant {black/woman/Jew/Italian/...}. (...) ‘I’m not ‘colored’ - I’m black!’, ‘I’m not a ‘gentleman of the Israelite persuasion’ - I’m a Jew!’” (Horn 1985: 133).

Many authors have suggested that the non-pejorative readings of the utterances in (6) do not speak in favor of the truth-conditional analysis of slurs, as they actually are instances of metalinguistic negation (see Anderson and Lepore 2013a: 28-29; Jeshion 2013a: 253, Bolinger 2015). However, they suggest such an alternative explanation “by analogy”, without providing arguments. The aim of the present section is not only to argue that a metalinguistic explanation of the non-pejorative reading of the utterances in (6) is possible, but that it is in fact correct. I show that the

non-projective readings of the utterances in (6) are the result of a metalinguistic interpretation of negation, by considering linguistic constructions in which metalinguistic negation is *not* available. I apply three tests put forward in Cepollaro and Thommen (ms): the prefixal incorporation of negation, the ‘it’s false that’ construction and the ‘without being’ construction.⁶²

(a) *Non-F*

Horn (Horn 1985: 140) observes that when one tries to incorporate negation prefixally, metalinguistic readings are blocked. Observe what happens with ‘Lin is not happy’. As we said, it is ambiguous between a propositional and a metalinguistic reading. In what follows, the material in parenthesis is meant to help the reader accessing the relevant readings.

- 7.
- | | |
|---|---------------------|
| a. Lin is not happy. | AMBIGUOUS (P vs. M) |
| b. Lin is not happy. (She is sad) | P |
| c. Lin is unhappy. (She is sad) | P |
| d. Lin is not happy. (She is ecstatic) | M |
| e. ?? Lin is unhappy. (She is ecstatic) | # |

The test shows that when negation is incorporated (7-c and 7-e), the only reading available is the propositional one (7-c).

Now consider the same test with a slur⁶³ instead of the predicate ‘happy’:

⁶² Horn proposes a similar test having to do with the distribution of polarity items (Horn 1989: 370, 374, 396). According to Horn, metalinguistic negation should not license negative polarity items (e.g. ‘any’), nor inhibit positive polarity items (e.g. ‘some’). However, I left this test aside, because of the controversial and conflicting results that it produces (see Geurts 1998: 278).

⁶³ Note that Eklund (2011) suggests a metalinguistic interpretation of negated thick terms and Väyrynen (2013) applies the incorporation test to negated thick terms.

8.

- | | |
|---|---------------------|
| a. Lin is not a chink. ⁶⁴ | AMBIGUOUS (P vs. M) |
| b. Lin is not a chink. (She is Russian) | P |
| c. Lin is a non-chink. ⁶⁵ (She is Russian) | P |
| d. Lin is not a chink. (She is Chinese) | M |
| e. ??Lin is a non-chink. (She is Chinese) | # |

Once more, (8)-a is ambiguous between a propositional and a metalinguistic reading of negation. Depending on how speakers interpret negation, the utterance is perceived as conveying a Asianophobic content or not. The prefixal incorporation of negation in (8)-c encourages the internal reading of negation, under which the slur conveys a pejorative content towards Chinese people. On the other hand, if we encourage a metalinguistic interpretation of negation when negation is incorporated,

⁶⁴ Note also that the comparison of noun slur ('chink') with an adjective ('happy') does not harm the point, as underlined in Cepollaro and Thommen (ms). The results replicate with adjectival slurs like 'queer' (but the well-known appropriative use of 'queer' introduces unneeded complexity):

- | | |
|--|--------------------|
| a. John is not queer. | AMBIGUOUS (P vs M) |
| b. John is not queer. (He is heterosexual) | P |
| c. John is non-queer. (He is heterosexual) | P |
| d. John is not queer. (He is homosexual) | M |
| e. ??John is non-queer. (He is homosexual) | # |

⁶⁵ English does not have a slur with an incorporated negation, but it is harmless to use a neologism in this case, as what matters is the contrast between the c- and the e-conditions. Moreover, imagine a situation in which speakers divide the world between people who are Chinese and people who are not. They could thus apply the predicates 'Chinese' and 'Non-Chinese'. Racists would similarly apply the pejorative predicates 'Chink' and 'Non-Chink'.

the result is deviant, as (8)-e shows. Overall, this is evidence that the non-pejorative reading of (8)-a is the result of a metalinguistic effect, rather than truth-conditional computations.

(b) *It's false that*

Constructions like 'it's false that' also block metalinguistic readings of negation (Horn 1989: 416):

- 9.
- | | | |
|----|--|---|
| a. | It's false that Lin is happy. (She is sad) | P |
| b. | ?It's false that Lin is happy. (She is ecstatic) | # |

From the ambiguous 'Lin is not happy', only the propositional reading survives when negation takes the above form, just like when it is incorporated. Now compare:

- 10.
- | | | |
|----|---|---|
| a. | It's false that Lin is a chink. (She is Russian) | P |
| b. | ?It's false that Lin is a chink. (She is Chinese) | # |

Again, an utterance of 'It's false that Lin is a chink' is interpretable only if the intended reading is propositional like in (10)-a; it is not felicitous with a metalinguistic reading, like in (10)-b. Just like above, from the ambiguous 'Lin is not a chink', the metalinguistic reading does not survive when negation takes the explicit propositional form. This is new evidence that non-pejorative readings of embedded slurs are not the result of a truth-conditional interaction, as maintained by Horn and May.

(c) Without being⁶⁶

‘Being F without being G’ is just another way of saying ‘being F and not G’. So when one says that ‘Lin is not happy, she is F’, one could as well say that ‘Lin is F without being happy’. This construction of negation also seems to rule out metalinguistic readings. Take the following pair:

11.

a. Lin is in a quiet mood, she is not happy. P

b. Lin is ecstatic, she is not happy. M

And put it under the ‘without being’ form:

12.

a. Lin is in a quiet mood, without being happy. P

b. ?? Lin is ecstatic, without being happy. #

Again, slurs pattern the same. Take the following slurring sentences:

13.

a. Lin works in China, she is not a chink. P

b. Lin is Chinese, she is not a chink. M

⁶⁶ Benjamin Spector (p.c.) suggested such a test, for which he credits Danny Fox.

And put them in the ‘without being’ form:

14.

a. Lin works in China, without being a chink.

P

b. ??Lin is Chinese, without being a chink.⁶⁷

#

Once more, an utterance of a slur under negation of the form ‘without being’ is interpretable only if the intended reading is propositional ((14)-a, which is racist); it is not interpretable under a metalinguistic reading ((14)-b). In other words, the metalinguistic reading of (13)-b does not survive the ‘without being’-transformation.

The three tests just presented show that non-offensive uses of embedded slurs pattern with metalinguistic uses of negation. Thus, data like the utterances in (6) do not show that the pejorative content of slurs is propositional, as it is not cancelled by propositional negation but rather by *non-propositional* metalinguistic negation.

Before moving to next section, let me mention some experimental work that recently investigated the way in which speakers perceive the pejorativeness of embedded slurs. Panzeri and Carrus (2016) asked participants to rate on a 7-points scale the offensiveness of a person who utters a sentence that contains a slur embedded under negation, antecedent of a conditional and question and compared it to the perceived offensiveness of unembedded slurs rated on a 7-points scale. They found that slurs are perceived as offensive when they are embedded in a question or in the antecedent of a conditional; however, they also found that, quite surprisingly, when the slur is embedded in negation it is not perceived as offensive as under the other kinds of embedding. The question at stake is what makes the negation different from the other types of embeddings. Panzeri and Carrus hypothesized that participants might access a *metalinguistic* interpretation of negation. To test such an hypothesis, they asked participants to provide a continuation of the utterances featuring negated slurs. They found two main patterns of continuation: metalinguistic ones and non-metalinguistic ones. For instance, for the sentence ‘Leo is not a

⁶⁷ I find the first two tests slightly more convincing than the ‘without being’ one. However, most speakers that have been asked, found (14)-b infelicitous.

faggot’, a non-metalinguistic continuation would be ‘He has a girlfriend’, since the negation would be interpreted as affecting the descriptive content (that is, that Leo is homosexual); on the other hand, a metalinguistic continuation would be ‘I don’t like these expressions’, where the negation is taken to affect the *use* of the slur rather than its descriptive content. In observing the ways in which participants continued the utterances, they found a high percentage of metalinguistic-like continuations, such as ‘I don’t like these expressions’. Panzeri and Carrus conclude that (i) there is a contrast in the ways in which speakers perceive the pejorative power of slurs embedded under negation on the one hand and slurs embedded under conditional and question on the other hand; (ii) such a contrast might be due to a metalinguistic interpretation of negation, revealed by the continuations provided by the participants. Panzeri and Carrus (2016)’s results provide further support for the claim that the data taken by Hom and May in support of a truth-conditional analysis of slurs are actually instances of metalinguistic uses of language.

7.2 Explaining *some* projection data: derogation and offense

As we saw above, according to the truth-conditional account of slurs, slurs under negation (and other operators) are not derogatory, in the sense that they do not *predicate* negative properties of a subject. However, Hom and May acknowledge that utterances with embedded slurs like (15)-(17) *sound* pejorative:

15. Madonna is not a chink.

16. There are no wops in my neighbor.

17. If Peter is a fag, then he is worthy of contempt.

Hom (2012) and Hom and May (2013) distinguish the phenomenon of ‘derogation’ – predication of a negative moral property based on the subjects belonging to a group – from another phenomenon, somehow similar in its effects. In addition to being derogatory, the use of slurs would generate what they call ‘offense’. As Hom puts it,

(...) [Derogation] is an objective feature of the semantic contents of pejorative terms. Derogation is the result of the actual predication, or application, of a slur or pejorative term to its intended target group. (...) [Offense] is a subjective effect of the semantic contents of pejorative terms in a context. Offense is a psychological result on the part of the discourse participants, and is a function of their beliefs and values. (Hom 2012: 397).

There are thus two possible factors responsible for the negative effects that slurs elicit: a truth-conditional component on the one-hand (derogation) and a psychological and pragmatically driven component on the other hand (offense). Among the critics of such a distinction, Jeshion (2013b) notices that the derogation/offense distinction makes the surprising prediction that ‘John is a nigger’ and ‘Is John a nigger?’ are disparaging for different reasons. Furthermore, it predicts that ‘John is not a chink, he is a nigger’ is derogatory (in the technical sense) towards African-Americans, but not toward Asian people. As we underline in Cepollaro and Thommen (ms), given these two parameters, an utterance U featuring a slur S in a context C can stand in four possible states depending on the linguistic environment and the conversational context, as illustrated in Figure 1. Whenever negative moral properties are predicated of a subject (like in U₁), there is *derogation* (cases a and b); on the other hand, *offense* is brought about only in contexts of utterance where a ‘psychological result’ is expected, i.e., when it is problematic in the context of utterance to suggest that the members of the target class are despicable (cases a and c); otherwise – for example in a deeply homophobic society where everyone, including the members of the target class, is homophobic – offense does not arise (cases b and d). The counterintuitive predictions of the derogation/offense distinction are that a question like U₂ ‘Were the last six Roman emperors fags?’ *does not* carry strictly speaking any derogatory content towards homosexual people and it is ‘offensive’ only if it is uttered in a non-completely homophobic scenario. In the hyperbolic-homophobic scenario C₂, the question U₂ is neither derogatory nor offensive (case d). Similarly, the assertion U₁ is derogatory, but not offensive (case

b).

Figure 1. Hom and May's derogation/offense distinction (based on Cepollaro and Thommen ms).

	Offense +	Offense -
Derogation	a) U ₁ : 'The last six Roman emperors were fags.' C ₁ : A journalist says that on TV today in Europe.	b) U ₁ : 'The last six Roman emperors were fags.' C ₂ : A journalist says that on TV in a deeply homophobic society where everyone - including members of the target class - is homophobic.
No Derogation	c) U ₂ : 'Were the last six Roman emperors fags?' C ₁ : A journalist says that on TV today in Europe.	d) U ₂ : 'Were the last six Roman emperors fags?' C ₂ : A journalist says that on TV in a deeply homophobic society where everyone - including members of the target class - is homophobic.

Hom and May provide another, finer grained, analysis of the *causes* of offense:

"Offensiveness can be linguistically triggered, because when speakers use predicates, they typically conversationally implicate their commitment to the non-null extensionality of the predicate". (Hom and May 2013: 310)

So the idea is the following: Speakers tend to use terms that they take to have a non-empty extension; this is true of any predicate ('bikes', 'trees', etc.). For example, if John asks Mary whether she ever speaks to angels, he will conversationally implicate that he believes that angels exist; Mary and bystanders will generate such an implicature.⁶⁸ Likewise, if John asks Mary whether she ever speaks to 'wops',

⁶⁸ Conversational implicatures are expected to be cancellable. Hom and May rely on the non-pejorative readings of utterances like those listed in (6) to argue in favor of the cancellability of the conversational implicature responsible for offense. 'There are no kikes, kikes don't exist', 'John is not a kike because there is no such thing' or 'No Jews are kikes'. We have seen in section 7.1 that these cases are better understood in terms of metalinguistic effects.

Mary and bystanders will tend to infer that John believes that ‘wops’ exist (i.e., in Hom and May’s analysis, that there are people who are bad because of being Italian). Any use of slurs, embedded or not, triggers an implicature of non-null extensionality, which is offensive, given the alleged meaning of slurs.

However, as we claimed in Cepollaro and Thommen (ms), non-vacuity inferences do not satisfactorily explain the offensiveness of embedded slurs. Take a construction where non-vacuity inferences are typically blocked, for example ‘there is no F’: People do not infer that the speaker believes in the existence of God from her utterance of ‘there is no God’. The same goes with ‘there are no vampires’. Now, note that although vampires do not exist, there are slurs for vampires (e.g. ‘fangs’), just like for other fictional entities⁶⁹ (‘pointyear’ for elves, ‘toaster’ for robots, ‘furface’ or ‘moondog’ for werewolves). Imagine that John is afraid of being bitten by a vampire on his way home and Mary wants to reassure him. She could utter (18) or (19):

18. Don’t worry! There are no vampires! They don’t exist.

19. Don’t worry! There are no fangs! They don’t exist.

Although neither (18) nor (19) trigger existential inferences because of the ‘There is no F’ form, the utterance of (19) still carries Mary’s negative evaluation of vampires. The negative evaluative content about vampires in (19) cannot be the result of non-emptiness inferences. Therefore, non-null-extensionality implicatures do not explain the projection of the evaluative content of slurs.⁷⁰

Wrapping up, in this chapter I presented and discussed the truth-conditional account of slurs. The main difficulty faced by this approach is to account for projection data, i.e. for the fact that embedded slurs are typically disparaging. I

⁶⁹ The ‘fangs’ example comes from the HBO tv-series True Blood; among gamers, there are slurs for all sort fictional entities (tvtropes.org/pmwiki/pmwiki.php/Main/FantasticSlurs).

⁷⁰ One could instead think that what triggers offense is not the inference of non-emptiness *per se*, but the inference of possible non-emptiness. ‘Wop’ triggers offense because it suggests that the speaker believes that it is *possible* that Italians ought to be the target of negative moral evaluation because of being Italian. This variation does not work either, as Mary could reassure John by saying ‘Don’t worry, there could be no fangs! It is simply impossible that fangs exist’ and this would not delete the pejorative content towards the fictional entities.

presented two kinds of alleged solutions to the problem put forward by Hom (2008, 2010, 2012) and Hom and May (2013, 2014, forthcoming): To take the non-offensive readings of embedded slurs as evidence of the truth-conditional nature of the pejorative content (7.1) and to introduce an additional pragmatic mechanism that is responsible for the perceived offensiveness of slurs (7.2). I showed that neither of these strategies is apt to solve the projection problem: The first one fails because the non-pejorative readings of certain negated slurs are elicited by metalinguistic readings of negation; the second one fails because the derogation/offense distinction provides unsatisfactory predictions if offense is interpreted as a psychological reaction caused by speakers' beliefs, while it provides wrong predictions when offense is interpreted as a psychological reaction caused by non-null extensionality implicatures.

I take the data discussed in this chapter to show that the derogatory content of slurs is not truth-conditional. Note moreover that rejecting the truth-conditional theory does not automatically mean to endorse the so-called identity thesis, according to which slurs and neutral counterparts are co-referential, as in order to establish reference, other dimensions of the linguistic machinery have to be taken into account, as discussed in section 3.3.

In the next chapter, I present and discuss the truth-conditional analysis of thick terms, put forward by Kyle (2013).

8. Kyle's semantic analysis of thick terms

This chapter is dedicated to the truth-conditional analysis of thick terms and I discuss in particular Kyle's approach. Kyle proposes a semantic approach for thick terms that resembles in some respects Hom and May's theory of slurs. He is neutral with respect to two questions concerning derogatory epithets: (i) whether slurs can be analyzed as thick terms and (ii) whether, if (i) is the case, his semantic account would apply to slurs. As we shall see, though, some of the difficulties encountered by Kyle's Semantic View on thick terms are similar to the problems met by Hom and May's truth-conditional account of slurs.

According to Kyle, the relation between a thick term and its associated evaluation is conceptual entailment: "The sentence 'A is generous' conceptually entails that 'A is

good in a way' is true relative to select contexts." (Kyle 2013: 6). Kyle's main point is that – *contra* Brower (1988), Blackburn (1992), Väyrynen (2009, 2013) – thick terms have evaluative meanings: He calls such a view "Semantic View". The account that I propose is in a way a version of such a semantic view, as I do claim that the lexical meaning of thick terms has an evaluative content. However, I disagree with Kyle on how such a semantic relation is best spelt out. Kyle claims that "One way to formulate the Semantic View is to say that sentences containing thick terms have evaluative truth-conditions." and that "(...) It makes little difference which of these formulations we choose" (Kyle 2013: 2). I shall argue that it makes a great difference *how* exactly one spells out the semantic view and that the truth-conditional analysis that Kyle puts forward fails to account for projection data. However, as we have already seen in Part I, I do embrace Kyle's main claim that the evaluation conveyed by thick terms belongs to their lexical meaning.

For Kyle (2013:3), (20) commits the speaker to the truth of (21) because the latter is part of the meaning of the former:

20. Nancy is generous.

21. Nancy is good.

An interesting aspect of Kyle's account is the stress on the context dependence of the thin terms 'good' and 'bad' (recall the discussion in sections 3.1.2 and 3.1.3 on how 'good' and 'bad' should be read as 'good/bad in a way *W*'). Kyle correctly underlines that 'good' and 'bad' cannot be taken *simpliciter* (Kyle 2013: 6). Take (20): According to Kyle, 'generous' contains the predicate 'good'. But not *any* logically possible interpretation of 'good'. 'Good' could in fact be interpreted in many trivial or non-relevant ways: Something can be good as an example not to imitate, good as an instantiation of how to misbehave, etc. All these examples are not relevant to the interpretation of 'generous'. According to Kyle, every time 'generous' is used, the context provides a set of domains that contain all and only the *ways* of 'being good' that are relevant in the context. It follows that for Kyle 'generous' does not exactly entail 'good', but 'good *in a way*'.

Moreover, Kyle takes (22)-(23) to be positive evidence in support of the truth-conditional view (like before, we use ‘*’ to signal redundancy and ‘??’ to signal infelicity):

22. * Nancy is generous and she’s good in a way.

23. ?? Nancy is generous and she’s not good in any way.

According to Kyle, (22) sounds redundant because the second conjunct is part of the meaning of the first one; for the same reason, (23) is strongly infelicitous. Before discussing (22)-(23), an apparently terminological remark: Kyle (2013) seems to use ‘semantic’ as a synonym for truth-conditional. This approximation is dangerous given the subtlety of the differences at stake. As a matter of fact, the relation between (20) and (21) may well be ‘semantic’ in a broad sense (and (20)-(21) can be rightly taken to support a ‘semantic’ view), but not necessarily a *truth-conditional* view. As a matter of fact, the data taken by Kyle as evidence in support of the truth-conditional view also supports my presuppositional view. Consider the behavior of standard presupposition triggers in similar constructs (take a factive like ‘stop’):

24. * Nancy stopped smoking and she used to smoke.

25. ?? Nancy stopped smoking and she never smoked.

The presuppositional theory correctly predicts both the redundancy of (22) and (24) and the strong infelicity of (23) and (25). Therefore, Kyle needs to rely on some other kind of evidence in order to defend a truth-conditional account rather than a presuppositional one. I will not go back to Kyle’s objections to the presuppositional account of thick terms, that are discussed in chapter 5.

As I mentioned at the beginning of this chapter, some of the difficulties that Kyle’s account encounters are related to the issues discussed in the previous chapter about slurs, as they concern the projection of the evaluative content of thick terms. If the evaluative content of thick terms belongs to their truth-conditional meaning as

Kyle is suggesting, then it would not survive under semantic embedding such as negation, antecedent of conditional, modal and question. In the previous chapter, we saw that Hom and May adopt a twofold strategy for slurs: They (i) support the availability of non-evaluative readings of embedded slurs (section 7.1) and (ii) appeal to pragmatic mechanisms (“offense”) to explain the perceived projection (section 7.2) of the evaluative content. Kyle’s strategy to defend the truth-conditional account of thick terms resembles in a way the latter move: He proposes to analyze embedded occurrences of thick terms as *pragmatically* implicating evaluative content rather than literally expressing it. In particular, he appeals to two kinds of conversational implicatures, namely negative strengthening and clausal implicature. Take ‘discreet’, a thick term that means something like ‘tending to avoid attracting attention and to keep secrets’ and conveys a positive evaluation. Consider an embedded occurrence of ‘discreet’ such as the following:

26. Nancy is not discreet.⁷¹

Väyrynen (2009) holds that utterances like (26) show that the evaluative content of thick terms survives under negation, as, he claims, a ‘discreet’-objector would be reluctant to utter (26). For Väyrynen, the fact that a speaker who does not endorse the evaluation conveyed by ‘discreet’ would not utter (26) is evidence that the evaluative content of ‘discreet’ survives the embedding under negation in (26). Therefore, it is not part of the truth-conditional contribution offered by ‘discreet’. Kyle, on the other hand, offers an alternative explanation for the projection intuitions concerning (26) that is compatible – unlike Väyrynen’s approach – with the truth-conditional account. Kyle acknowledges that ‘discreet’-objectors are reluctant to utter (26). For Kyle, the reason however is not that the evaluation associated with ‘discreet’ survives under negation *qua* non-truth-conditional content (as Väyrynen claims), but that (26) typically implicates (27). This phenomenon, called ‘negative strengthening’, is discussed by Levinson (2000: 127 ff.): When a speaker says ‘Nancy is not happy’,

⁷¹ The original example in Väyrynen (2009) and Kyle (2013) involves ‘chaste’; however, the phenomenon of negative strengthening concerns gradable adjectives and predicates. I believe that Kyle’s examples are flawed by the appeal to negative strengthening in the case of a (possibly) non-gradable term such as ‘chaste’. For this reason, I choose ‘discreet’ as a working term. This will require the reader to put herself in the shoes of a ‘discreet’-objector.

she conversationally implicates something stronger than what she actually said, i.e. ‘Nancy is unhappy’. According to Kyle, if a speaker utters (26), then she would be implicating (27). He also assumes that it is plausible that a ‘discreet’-objector, reluctant to utter (26), would be equally reluctant to utter (27). Therefore, for Kyle, because (26) conversationally implicates (27) and the ‘discreet’-objector would *not* be happy with uttering (27), she is reluctant to utter (26) in the first place.

26. Nancy is not discreet.

27. Nancy is indiscreet.

My strategy to reject Kyle’s argument focuses on the constructs that are able to block negative strengthening. Since negative strengthening is a type of conversational implicature, it is cancellable. One way to block the implicature is to explicitly deny it. If we find that, when negative strengthening is blocked, embedded thick terms still convey their associated evaluation, then Kyle’s argument does not satisfactorily explain projection under negation. Consider (28):

28. Nancy is not discreet, but neither is she indiscreet.

In (28) there is no negative strengthening: The first conjunct, ‘Nancy is not discreet’, clearly does not imply ‘Nancy is indiscreet’, since that is explicitly denied in the second conjunct. Nevertheless, the occurrence in (28) of the thick terms ‘discreet’ and ‘indiscreet’ conveys their standard evaluative content: The embedded occurrences of ‘discreet’ and ‘indiscreet’ convey their standard evaluative content *despite* the fact that negative strengthening is blocked. Therefore, *contra* Kyle, I argue that negative strengthening does not explain the intuition that the evaluation associated with ‘discreet’ survives under negation.

My account of hybrid evaluatives, on the other hand, can easily explain why a ‘discreet’-objector would be reluctant to utter utterances like (28), namely because (28) presupposes an evaluation that the speaker does not endorse, since the negation only affects the *descriptive* meaning of ‘discreet’. We interpret (28) as saying that

Nancy is somehow halfway between tending to attract attention and giving away secrets (which is assumed to be bad) and tending not to attract attention and keeping secrets (which is assumed to be good).

Moreover, in order to explain the perceived projection of the evaluative content of thick terms in the case of conditionals, Kyle appeals to another kind of conversational implicature, namely “clausal implicature” (Gazdar 1979: 59-62; Levinson 2000: 108-109).⁷² The idea is that when a speaker utters a conditional rather than a conjunction, she is conversationally conveying epistemic uncertainty about the truth of the antecedent. That is, the utterance of (29) implicates (30).

29. If abstaining from extramarital sex is chaste, then so is refraining from desiring extramarital sex.

30. Abstinence from extra-marital sex may be chaste, or it may not.

The appeal to clausal implicatures is not fully convincing either. For one thing, the same line of reasoning should be also applicable to the ‘chaste’-fan. Kyle claims that the ‘chaste’-objector is reluctant to utter (29) because it typically implicates (30); since the ‘chaste’-objector does not believe that abstinence from extra-marital sex *may* be ‘chaste’, she is reluctant to utter (29) in the first place. If Kyle’s line of argument is correct, then we would also expect that a ‘chaste’-user would be reluctant to utter (29), as it implicates that abstinence from extra-marital sex *may not* be ‘chaste’, implicature that the ‘chaste’-user would not endorse. Moreover, clausal implicatures (just like typical conversational implicatures) can be blocked by denying that the speaker is uncertain whether the predicate ‘chaste’ can be applied to the subject. As before, if we find that the evaluation associated with thick terms survives in the antecedent of conditionals when the clausal implicature is blocked, then clausal implicatures cannot explain the intuition of projection in (29).

⁷² Kyle would plausibly rely on clausal implicatures also to explain the projection under questions and modals.

31. I am sure that Nancy is not chaste, but if she is chaste, then she is not having extramarital sex.

The standard positive evaluation associated with ‘chaste’ seems to be conveyed by (31). It is even more clear when a gradable thick term is employed:

32. Nancy is not discreet (nor indiscreet either). But if she is discreet, you can be sure she will not tell anyone your secret.

In (32), the first conjunct (‘Nancy is not discreet’) blocks the clausal implicatures, while the parenthetical (‘nor indiscreet either’) blocks the negative strengthening to the effect that both the negative strengthening and the clausal implicatures are blocked at the same time and yet ‘discreet’ still conveys its associated positive evaluation. This is strong evidence that the intuitions of projection – the idea that thick terms convey evaluative content also when they are embedded – cannot be explained in terms of conversational implicatures and related mechanisms.

All in all, the conversational mechanisms invoked by Kyle such as negative strengthening and clausal implicatures cannot be responsible for the evaluative content conveyed by embedded thick terms. We have seen that such conversational implicatures, if present, can be blocked like in (28) and (32) and yet, the evaluative content survives.

8.1 Repeatability and redundancy

As we saw, Kyle relies on data concerning redundancy in order to support his truth-conditional analysis of thick terms over competing approaches. However, he neglects some important factors. First, expressive content has – in general – some peculiar features, such as the repeatability (see Potts 2005). One can say ‘The damn engine of the damn car broke down in the middle of the damn street’ without redundancy: The repetition of expressive content strengthens the expressive effect rather than being redundant. In the case of thick terms, we find that repeating the evaluative content sounds slightly less redundant than repeating the descriptive one. We therefore expect (34) to be slightly more redundant than (33).

33. Nancy is generous and she is good in that respect.

34. *Nancy is generous and she is willing to share her resources.

Moreover, Kyle uses the alleged oddity of (33) to discard the conventional implicature explanation of thick terms, but we shall briefly see in chapter 10 that there are more problematic data against the CI analysis of thick terms. Moreover, the redundancy issues raised by Kyle do not discard the CI strategy, for at least two reasons: (i) first, not all conventional implicatures are reinforceable: Some are, some are not and *if* one finds that (33) is odd (and I do not agree on that), one could just conclude that thick terms are the kind of non-reinforceable conventional implicature and (ii) it is not entirely clear that a sentence like (33) or its variant ‘Nancy is generous and she is good in *this* way’ are infelicitous or redundant. Finally, as mentioned in section 5.1, intuitions of redundancy can be shaky because different factors are at play (prosody, intonation, stress, contextual relevance of certain aspects, focus of the utterance, etc.). I shall provide sounder arguments against the CI approach in the next chapter.

To conclude, in chapters 7-8 I discussed the truth-conditional analysis of slurs and thick terms respectively, by discussing the proposals of Hom (2008, 2010, 2012) and Hom and May (2013, 2014, forthcoming) on the one hand, and of Kyle (2013) on the other. The main problem with the truth-conditional approach is to account for the projection of the evaluative content. I have shown that the strategies employed to explain projection without appealing to non-truth-conditional meaning fail to explain the phenomena. In particular, I (i) showed that Hom and May’s data in support of a non-pejorative readings of embedded slurs actually stem from metalinguistic readings and (ii) showed that pragmatic mechanisms (both Hom and May’s *offense* and Kyle’s conversational implicatures) cannot explain projection. Therefore, the truth-conditional analysis of hybrid evaluatives fails to explain basic phenomena of how these terms work and need to be abandoned. However, as I shall discuss in the next chapter, I endorse the following claim of the truth-conditional accounts of

hybrid evaluatives: The evaluative content is linguistically encoded in the meanings of such terms and it does not arise only on the basis of pragmatic mechanisms.

9. Deflationary accounts of slurs: It's just a matter of pragmatics

The peculiar embedding pattern of slurs, that we observed in Part I, typically motivates a 'thick semantics', i.e. a semantics that accounts for the pejorative content associated with slurs, however this can be analyzed (truth conditions, conventional implicatures, presuppositions). Nevertheless, some authors challenge the very idea that a thick semantics is needed. In this chapter, I will present and discuss the so-called deflationary accounts of slurs (following the label that Anderson and Lepore (2013a) gave to their own approach), according to which slurs have a fairly unexceptional semantics. I will include in this category all those accounts according to which the derogatory content associated with slurs is *not* part of their encoded meaning (again, broadly intended: not just truth conditions, but also conventional implicatures, presuppositions and the like). I will focus on three main approaches⁷³ developed by Anderson and Lepore (2013a, 2013b), Bolinger (2015) and Nunberg (forthcoming), focusing on three notions respectively: taboo, contrastive preference and manner implicatures.

9.1 Anderson and Lepore: violating prohibitions and taboos

Anderson and Lepore (2013a, 2013b) defend a non-content-based view, according to which the reason why slurs display an offensive power is not because of *any* content they get across. The offensiveness of derogatory epithets has to do with certain edicts and prohibitions about language that are meaning-independent. Because slurs are prohibited terms, using or mentioning them constitutes a violation of a restriction and such a violation is responsible for the perceived offense. They say that "When a word is prohibited, then whoever violates its prohibition risks offending those who respect it" (Anderson and Lepore 2013a: 38). Anderson and Lepore do not draw a strong divide between mention and use: For the restriction to be violated it suffices for the

⁷³ I will not focus on Rappaport (ms), that analyzes the pejorative content of slurs in terms of the relevance-theoretic notion of 'showing': By using a slur, speakers *show* (as opposed to *mean*) that they share a negative attitude towards the target group. I believe however that most of the objections that I raise for Bolinger (2015) could also apply to Rappaport's account.

slur to *occur* (Anderson and Lepore, 2013b: 353). If the mere occurrence of a slur warrants offense, then the effects elicited by slurs cannot be ‘sealed off’ by scare quotes. In this sense, they expect mention, not just use, to be responsible for triggering offense.

In criticizing Anderson and Lepore’s view I will focus on the explanation/explanandum order and the notion of offense and I will propose two remarks concerning the notion of taboo and what I call the imagine-to-be-overheard argument.

The main problem for Anderson and Lepore (2013a, 2013b) has to do with the explanation/explanandum order. As we said, in their view a term is offensive *because* it is forbidden, even though the most intuitive thing to say seems that the reason why slurs are prohibited in many contexts in many societies is that they are offensive. Anderson and Lepore (2013b: 355) tried to assess such a worry by underlining that there can be a taboo on a word even if there is nothing bad in the expression itself: In some cultures, they argue, there is a taboo on expressions which are not offensive *per se*. The pronunciation of the tetragrammaton (i.e., YHWH, a name of God) was forbidden in Jewish culture, but this does not mean that there is anything offensive in the name of God, it is the violation of the norm to be offensive. Nevertheless, this comparison does not really solve the explanation/explanandum problem, as within religion there can be found a reason why the name of God should not be pronounced besides the fact that it is prohibited (one should not invoke God when it is not to praise him: That would be ‘saying his name in vain’). In the deflationary account of slurs, there seems to be no meaning-independent reasons for the prohibition of slurs but their offensiveness; however, the offensiveness of slurs cannot be the *cause* of the prohibition, as it is the *result* of the violation of this prohibition. Anderson and Lepore (2013b: 351) try to provide such a meaning-independent reason for prohibiting slurs by suggesting that a slurring term S that targets a certain group G is typically not introduced by G itself, but by some other group. According to them, it would then be forbidden to use a term that is not approved by the people it refers to: To be called with a name that the target did not choose is allegedly offensive, as it undermines their right to self-determination. This strategy to explain why slurs are prohibited is not convincing either, as the fact that the targets do not choose their

name does not give rise to derogation prohibition in general. Inter-linguistically, it is just standard that speakers choose non-previously-approved terms to call the others. Moreover, there would be no contrast in this respect between slurs and their neutral counterparts. German people are called ‘German’ by English speakers and no one ever asked if they liked it; at the same time, some germanophobic English speakers call them ‘Boche’. It is not the case that the former term was accepted by German people and the latter did not. Intra-linguistically, it is the same: Toddlers were called so by non-toddler speakers without approval, but it does not constitute an insult to them. In fact, Anderson and Lepore (2013a: 39) themselves acknowledge that those who declare a term a slur do not have to be the targets themselves.⁷⁴ In sum, Anderson and Lepore do not provide a convincing meaning-independent reason of why slurs would be the object of prohibition and this grounds the explanation/explanandum problem.

Secondly, note that for Anderson and Lepore slurs do not have any *derogatory content*; the pejorative potential is only determined by the violation of a rule. Recall: “When a word is prohibited, then whoever violates its prohibition risks offending those who respect it” (Anderson and Lepore 2013a: 38). If we take a racist context where speakers systematically violate the prohibition, uttering a slur does not warrant offense and this is a plausible prediction if we understand offense to be a psychological notion, a state of mind of the participants to a conversation. But recall that the violation of the prohibition is all there is to explain the pejorative power of slurs: It is not obvious how Anderson and Lepore can account for the intuition that a slur is still derogatory even if it is uttered in a fully racist environment, given that there is no one to take offense.

A further objection to Anderson and Lepore’s account has to do with how terms come to be slurs. I will leave the discussion of this issue for the next sections, as the problem also concerns Bolinger (2015)’s and Nunberg (forthcoming)’s strategies. Before moving to other deflationary accounts, I would like to make a couple of remarks. The first concerns the notion of taboo. Supposing for the sake of the argument that the prohibition is all that is required to explain the offensiveness of

⁷⁴ “We can imagine slurs for infants or the severely mentally disabled”, Anderson and Lepore (2013a: 46).

slurs, it has to be specified what the prohibition is *really* about. Let's consider three candidates that could be the object of the edict: (i) the meaning (ii) the mention (iii) the mere phonetic realization of slurs. We can immediately exclude (i), as for Anderson and Lepore the meaning of slurs does not differ from the meaning of their neutral counterparts: Since speakers use neutral counterparts of epithets without triggering offense, the edict has not to do with the meaning of slurs. We can also exclude (ii), according to which slurs are prohibited words in the sense that it is forbidden to make reference to them: Speakers often adopt euphemisms such as the 'N-word' to make reference to slurs without being offensive. Therefore, the mere mention of slurs is not what generates offense. We are left with (iii), the phonetic realization, which seems *prima facie* the most convincing candidate and which is also closer to the notion of 'occurrence' invoked by Anderson and Lepore. Nevertheless, (iii) presents problems. If it was the mere phonetic realization that is responsible for offense, then we should find that the occurrence of terms which are homophonic (and possibly homographic) with slurs is offensive while it is not the case. Take the case of the term '*finocchio*' in Italian. It is a neutral term that means 'fennel', but it is also a homophobic slur that translates the English 'faggot'. The same goes for French: the term '*nègre*', which translates the slur 'nigger', also means 'ghost writer'. The allegedly intuitive idea according to which the object of the prohibition is the phonetic realization of slurs is challenged by those cases, where the same phonetic realization gives rise to offensive and non-offensive readings according to the meaning. If Anderson and Lepore want to maintain that the phonetic realization of slurs warrants offense, they may conceive the edict as a prohibition about the combination of a certain phonetic realization (which is not offensive *per se*) and a certain meaning (which again is not offensive *per se*): In this sense, they need – at least – to weaken the claim that the meaning of the term is in no way responsible for the derogatory power of slurs. We shall see in the next sections how Bolinger's and Nunberg's accounts can avoid the issue of specifying what the object of the taboo is, as they do not take mention to warrant offense.

The second remark has to do with *the way* in which Anderson and Lepore support their claim that merely mentioning a slur is offensive. The claim is a very strong one on Anderson and Lepore's side. First because it is a radical position to hold about

how slurs work, as slurs would be different from any other term in this respect; and second because it makes most scholars working on pejoratives (notably including the two authors) responsible for offending every member of every group that slurs target. Against the fact that scholars do mention explicitly slurs in their work, Anderson and Lepore suggest a sort of test that is supposed to destabilize our intuitions: Every time the occurrence of a slur seems non-offensive (for example when mentioned), imagine what would happen if someone – especially a targeted member – overheard you uttering it. While acknowledging that caution is required when slurs are mentioned, this does not seem to me as a strong argument. What is fine to say when we think we are over-heard depends on several and complicated factors and we would better avoid taking such a parameter as a test for offensiveness and the like. Based on the imagine-to-be-overheard argument one could claim that the offensiveness of the non-slurring insult ‘wanker’ cannot be affected by negation: I would probably avoid uttering (35) when my granny can overhear me, but this should not be taken as evidence that the offensive content survives under negation in (35).

35. John he is not a wanker.

The imagine-to-be-overheard kind of argument should be avoided, as it brings up many more issues than needed.

Wrapping up, Anderson and Lepore’s explanation of why slurs are offensive relies on the fact that they are prohibited, but their explanation of why slurs are prohibited in the first place is not satisfactory: They cannot appeal to what slurs *mean* – or their account would not be deflationary anymore – nor can they appeal to the fact that they are offensive, as this is the datum to be explained. The notion of *taboo* is surely very relevant for slurs, but it should be investigated more, as the mere phonetic realization is not a good candidate for being the object of the prohibition. We shall see in the next sections how the other deflationary accounts reinterpret the claim that the offensiveness of slurs does not depend on what slurs *mean*, by leaving aside the notion of taboo and focusing on the very complex and debated notion of markedness.

9.2 Bolinger: co-occurrence expectations and contrastive preferences

Renée Bolinger (2015) draws a parallel between slurs and impoliteness behavior. The core intuition is that the offensiveness of slurs is the result of the pragmatic mechanisms involved in lexical contrastive preferences. Slurs-users choose to use a slur *rather than* its neutral counterpart and their choice therefore signals that they endorse the contents typically associated with the slur.

Bolinger makes reference to the work of two theorists working on rude and impolite behavior: Marina Terkourafi (2005a, 2005b, 2005c) and Jonathan Culpeper (2011). In particular, Terkourafi characterizes politeness and impoliteness in terms of co-occurrence expectations: Certain terms or expressions are typically associated with politeness (or impoliteness) because of “the regular co-occurrence of particular types of contexts and [these] particular linguistic expressions.” (Terkourafi 2005a, quoted in Bolinger 2015). The idea is that there are regularities between contexts characterized by certain extra-linguistic features and the choice of certain expressions over semantically equivalent alternatives. Such regularities determine co-occurrence expectations and the expressions associated with politeness become conventional signals of politeness (or impoliteness). According to Culpeper, on the other hand, co-occurrence and conventionalization are not enough to account for impoliteness (as opposed to politeness). In Culpeper’s account, impoliteness arises from the speaker’s choice to *flout* the expectations of politeness relative to the context. Moreover, being ‘impolite’ is different from being ‘offensive’; two conditions are needed for the impolite behavior to contextually determine offense: (i) the chosen option contrasts negatively with the default minimally polite behavior and (ii) this is a well-known fact in the linguistic community. What these two accounts have in common according to Bolinger (2015) is that “the offensiveness of impolite/rude behavior results from content signaled by the speaker’s decision to perform that particular behavior, rather than a comparatively polite alternative”. Bolinger proposes to apply such an analysis of impoliteness to slurs and she formulates her analysis as follows:

“For some content ϕ , when it is common knowledge in the linguistic community that

- (i) α is an expression for ψ associated with ϕ , and
- (ii) β is an expression for ψ not associated with ϕ

then in situations where the choice of expression is not forced, and the speaker is aware of (i) and (ii), selecting α in contrast to β signals that the speaker endorses or shares ϕ .” (Bolinger 2015)

Moreover, Bolinger makes two assumptions:

“(a) slurring terms are marked as the dispreferred option for achieving reference to the target group.

(b) they are associated with the derogation of their targets.” (*Ibidem*)

Bolinger’s idea is the following: Discriminatory (racist, homophobic, etc.) situations tend to co-occur with the choice of slurs over their neutral counterparts, resulting in the conventionalization of the slur as the impolite way to refer to their target. Using a slur becomes associated with derogatory contents about the targets. On the other hand, non-discriminatory situations co-occur with the use of the neutral counterparts, that conventionally become the minimally polite way to refer to the targets and they therefore typically pass unnoticed.

According to Bolinger, the contrastive choice account can provide a unified analysis of a very general phenomenon, involving not only slurs, but many other expressions that typically signal the attitudes of the speakers: For example, the use of ‘cisgender’ to refer to people whose gender corresponds to their biological sex, unlike expressions such as ‘non-trans’ or ‘normal’ people, affiliates the speaker to the trans community and its allies. However, this affiliation information is *not* encoded in the meaning of the term ‘cisgender’, it is not part of its meaning, just like the derogation of the target is not part of the meaning of slurs.

The first issue to discuss is that there is a contrast concerning how to conceive default contextual expectations between the kind of approach suggested by Terkourafi (2005a, 2005b, 2005c) and Culpeper (2011) on the one hand and the approach put forward by Bolinger (2015), on the other hand. For Terkourafi and Culpeper, what is to be expected as contextual default depends on each context: Co-

occurrence regularities determine the expected lexical choice. That's how Bolinger describes the terkourafian mechanism:

“Default frames for a context are generally minimally polite, so any expression that deviates from that—either impolite or overly polite—is a marked alternative. Use of a marked expression signals that the default frame does not fit the actual context, and use of a conventionalized signal activates the frame for the expression's associated context-type.” (Bolinger 2015)

The idea would be that the use of a slur (for example, ‘wop’) signals that the context is not one of those associated with minimal politeness: The default frame turns out to be inappropriate and the slur activates a frame that fits the ‘wop’-context, that is, a racist one. In such contexts, the co-occurrence expectations make the slur ‘wop’ the most natural way to refer to the target in the racist context, the way of reference that passes unchallenged and unnoticed. “Politeness resides, not in linguistic expressions themselves, but in the regularity of this co-occurrence. To the extent that these expressions go by unchallenged by participants, they are polite” Terkourafi (2005a: 248). On the same lines, according to Culpeper something is impolite when it signals the agent's choice to violate the *contextual* expectations, that is, to choose a certain expression over a more polite one (*contrastive choice*). Moreover, for an impolite expression to warrant offense, speakers must know that it contrasts negatively with the default behavior.

Bolinger, on the contrary, claims that slurs are (always?) the dispreferred option to refer to the target as they are negatively marked and any competent speaker must know that. Let's consider both ways to conceive the impoliteness of slurs in terms of co-occurrence expectations and contrastive choice and see how they would account for the behavior of slurs. Let us start from the terkourafian-culpeperian approach. Take a deeply racist context like a Ku Klux Klan meeting and imagine that a participant uses the slur ‘nigger’. In such an environment, the slur activates the frame appropriate for the racist context, where the default expression to refer to black people is a slur. This way of characterizing contextual expectations accounts for the

intuition that what is standard, unnoticed and unchallenged in a racist environment is not the same lexicon that is standard, unnoticed and unchallenged elsewhere. Therefore, if a KKK member talks about ‘African-American people’ in a context where ‘nigger’ is the standard term for black people, she is choosing to flout the contextual expectations, thus possibly signaling her endorsement of non-racist contents. While doing justice to this kind of intuition, the terkourafian-culpeperian analysis of slurs has a hard time explaining the intuition according to which the slur ‘nigger’ is derogatory *also* when used in the KKK meeting. Such a use would not count as impolite in Terkourafi’s and Culpeper’s sense, as the contextually expected term that goes unnoticed in the KKK meeting is precisely ‘nigger’. It seems hard to explain why the KKK use of ‘nigger’ is derogatory by only relying on the speaker’s choice to flout contextual politeness expectations.

The account of offense is *prima facie* more convincing: Impolite behavior warrants offense when a speaker (i) chooses to use an expression that contrasts negatively with the default minimally polite behavior and (ii) the contrast is a well-known fact in the linguistic community. This accounts for the fact that in the KKK community where everyone endorses racist beliefs and racist slurs are the standard way to refer to targets, no one takes offense for the use of a slur; on the contrary, racists mutually reinforce their attitudes by employing slurring terms. Nevertheless, imagine that a black person is present at a KKK meeting. Again, in the frame that fits the racist context the default expression to refer to black people is a slur. However, the utterance of ‘nigger’ does in fact warrant offense, at least for the witness, independently of the fact that in the KKK meeting the slur is the expected term and everyone knows it. In sum, the speaker’s choice to flout *contextual* politeness expectations *à la* Teroufaki-Culpeper is not enough to account for the derogatory power of slurs nor for how they warrant offense.

As a matter of fact, Bolinger seems to have in mind another way to conceive impoliteness (and thus the behavior of slurs) when she spells out her assumption that slurs, *qua* negatively marked expressions, are *tout court* the dispreferred option to refer to targets, independently of the context in which they occur. This second option accounts for the derogatory power of slurs across different contexts, but it meets difficulties, too. Let’s consider the relation between slurs and the associated

derogatory contents. If we conceive this stable relation as linguistically encoded (however we mean ‘encoded’: At the level of truth conditions, or presuppositions or conventional implicatures, etc.), we would end up with a non-deflationary account of slurs. So there must be for Bolinger another way to conceive a stable conventional relation between certain terms and certain associated contents that does not require linguistic encoding. Bolinger provides such an example when she talks about the term ‘cisgendered’. Her point is that the well-known association between certain terms and certain contents on the basis of co-occurrence does not only involve slurs, but also other sorts of affiliating expressions, for which the speaker’s lexical choice generates implicatures about her beliefs and attitudes. However, Bolinger neglects a crucial contrast between affiliating expressions like ‘cisgendered’ and slurs, a contrast that reveals how the contrastive choice account, however convincing, elegant and interesting it may be, is not apt to account for the derogatory content of slurs. As a matter of fact, it is possible to use the term ‘cisgendered’ literally, without affiliating oneself to the trans community and its allies: One can even use ‘cisgendered’ while explicitly endorsing discriminatory practices against the trans community. On the contrary, it is not possible to use a slur without affiliating to racism.⁷⁵ Compare:

36. Only cisgendered people should have the right to get married.

37. I have nothing against faggots, I’m friend with lots of them.

The contrast in (36)-(37) shows that while it is true of ‘cisgendered’ and ‘faggot’ that they are *typically* used by a certain group (trans-community allies and homophobes, respectively), the associated contents are easily cancelable for ‘cisgendered’ and hardly cancelable for ‘faggot’.⁷⁶ I take this contrast to suggest that while the pro trans-community contents are not linguistically encoded in the meaning of ‘cisgendered’, the homophobic contents are in fact linguistically encoded in

⁷⁵ Note that both occurrences (or ‘cisgender’ and ‘faggot’) are *literal* ones; no irony or friendly mockery is involved. For non-offensive uses of slurs that signal intimacy, see Appropriation, Part III.

⁷⁶ The infelicity depends on a property named ‘Derogatory Autonomy’, according to which “the derogatory force for any epithet is independent of the attitudes of any of its particular speakers” (Hom, 2008: 426). See also Anderson and Lepore (2013a: 33) and Bolinger (2015).

‘faggot’. Note that for Bolinger the relation between affiliating terms α and the associated contents φ is to be understood in terms of high probabilities that the α -user endorses some or all the attitudes of the cluster φ . In particular:

“The more well-known the association between α and φ is, the higher the information content of the signal, and thus the more strongly the contrastive choice signals the speaker’s endorsement of φ ” (Bolinger 2015)

Formulating the association between α and φ in terms of probabilities makes the proposal compatible with the felicity of utterances like (36): Speakers might fail to hold some of the relevant attitudes associated with the term without sounding contradictory or infelicitous. However, this move does not save the account from the objection just raised, as it is not the case that the association between ‘cisgendered’ and the pro-trans community contents is weaker than the association between the slur ‘faggot’ and the homophobic contents. As a matter of fact, one could claim it is the other way around: While nearly no speaker who is not a trans-community ally employs the term ‘cisgendered’, the term ‘faggot’ has, on the contrary, many non-homophobic uses, among which appropriated uses are the most clear cases (for an analysis of Appropriation, see Part III). In Bolinger’s account there is no clear reason why the relation between ‘cisgendered’ and its associated contents should be weaker (and easier to withdraw without infelicity) than the relation between ‘faggot’ and its associated homophobic contents.⁷⁷ The presuppositional account, on the contrary, predicts that while slurs trigger a presupposition, affiliating terms like ‘cisgendered’ do not, as shown by the felicity of utterances like (21). However, the co-occurrence expectations described by Marina Terkourafi (2005a, 2005b, 2005c) and Jonathan Culpeper (2011) and developed further by Bolinger (2015) might in fact account for the behavior of affiliating terms such as ‘cisgendered’: Certain expressions pragmatically implicate certain contents to which they are usually associated and that are not linguistically encoded. Along the same lines, the mechanism of flouting

⁷⁷ One might think that the contrast could depend on an asymmetry between positive and negative associated content, where for some reasons the former would be easier to withdraw than the latter. This does not seem to be the case: I asked Italian native speakers to judge pairs like (36)-(37) containing affiliating-terms *à la* ‘cisgendered’ that are associated with negative contents just like slurs (namely ‘*ideologia gender*’) and I found similar results.

contextual expectations described by Culpeper is available also for non-deflationary accounts (such as the presuppositional one) to explain why, in certain cases, it looks as if the use of a non-loaded term *rather* than a slur is in fact communicating the speaker's dissociation from the widespread discriminatory attitude. The presuppositional account would rely on contrastive preferences to explain correction patterns (see section 2.2) like the following:

38. You think he is a fag.

39. You think he is gay.

In addition to the cases of correction, another interesting example of how the contrastive preference communicates the speaker's endorsement of a non-discriminatory attitude, is the introduction of non-loaded terms when the most common expression for a group is perceived as derogatory. This happened with the term 'handicapped', that got gradually substituted by the expression 'person with disability' by those who wanted to signal that they do not endorse the kinds of discriminatory attitudes often associated with the term 'handicapped'.

Note however that there is a crucial difference – underlined by Stojanovic (p.c.) – between prototypical slurs and affiliating expressions: Slurs convey an evaluation that is *about* their target, whereas the additional content of affiliating expressions need not be about the target. While the additional content of 'wop' is *about* its reference – Italians –, the additional content of 'cisgender' is not about its target – namely, non-trans people (anyone whose gender identity matches the sex that they were assigned at birth) –. This divergence could be responsible for the different behavior displayed in (36)-(37). However, Bolinger would not take such a difference as relevant, as in her account slurs and affiliating expressions are analyzed along the same lines.

Going back to Bolinger's proposal, another problem concerns an interesting phenomenon that is not accurately taken into account by Bolinger, i.e., *complicity*. The phenomenon of complicity was already discussed in chapter 4 and it amounts to the fact that when a slur is used in a context, if speakers do not protest, they are

responsible for accepting the slur and the associated derogatory content. In section 4.1, I showed that the phenomenon of complicity is easily explained by the presuppositional account, because for any presupposition, in the absence of objections, speakers are taken to endorse any content that is accommodated in the common ground, no matter what they think or what their intentions are. On the contrary, it is not entirely clear how the contrastive preference account could explain complicity. According to Bolinger's proposal, a slur is the dispreferred option to refer to the target group, while the neutral counterpart is the preferred one. By choosing 'wop' over 'Italian', the speaker is signaling her own endorsement of the associated pejorative content, but it is not clear how this would make anyone else responsible for it. Note that the notion of 'being responsible' is crucially different from 'being bothered/pleased': Bolinger can perfectly account for why the use of a slur can be bothering, enflaming or enraging for the non-racist audience (or pleasing for the racist audience), but it is hard to see how they could be responsible for the utterance of a slur in their presence. A more fine-grained account of how the derogatory content of slurs interacts with conversational dynamics is needed to explain the complicity phenomenon.

In sum, according to Bolinger speakers pragmatically convey contents about their beliefs and attitude via contrastive preferences and this widespread mechanism is able to explain a variety of linguistic and communicative phenomena, some of which involve slurs and their counterparts; nevertheless, the contrastive preference mechanism does not provide satisfactory predictions for the behavior of slurs, which appeared to differ from the other expressions. On the other hand, the presuppositional account correctly predicts that slurs systematically trigger derogatory contents and it can rely on the contrastive preference account to explain the peculiar cases where a neutral counterpart generates an anti-derogatory implicature. The presuppositional account is therefore less parsimonious, but crucially not *ceteris paribus*, as, unlike the contrastive preference account, it properly captures the data.

9.3 Nunberg: markedness, affiliation and manner implicatures

The core of Nunberg's proposal is that slurs are offensive because they are used by bigoted people (racists, homophobes, etc.). This line of explanation was also

suggested in passing by Anderson and Lepore (2013a: 39): “It may have become prohibited because of who introduced or used it”. Nunberg’s account involves notions such as markedness, metadata knowledge and ventriloquistic implicatures. Before presenting Nunberg’s account, note that the notion of slur that I employ is not the same as Nunberg’s. According to Nunberg,

(...) a derogative word qualifies as a slur only when it disparages people on the basis of properties such as race, religion, ethnic or geographical origin, gender, sexual orientation or sometimes political ideology—the deep fatalities that have historically been the focus of discrimination or social antagonisms that we see as rents in the fabric of civil society. Sailing enthusiasts deprecate the owners of motor craft as ‘stinkpotters’, but we probably wouldn’t call the word a slur—though the right-wingers’ derogation of environmentalists as ‘tree-huggers’ might qualify, since that antipathy has a partisan cast.

I depart from Nunberg in that in my account it does not matter whether the evaluation conveyed by slurs has a partisan cast or not. I believe that the object of investigation in providing an account of slurs is a general linguistic mechanism, while whether the evaluations conveyed are appropriate or not is a matter of social facts and moral theory (that might have linguistic consequences, too: cf. section 3.3). Of course most slurs are more offensive than ‘stinkpotter’, but ‘stinkpotter’ counts however as a slur, as long as it displays the relevant features that I defined in chapter 1.

For Nunberg what distinguishes slurs and non-loaded terms is the affiliatory content: In using a ‘marked’ term (‘marked’ with respect to what is standard in a certain group, not with respect to the neutral counterpart), the speaker affiliates herself with the a group “whose norms wouldn’t ordinarily govern linguistic choices in the speech-situation” (Nunberg forthcoming). Nunberg postulates a submaxim of Manner of the form “Use appropriate language”, whose violation would generate an implicature. The choice of the marked term generates routinized conversational implicatures, as by choosing a slur over its neutral counterpart, speakers violate of

Grice's Maxim of Manner. The offensiveness of slurs is the result of the affiliation to discriminatory groups; such an affiliation is the result of a generalized implicature generated by the exploitation of a Manner submaxim. Nunberg calls this kind of implicature 'ventriloquistic', as when a speaker uses a term that typically belongs to a certain group, it is as if the speaker was impersonating a member of such a group.

Nunberg's theory of pejoratives (or "prejudicals", as he calls them) heavily relies on the notion of *affiliation*. Jeshion (2016b) argues that in Nunberg's account there is no room for the speaker's attitude. The spectrum of attitudes associated with slurs – contempt in the first place, but also disgust, dismissiveness, derision, fear, patronage and superiority – are absent from the affiliation mechanism described by Nunberg. What the use of a slur expresses via a conversational implicature is the speaker's affiliation to a group, not something about her own mental states. Of course one can further infer from the display of a certain affiliation on the side of the speaker that she might also endorse the prototypical attitudes of the group typically using the slur. Nevertheless, according to Jeshion, this explanation fails to account for a crucial feature of slurs, that is the fact that they primarily express a negative attitude. In Nunberg's account the derogation associated with slurs is, in a sense, a second order inference, as it is inferred from an affiliation implicature generated by the use of a slur. In addition to the 'expressivist' perplexity, Jeshion raises another objection about the affiliation mechanism. As Nunberg highlights, the way speakers use language can affiliate them to one group or the other. So for example the mere choice of the expression 'free enterprise' affiliates the speaker to political right, without there being anything in the linguistic meaning of the term that encodes such information. The same idea was suggested in one of the first works about hate speech by Jennifer Hornsby (Hornsby, 2001) who compares slurs and young teenage slang. While acknowledging that the use of a slur can affiliate the speaker to a racist community, Jeshion (2016a, 2016b) challenges the idea that this is always the case: Not for every slur is there a community the slur-user can self-identify with. Nunberg's proposal works well for a slur like 'nigger', as there are movements like the KKK, but it is not clear with what community a 'wop'- or a 'boche'-user would be affiliating. Note in passing that Bolinger's proposal would not raise either of the

these objections, as according to her, slurs are directly associated with the derogation of their target, without inferring such a content from the group affiliation.

In addition to Jeshion's perplexities, there is a further remark to be made about Nunberg's affiliation mechanism, already mentioned in the previous section. There seems to be a contrast between the case of slurs and the case of other affiliating expressions, like 'free enterprise' or the teenage slang. In order to affiliate oneself with right-wing parties or teen-agers, it is *not* enough to use a typical term: There are some other requirements such that if they are not met, the affiliation might fail to be communicated. For example, a leftist can use the term 'free enterprise' with a critical attitude toward the political right without infelicity: "I am a critic of the free enterprise system". No affiliation is accomplished. Similarly, in the case of teenage slang: A grown-up calling a movie 'wicked' just sounds lame or funny.⁷⁸ In the case of derogatory epithets, it seems that the speaker sounds racist the minute she uses slurs (when they are used literally). The contrast between the affiliating expressions and slurs suggests that when one utters a slur there is more going on than just pragmatically signaling affiliation, to the extent that not only using a slur is usually enough to convey derogatory contents, but it is difficult – once one used a slur – to de-affiliate oneself from racism. A similar case was discussed in the previous section, where I compared the behavior of affiliating expressions like 'cisgendered' and slurs and I showed that the asserted content can felicitously withdraw the affiliation for 'cisgendered'-like cases, but it cannot do the same with slurs.

Finally, the *genesis* issue constitutes a problem for all the deflationary accounts that I presented (or "slur-introduction" problem, as Jeshion 2016b calls it). While it might sound convincing that the reason why 'Jap' is offensive and 'Japanese' is not just depends on *who* typically uses the term, other slurs clearly suggest that some pejoratives are coined with the intent to derogate their target. Compounds are a good example: the slur 'spaghetti-eater' for Italians or 'carpet-muncher' for homosexual women were originally coined precisely in order to derogate their target. Deflationary strategies cannot account for the derogatory content associated with 'spaghetti-eater' and 'carpet-muncher' by relying on the group typically associated

⁷⁸ We can also imagine a case of a mocking use; this would probably count as ironic and therefore it wouldn't be a literal use of 'wicked'. We are interested in the case where the term is used literally and this is not enough to accomplish any affiliation.

with the term, as the term was derogatory the moment it was created, long before anyone could develop any kind of metadata knowledge about who typically uses the term, which is the preferred option etc.

Something should be said, though, about how Anderson and Lepore's account could address the genesis issue. Note that Anderson and Lepore (2013a: 39) do not aim to provide a uniform account of slurs:

“We agree with Hornsby, who declares that ‘[a] unified account of a [slur] cannot be achieved by identifying a pragmatic ingredient to be added to a semantic one given by the word's neutral counterpart, because only the word itself provides the outlook from which one can make sense of the variety of associated speech acts’ (Hornsby, p. 135; our emphasis).”

When they consider slurs that are clearly coined with the intent of derogate their target, they acknowledge that some slurs can have a derogatory meaning. “It might even be coined to signify something abusive or offensive about a particular group, and so, prohibited because of its explicit meaning”. So unlike Bolinger's and Nunberg's, Anderson and Lepore's approach can account for compound-like cases, as they welcome the idea that different mechanisms are responsible for the pejorative power of different slurs (we shall see in the next chapter that Väyrynen 2013 suggests something similar for thick terms); however, this maneuver might change the nature of Anderson and Lepore's deflationary project. One of the main advantages of such an approach is that it is very parsimonious, it does not posit any extra component in the meaning of slurs to explain their pejorative power, but some social edicts about language; nevertheless, it looks like a less parsimonious account is needed, as they have to acknowledge that in some cases slurs are offensive in virtue of what they mean (at least some slurs, especially certain compounds).

9.4 A note on markedness

Before concluding, I introduce a brief digression about a crucial notion that was employed by both Bolinger and Nunberg, namely markedness. The aim of my digression is not to list and compare all the ways in which one can understand the

notion of markedness, nor to establish how to best conceive it. Rather, I show how certain ways to understand markedness can help explaining certain behaviors of pejoratives. I will conclude that markedness is very relevant notion to account for many uses of slurs, but it does not seem enough to account for the conventional nature of the evaluative content of pejoratives.

The notion of markedness, however widespread in linguistics, is far from being univocally understood. I will make reference to Martin Haspelmath (2006)'s classification of twelve ways in which markedness has been conceived in 20th century linguistics, grouped in four families of similarities. I will consider two families that seem more promising as far as slurs are concerned: complexity and abnormality.

Starting with complexity, let us leave aside phonological and morphological complexity⁷⁹ and let us consider *semantic* complexity: Haspelmath's example is the pair 'dog'/'bitch'. 'Bitch' can only refer to female dogs, while 'dog' can be either referred to dogs in general or to male dogs. In the case of slurs, this would correspond to the idea that the slur is systematically associated with derogatory contents, but its neutral counterpart is compatible with both discriminatory and non-discriminatory attitudes about the targets. A similar notion was introduced by Greenberg (1966: 28): The idea is that the marked category is optional (for instance 'lioness'), while the unmarked category ('lion') occurs both with a general sense ('lion in general') and in a specific sense opposite to the marked sense ('male lion').

However, there is something non-convincing in the parallel between this kind of markedness and the markedness of slurs. The information about sex that distinguishes 'dog' from 'bitch', unlike the derogatory content that distinguishes neutral counterparts and slurs, does not scope out semantic embedding. Imagine someone wants to adopt a female dog from an adoption center. When she points at an animal in particular, she is told 'Sorry, this is not a bitch. This is a male dog'. When 'bitch' is negated, the information about the animal being female is blocked. On the contrary, we have already seen how negating a slur does not prevent the derogatory

⁷⁹ Take 'Jap', the slur for 'Japanese' or 'Paki', the slur for 'Pakistani': They are not phonologically or morphologically more complex than their neutral counterparts. We can thus discard the idea that slurs are marked because they are phonologically or morphologically more difficult than their neutral counterparts.

content from projecting (section 2.1). So the way in which ‘bitch’ is semantically marked with respect to ‘dog’ cannot account for the pejorative contents associated with slurs.

Another way to understand markedness is as *abnormality*, in the sense of deviation from a default parameter setting. One way to understand the notion of default here is relying on text frequency. In the case of slurs, it is unlikely that frequency is responsible for the alleged markedness of slurs, for at least two reasons. The first one is that for certain slurs, the neutral counterpart is a rather complex expression that is unlikely to be more frequent than the slur. Take the slur ‘JAP’, whose counterpart is roughly ‘Jewish young women with privileged backgrounds’; or ‘twink’ whose counterpart is ‘gay man in his late teens or early twenties’. In general, for both slurs and thick terms, whether the non-loaded counterpart is a common expression or not will depend on the interest that a certain linguistic community has to refer to a certain group identified as such. Sometimes the only speakers who have a word to call what *they* consider a group are those who use a slur and on the other hand the neutral counterpart is a complex paraphrase that people do not happen to employ. The mere fact of having a word to call a group of people reflects some sort of interest in referring to certain people while identifying them as a group: The majority of English speakers displaying an interest in talking about and referring to Jewish young woman with privileged backgrounds are those who also derogate them by using a slur (about the information conveyed by the mere fact of having a word for a group, see chapter 6). The second reason to be cautious with the frequency explanation is that problems arise even if we only focus on the most paradigmatic cases of slurs with non-complex counterparts. It is conceivable that for a speaker that grows up in a homophobic environment the slur ‘faggot’ might be relatively more frequent than its neutral counterpart ‘homosexual man’ and yet it is clear to the speaker that ‘faggot’ is associated with the derogation of homosexual men, while the neutral counterpart is not. Thus, either one acknowledges that the markedness of slurs is not what determines their derogatory power, or, if one wants to maintain that, either one grants that it is not frequency that determines the markedness of slurs, or one has to grant that ‘faggot’ is perceived by a speaker as offensive compared to ‘homosexual man’, even if the neutral counterpart is not in

fact more frequent than the slur in the particular language to which speakers are exposed.

According to the context, there are default expectations which can depend on frequency. So it is in fact true that the term ‘African-American’ is not to be expected in a KKK meeting. Its occurrence is likely to generate conversational implicatures (as it was discussed in the previous section). But this is very different from the case of slurs, whose pejorative power does not depend on an implicature. When ‘African-American’ is marked, that is, when it departs from the default expected standard, it does in fact strike as weird: Using ‘African-American’ at a KKK flouts the expected racist attitude. Nevertheless, as we saw, such an implicature or associated content is quite easy to cancel. The case of slurs is crucially different. Consider (41): While the literal content of the sentence advocates for equal rights, the occurrence of a racist slur makes the whole utterance racist. On the other hand, (40) expresses a racist content (it argues against equal rights), even if it features an expression, ‘African-American’), that is *not* associated with racist contents. However, while the occurrence of a racist expression is enough to make (41) racist, the occurrence of a non-racist expression like ‘African-American’ is not enough to make (40) non-racist. If (41) is uttered in a non-racist situation, the survival of the racist associated content is stronger than the survival of the non-racist content allegedly associated with ‘African-American’ that we would have if (40) was uttered at a KKK meeting. (40) might sound a bit odd to the KKK members, but less problematic – or so I argue – than how problematic (41) would be if pronounced among non-racist people. I argue that the contrast between (40) and (41) depends on the fact that the pejorative content of ‘nigger’ is linguistically encoded and does not depend on the slur being marked.

40. [in a KKK meeting] African-Americans should not have the right to vote.

41. [among non-racists] Niggers should have equal salary for equal work.

Note that the presuppositional account can rely on the notion of markedness to explain certain phenomena involving pejoratives: As we saw, according to the

presuppositional account, slurs trigger an additional content that is responsible for their derogatory power; so the presuppositional account *can* appeal to the idea that slurs are semantically more complex and that they might be cognitively more costly than their counterparts. Moreover, it is compatible with the presuppositional account that each context determines certain default expectations whose violation generates implicatures. In sum, markedness may be relevant for many uses of slurs, but it does not seem enough to account for the conventional nature of the evaluative content of pejoratives. The presuppositional analysis can explain certain phenomena by appealing to markedness, but it does not put all the explanatory weight on this notion.

9.5 Conclusion

Deflationary accounts of slurs underline how the offensiveness of slurs is (at least partly) determined by the mere fact of choosing the slur word rather than its neutral counterpart. I showed that, however promising and parsimonious this strategy is, it fails to account for crucial aspects concerning slurs. The mechanisms described in various ways by the deflationary account defenders (taboo effect, contrastive preference, metadata knowledge) are all relevant to explain the behavior of slurs. Nevertheless, the question is whether such mechanisms are all that is required to account for the behavior of epithets. In this chapter, I showed that this is not the case and that in fact a less parsimonious theory is needed to properly explain the data. However, the deflationary accounts are available to non-deflationary strategies to explain certain features of slurs and evaluatives in general. Moreover, they also account for the intuition that not only slurs, but also non-loaded terms can trigger inferences about the speaker's beliefs and feelings.

We shall see in the next chapter a deflationary account of thick terms, developed by Pekka Väyrynen.

10. Deflationary accounts of thick terms

The pragmatic analysis of thick terms was developed independently from those of slurs and is due to the work of Pekka Väyrynen. Väyrynen is among the first to systematically apply the tools of linguistics and philosophy of language to the

analysis of thick terms. To this extent, Väyrynen closely presents and discusses alternative accounts of thick terms, in order to show the superiority of the pragmatic approach. Väyrynen's main opponent are the theories that embrace the tenet that "the meanings of thick terms (...) somehow or other contain evaluation" (Väyrynen 2013: 9), claim that the author calls 'Inherently Evaluative' (IE). The IE approaches discussed in Väyrynen (2013) differ from each other in the way in which they take the evaluation to be encoded, that is, in the way they interpret "meaning" and "somehow or other" in the above definition. This kind of strategy includes theories that analyze evaluation as part of the truth-conditional meaning of thick terms, or as conventional implicatures, or as presuppositions (construed as lexically encoded content). The bulk of Väyrynen's own proposal is to abandon IE and hold instead that thick terms do not linguistically encode evaluation: The evaluation associated with thick terms (*T-evaluation*, in Väyrynen's terms) conversationally arise as pragmatic implication. These implications are not part of the asserted content, nor are they the main point of the utterance in their literal uses in normal contexts: In other words, they are typically backgrounded, not at-issue. The fact of being usually backgrounded might suggest a presuppositional approach. However, according to Väyrynen, the fact that the uses of thick terms systematically trigger certain evaluations does not automatically show that evaluation is systematically or lexically encoded: Evaluation can be triggered in virtue of some shared extra-linguistic beliefs/knowledge concerning what sort of evaluative perspective is normally held, and commonly known to be held, by speakers who recognize a given thick term as 'one of their words'. Note moreover, that there is no need for a uniform pragmatic mechanism: There can be different mechanisms at play each time to explain the emergence of evaluative contents.

As we have mentioned in section 1.3, even though his work mainly concerns thick terms, Väyrynen (2013: 150-156) briefly discusses the option of treating pejoratives⁸⁰ as objectionable thick terms and Väyrynen (2016a) leaves room for the possibility that his account of thick terms might be applied to slurs. In particular, he suggests that the evaluative content of thick terms could be analyzed in Bolinger's

⁸⁰ Väyrynen (2013: 151) takes the class of pejoratives to be a "mixed bunch", as he does not distinguish between slurs and the rest of pejoratives. Therefore not all of his arguments against the uniform treatment pejoratives and thick terms apply to the uniform analysis of slurs and thick terms.

terms of contrastive choice. I shall add that Väyrynen's approach seems *prima facie* compatible with Nunberg's theory of slurs as well. The main contrast is that for Väyrynen there can be different conversational mechanisms at play each time, while in Bolinger's analysis there is, as we saw, a uniform mechanism at play, the one based on contrastive choice and co-occurrence expectations. However, it is interesting for us to note how scholars from (relatively) different fields developed accounts for slurs and thick terms respectively that analyze these evaluative terms along similar lines.

Väyrynen supports his view by (i) showing how it captures the data and (ii) challenging the rival theories. Let's consider this second line of arguments and then come back to the first one. Väyrynen's arguments against IE theories that analyze T-evaluations as part of the truth conditions of thick terms resemble the kind of arguments that we put forward when discussing the truth-conditional approach to slurs (chapter 7) and Kyle's theory of thick terms (chapter 8). The arguments against the IE theories that analyze T-evaluation as a non-truth-conditional component of meaning (namely, conventional implicatures and lexical presuppositions) are based on the *defeasibility* of the evaluative content of thick terms. Väyrynen takes this feature to show that T-evaluations are not part of the conventional meaning (however broadly conceived) of thick terms and are better analyzed as pragmatic implications. I will consider the data supposedly in support of the pragmatic analysis and select the cases that I recognize as genuine instances of defeasibility. I show that my presuppositional analysis of thick terms, unlike a conventional implicature approach, is apt to account for such cases.⁸¹ Note moreover that Väyrynen also takes into consideration the *Variability* of the evaluative content associated with thick terms (see Part III) as an argument against IE theories, but he acknowledges that different IE approaches can develop a bunch of strategies to account for such cases. In Part III I will propose an analysis of the suspendability and reversibility of the evaluative content of thick terms by relying on notions such as irony and non-literal uses of language, that not only answers the Variability concern, but suggests once more a uniform treatment of slurs and thick terms.

Let's now consider Väyrynen's defeasibility data:

⁸¹ This will possibly answer the objections raised in Väyrynen (2016b).

42. ⁸² Whether or not Madonna's stage show is lewd, it would be in no way bad for that. (Väyrynen 2013:71)

43. ?? Madonna's show is lewd, but I don't mean to imply that it is bad in any distinctive way for its overt and transgressive sexual display. (Väyrynen 2013: 104)

44. Whether or not Madonna's show involves explicit sexual display, it would be in no way bad for that. (Väyrynen 2013:71)

45. ? Whether or not life is short but sweet, there would be no contrast between life being short and life being sweet. (Väyrynen 2013: 104)

According to Väyrynen, (42) and (44) are both felicitous utterances, whereas (43) and (45) are not. Moreover, for the author, the contrast between the felicity of (42) and the infelicity of (45) shows that, while the conventional implicature content activated by 'but' is non-defeasible, the evaluative content associated with 'lewd' is defeasible. While I agree with Väyrynen on the felicity of (44) and the infelicity of (43) and (45), I strongly doubt the felicity of (42). It is unlikely that one utters something like 'Whether or not Madonna's stage show is lewd', if one believes – as stated in the second clause – that something *cannot* be bad for being sexually explicit beyond conventional boundaries (i.e. if one is in fact a 'lewd'-objector). Väyrynen suggests that one might utter (42) as a polite strategy to sidestep the issue of whether Madonna's show can be 'lewd' or not, in order to focus (in the second clause) on the rejection of the associated negative evaluation. According to Väyrynen, (42) does not count as a misuse of 'lewd', nor is the speaker contradicting herself. However, a politeness case is not a good example of "literal use" in "normal context" that is apt to argue in favor of the felicity of (42), which otherwise sounds very odd. I acknowledge, though, that there is some contrast between the infelicity of (42) and

⁸² I changed the original numbers of the examples to adapt them to the numerical sequence of this chapter.

that of (45), in the sense that (42) appears to be *less* infelicitous than (45). Such a contrast could be explained in terms of the availability of the pure descriptive content (discussed in chapter 6): If the utterer of (45) did not think that there is a contrast between life being short and life being sweet, she could easily use ‘and’ instead of ‘but’. On the other hand, the ‘lewd’-objector does not have an exact evaluation-free counterpart of ‘lewd’: She needs to come up with a paraphrase. This unavailability of a non-loaded counterpart can explain why (42) is less infelicitous than (45): The participants to the conversation are more charitable in making sense of the speaker’s utterance and they interpret it as a kind of ‘loose’ talk, where what the speaker meant to express was more precisely (44). Hare (1963) suggests something along the same lines, by claiming that when there is no neutral counterpart, speakers would have to go for a longer and complicated paraphrase or – alternatively – they would employ the evaluative term (whose evaluative content they do not endorse), as if it was mentioned:

“(...) a person who is not inclined to despise people just because they are negroes⁸³ will reply that (...) he prefers not to use the word ‘nigger’. (...) He substitutes, let us suppose, the neutral word ‘negro’. (...) If there were a person who was not in the least disposed to commend those who preserved the safety of others by disregarding their own, then he could say, as before, ‘I prefer not to use the word ‘courageous’, just because it incapsulates this attitude to which I do not subscribe. I prefer the longer, morally neutral expression, ‘disregarding one’s own safety in order to preserve that of others’. (...) It is true that there is no single evaluatively neutral word, like ‘negro’, which in the present case can be used to describe such actions without committing the describer to any evaluation; but we could have such a word. What I shall actually do, in default of an invented word, is to use the same word ‘courageous’, but to make it clear by my tone of voice or by putting quotation marks round it, that I am (...) implying thereby no commendation whatever”. (Hare 1963: 188-189)

⁸³ As we have seen in section 6.2, Hare, who is writing at the beginning of the Sixties, takes the term ‘negroes’ to be a descriptive non-derogatory term.

Going back to the politeness scenario suggested by Väyrynen, I argue that such a scenario could also produce a reading where (43), that Väyrynen recognizes as *strongly* infelicitous, would be relatively fine: Strictly speaking it would be a misuse of a term, but speakers make sense of the utterance nevertheless in virtue of the unavailability of the non-loaded equivalent word. But again, the politeness case does not count as a “literal use” in “normal context”. If we stick to literal uses of language, I argue, (42) and (43) are infelicitous, while only (44) is a perfectly fine utterance.

On the other hand, note that my point is not to state that T-evaluations cannot be suspended. For example, there are contexts where the evaluation conveyed by thick terms interacts with semantic embedding and is genuinely suspended under certain predicates. In the following examples, the evaluation is not ascribed to the speaker, i.e. it does not project:

46. Back then, people thought it was lewd to paint a naked person. Nowadays, everybody thinks there is nothing bad in things that are sexually explicit beyond conventional boundaries.

47. Lucy praised her friend Sue and told me that she is generous, but I don't think there is anything good in sharing one's resources without expectation of receiving anything in return.

I find (46)-(47) much more convincing cases of defeasibility of T-evaluation than (42). These utterances featuring embedded thick terms, can be accounted for in terms of presuppositions that get plugged by attitude verbs, just like standard presuppositions in (48)-(49). On the contrary, in (50)-(51) we observe that conventional implicatures are typically non-defeasible, not even when embedded.

48. Back then, people thought that the accused stopped paying his debts. Nowadays, everybody knows he never paid them.

49. Lucy told me Mary that stopped smoking, but I don't think that Mary used to smoke in the first place.

50. ? Back then, people thought that Mary was rich but honest. Nowadays, everybody knows there is no contrast between being honest and being rich.

51. ? Lucy told me that Mary is rich but honest, but I don't think that there is any contrast between being honest and being rich.

I agree with Väyrynen about the fact that the defeasibility of the evaluative content of thick terms speaks against a conventional implicature account, since conventional implicatures are typically not defeasible, as shown in (50)-(51) (Grice 1989, Potts 2005). On the other hand, I disagree on how Väyrynen discards the presuppositional approach. According to him, a presuppositional account of thick terms that analyzes their evaluative content in terms of lexically encoded presuppositions, cannot account for utterances like (52), as the contrast between (52) and (53) is supposed to show:

52. Whether or not this is good in any distinctive way, John can be truthfully and neutrally described as being considerate. (Bergström 2002: 5, quoted in Väyrynen 2013: 84, 115)

53. ?? Whether or not there is a King of France, we can truthfully say that the King of France is bald. (Väyrynen 2013: 115)

For Väyrynen, it is possible to felicitously apply a thick term while signaling agnosticism about the relevant T-evaluation. I am not sure about the felicity of (52): (52)-like constructions appear particularly odd when used for objectionable thick terms:

54. ? Whether or not this is bad in any distinctive way, John's show can be truthfully described as being lewd.

The presuppositional example provided by Väyrynen – (53) – is supposed to show a contrast between standard presuppositions and thick terms in (52)-like constructions. But (53) only proves that *some* presuppositions are infelicitous in (52)-like constructs. It is broadly accepted in the literature that presuppositions triggered by different triggers display different projective behaviors (Abbott 2006, Abusch 2002, 2010, Abrusán 2016).

As I was suggesting before, it is in the first place dubious whether (52) (and (54)) should count as felicitous and I am much inclined to say that they do not. In order to ‘fix’ (52) and (54), I would reformulate them as (55) and (56) respectively, which sound far more felicitous:

55. I don’t know whether it is good in any distinctive way to show regard for another’s feelings. But if it is, we can truthfully describe John as being considerate.

56. I don’t know whether it is bad in any distinctive way to be sexually explicit. But if it is, we can truthfully describe John’s show as being lewd.

These constructions pattern again like standard presuppositions:

57. I don’t know whether John used to smoke in the past. But if he did, then we can truthfully say that he stopped.

As we saw, the bulk of Väyrynen’s arguments against a IE presuppositional view relies on the defeasibility of T-evaluations, but defeasibility – at least in convincing literal cases – can be accounted for in terms of presuppositions. If one agrees with me that uses of thick terms *à la* (52)/(54) are *not* actually felicitous, then Väyrynen’s argument loses its main source of support, as the felicitous reformulations *à la* (55)-(56) can be accounted for in terms of presuppositions.

Väyrynen advocates a parsimonious account of thick terms: A semantic view is theoretically more costly than his pragmatic approach, as it requires building

evaluation into the meaning of thick terms, providing them with a complex semantics. However, parsimony arguments only make sense in *ceteris paribus* conditions. The question is: Are conditions *pares*? I hope I showed that they are not, as the reasons to go for a pragmatic account involve analyzing cases like (52) or (54) to be just as felicitous as utterances like (55) or (56).

We shall see in Part III that the Variability of the evaluation associated with thick terms – typically used as an argument in support of the pragmatic account of thick terms – can be analyzed in terms of echoic uses of language. I will argue that this speaks in favor of a uniform analysis of slurs and thick terms.

11. Conventional Implicature: an alternative hybrid account

In the previous chapters of this part, I discussed the so called one-source accounts of slurs (after Jeshion forthcoming). Unlike the truth-conditional and the pragmatic accounts, the conventional implicature (CI) analysis is a hybrid approach: It agrees with truth-conditional theories that the derogatory content of epithets is linguistically encoded, but it shares with the deflationary approaches the idea that the pejorative content does not belong to the truth-conditional dimension of meaning. Such an approach, that analyzes the derogatory content of slurs in terms of conventional implicatures, was proposed by Christopher Potts (Potts 2005, 2007) and developed by Eric McCready (McCready 2010) and Daniel Gutzmann (Gutzmann 2011). It was also supported *i.a.* by Williamson (2009) and Whiting (2013). While the previous approaches I presented have been independently developed for both slurs and thick terms, the conventional implicatures view was proposed for slurs only and we shall see why.⁸⁴

11.1 Pott's proposal

Potts (2005)'s seminal work provides a logic for conventional implicature (CI) and the fifth chapter is dedicated to the treatment of expressives and epithets in terms of

⁸⁴ As I mentioned in the Introduction, Copp (2001, 2009) proposed an analysis of moral terms based on conventional implicatures. However, he characterizes conventional implicatures in such a way that they turn out to be very different from Grice's ones and from Pott's ones. Copp himself decided to distinguish the inferences he had in mind from conventional implicatures by introducing the label of 'conventional simplicature'.

conventional implicatures. As Potts (2007a) states, “The history of conventional implicatures is rocky, their current status uncertain”. However, Potts provides a minimal definition:

- “i. p is a conventional (encoded) property of a lexical item or construction in S.
- ii. p is entailed by S.
- iii. p’s truth or falsity has no effect on the at-issue content of S.” (Potts 2015)

According to Potts, slurs (together with other expressives) introduce descriptive content at the at-issue level and expressive content at another dimension of meaning, namely the one of conventional implicatures. Conventional implicatures are lexically encoded, they are part of the conventional meaning of words. Because the CI content has no relation or interaction with the at-issue content (it is compositionally independent), in Potts’ analysis the expressive content is completely independent from the descriptive one. Another crucial point is the non-displaceability of the expressive content. Expressive content is always referred to the situation of utterance: Even when syntactically embedded, the expressive content of slurs is semantically unembedded (which is not surprising, given the fact that it cannot interact with at-issue content). A consequence of non-displaceability is that the expressive content is typically speaker-oriented, that is, it is typically referred to the speaker at the time and place of the utterance.

“Another characteristic distinguishing expressive meaning from propositional meaning is that it is valid only for the utterer, at the time and place of utterance. This limitation it shares with, for instance, a smile, a frown, a gesture of impatience [. . .]” (Cruse 1986: 272, quoted in Potts 2007c: 169)

McCready (2010) and Gutzmann (2011) suggested some changes to Potts (2005)'s proposal, by challenging two of his claims: (i) that there are no expressive modifiers (i.e. expressions that can have expressives as their argument) and (ii) that there are no mixed expressives, that is, terms that contribute descriptive and expressive content *at the same time*. However, for our present purposes, we can stick to Potts (2005) formulation of the theory.

For Potts, an utterance like (58) is true iff Bianca is Italian:

58. Bianca is a wop.

Moreover, because of the occurrence of the lexical item 'wop', (58) conventionally implicates a certain kind of attitude of the speaker at the time of the utterance. This content about the state of mind of the speaker is computationally and logically independent from the at-issue content, that is 'Bianca is Italian'. Moreover, Potts underlines two features of expressive terms: The repeatability (i.e. the fact that when a speaker repeats an expressive term, the repetition typically strengthens the effect rather than give rise to redundancy) and the ineffability (i.e. the fact that when speakers try to paraphrase expressives in non-expressive terms they typically get a frustrating feeling of non-satisfaction).

One might wonder why the CI account was developed for slurs and not for thick terms.⁸⁵ As we briefly discussed in chapter 10, the non-defeasibility of conventional implicatures suggests that the evaluative content associated with thick terms cannot be analyzed as conventional implicature. As we saw, the evaluation conveyed by thick terms seems to be embeddable under *verba dicendi* and other constructs that typically plug presuppositions:

59. I don't know whether it is bad in any distinctive way to be sexually explicit beyond conventional boundaries. But if it is, we can truthfully describe John's show as being lewd.

⁸⁵ Eklund (2011: 39) considers and discards the possibility of analyzing thick terms as lexical items that activate conventional implicatures.

60. Back then, people thought it was lewd to paint a naked person. Nowadays, everybody thinks there is nothing bad in pieces of arts that are sexually explicit beyond the conventional boundaries.

As we shall see, it is less clear for slurs whether their evaluative content can interact with semantic operators.

11.2 Challenges

Pott's account and the one that I am proposing are similar under several points of view: In analyzing the pejorative content of slurs as not-at-issue, in requiring a hybrid analysis to account for the behavior of slurs, in postulating some form of pejorative meaning linguistically encoded in the lexical entries, etc. Why then going for my account over the other? I will consider two sets of reasons concerning on the one hand the linguistic data and on the other the effects and powers of slurs. The first reason to prefer the presuppositional account is that, as we saw, the CI account, as Potts conceives it, does not leave open the possibility for the derogatory content of slurs to be affected by the at-issue content. Conventional implicatures are by definition *totally* independent from the at-issue content, which means that they cannot be the arguments to attitudinal predicates, quantifiers, etc. As Potts acknowledges,

“This severely limits the class of true conventional implicatures. Indeed, following Bach (1999), I am forced to conclude that many of the items [usually considered conventional implicatures] are just secondary at-issue entailments.” Potts (2015)

Any case where there seems to be a semantic interaction between the at-issue content and the not-at-issue content of expressives would constitute a problem for Pott's account and for a CI account of slurs. Kratzer (1999) famously points out that certain expressives embedded in reported speech can be ascribed to the reported speaker rather than the actual speaker:

61. My father screamed that he would never allow me to marry that bastard Webster. (Kratzer 1999)

Without committing ourselves to the idea that slurs should get the same analysis as expressions like ‘bastard’, we can reformulate Kratzer’s examples with slurs:

62. My father screamed that he would never allow me to marry a wop.

In the same spirit, Schlenker (2003, 2007) argued for the possibility for the derogatory content of slurs to be plugged under a *verbum dicendi*:

63. I am not prejudiced against Caucasians. But John, who is, thinks / claims that you are the worst honky he knows. (Schlenker 2003)

The examples *à la* Kratzer are therefore hard to explain for Potts. On the contrary, the presuppositional account admits a wider range of possibilities. Presuppositions are typically blocked when the trigger occurs under the scope of non-factive attitudinal predicates and *verba dicendi*. However, Karttunen (1973) noticed that plugs tend to “leak”, in the sense that presuppositions can often survive. This is compatible with the idea that when for an utterance of the form “x said that y is an S”, where ‘x’ and ‘y’ are agents and ‘S’ is a slur, there are at least two readings: one according to which it was x who uses a slur without y’s endorsement and one according to which y at least endorses the attitude associated with ‘S’. The presuppositional theory has less restrictive requirements in this sense, because it is compatible with both interpretations. Lasnik (2007) has speculated on the reasons for which plugs like the *verba dicendi* are typically ‘leaky’ when slurs are involved:

“Because expressives are so emotionally charged, and because their use can carry a significant social risk, I suspect that speakers are especially

cautious about using them in embedded contexts where there is a chance of their content ‘leaking’ – except, of course, if the speaker does agree with the content of the expressive and is willing to make this agreement public.” Lasersohn (2007)

As we said, Potts holds that the fact that conventional implicatures are independent not only means that they scope out of presupposition holes, but also that “they should invariably evade plugs and filters, because any plug behavior would intermingle the at-issue and conventional implicature content”. This claim is challenged by several scholars, among which Wang et al. (2005), Amaral et al. (2007) and Schlenker (2007). Moreover, some recent experimental investigation on slurs provided some data that seem to speak in favor of a presuppositional account. Panzeri and Carrus (2016) compared the offensiveness of embedded slurs (under negation, antecedent of conditionals, questions and indirect reports) and the offensiveness of slurs in isolation (not only unembedded, but also out of context). What they found is that while the pejorative content of slurs typically scopes out of semantic embedding (especially antecedents of conditionals and questions),⁸⁶ results are less strong in the case of indirect report. The reporter is perceived as offensive, but, crucially, to a lesser degree than the person uttering an unembedded slur or a slur embedded in the antecedent of a conditional or in a question. They interpret their findings as speaking in favor of the presuppositional approach:

“We did not get clear results on what happens in indirect reports, since the degree of offensiveness of a person who reports a slur uttered by someone else is perceived offensive to a lesser degree. Within a presuppositional approach to the meaning of slurs, this might be interpreted as indirect evidence that *verba dicendi* do constitute plugs (that is, they block the offensiveness of the slur, that is thus attributed only to the person who actually uttered the slur), and that the person who reports the slur

⁸⁶ See a brief discussion of Panzeri and Carrus’ data on negation and metalinguistic negation in chapter 7.

uttered by someone else is felt to be less accomplice of the choice of that word” (Panzeri and Carrus 2016)

Note however that it is not the case that *some* subjects interpreted the reported slurs as offensive whereas others as non-offensive: Slurs embedded in indirect reports were perceived as offensive, but crucially less than slurs embedded under other operators. One way to interpret the data is to say that each subject could somehow access both readings: the leaking one, on which the report is offensive, and the plug one, on which the report is not offensive. The upshot of such a mixed intuition is a judgement of offensiveness that is somehow halfway between offensive and non-offensive.

Moreover, appositives typically cannot take a quantificational subject (Potts 2005, Wang et al., 2005), whereas predicative presuppositions have no constraints as to which quantifier they can take. As a matter of fact, while (64) is typically taken to be infelicitous, (65) is felicitous. What is interesting for us is that in (66) we see that slurs pattern like the predicative presuppositions rather than like appositives:

64. ? Most older people on the march, who, incidentally, left after Jesse Jackson, got home without too much trouble. (Potts 2005: 127, quoted in Sudo 2012: 38)

65. Most older people on the march stopped shouting. (Sudo 2012: 38)

66. Most professors are kikes.

The case of quantified subjects suggests that slurs are more similar to presuppositions than to CIs in this respect. The second kind of reasons to prefer my presuppositional account over the CI one has to do with the explanatory power of the presuppositional account. The theory of presupposition offers a rich set of tools to deal with the dynamics of conversation. In chapter 4, in particular, we discussed

how, with the use of a hybrid evaluative, speakers typically take for granted certain perspectives and tend to impose such a perspective on the audience. In my account, the derogatory content of slurs tends to slip into the common ground unless it is objected to. This is why speakers have a feeling of complicity when slurs are used in their presence; the feeling that, if the use of pejoratives is not challenged, this might in a way legitimate future uses as well and maybe even certain kinds of discriminatory behaviors. The CI account has no such direct way to explain the complicity phenomenon. The CI content informs on the emotional state of the speaker and should therefore have a more ‘subjective’ status (Camp 2013: 333 raised a similar objection). It is not quite straightforward to pass from a speaker *personal* emotional state to a social dimension, where the conversation participants are responsible for the use of slurs.

Another aspect for which the presuppositional account provides a more satisfactory answer concerns failure (section 3.3). As a matter of fact, speakers who do not endorse the discriminatory attitude associated with a slur, typically are reluctant to judge the truth value of slurring utterances or to answer questions involving slurs. As mentioned already, the main feature of CIs (present both in Grice 1989’s and Potts 2005’s characterization) is that they are truth-conditionally irrelevant, so we should expect speakers to be able to easily separate the two supposedly independent dimensions of meaning. On the contrary, my proposal holds that to employ a hybrid evaluative is to take for granted some kind of perspective. A perspective where certain descriptive properties warrant certain evaluations. If one does not share or accept such perspectives, the utterances featuring hybrid evaluatives are simply defective, in the sense that speakers would understand them as making wrong assumptions. On similar lines, the presuppositional account offers a better explanation of the *inappropriateness* of using a slur in a non-racist context: The inappropriateness consists in doing as if a certain attitude was shared and granted, when it is not (see chapter 5).

Note however that it is not impossible for a CI account to formulate answers to the perplexities I raised. Take the projection phenomena, for example, and the apparent embeddability of slurs (*à la* Kratzer cases). A CI-defender could appeal to the notion of ‘perspectival shift’. We shall discuss at a greater length in Part III the

issue of when and how the evaluative content of hybrid evaluatives can be suspended or reversed. The proposal that I will put forward is largely independent of which theoretical approach to slurs and thick terms one favors. However, I shall mention some of the issues that are going to be discussed, in order to present Potts' strategy to deal with the non-offensiveness of certain embedded slurs such as those we observed in (62)-(63). Harris and Potts (2009) put forward the notion of "perspectival shift". The idea is that a speaker can take for a moment someone else's perspective, to the effect that expressions typically preferring speaker-oriented interpretations (expressives and appositives) receive a shifted interpretation. They are not interpreted as attributed to the speaker, but to another doxastic agent. Potts' rationale is to explain the cases of apparent semantic embedding in terms of pragmatic effects. If the fact that the pejorative content of slurs in certain cases does not project out of embedding (*à la* Kratzer cases) is indeed due to pragmatic effects, then Pott's claim that the expressive content does not interact with at-issue content is safe. Harris and Potts' examples involve appositives (that Potts analyzes in terms of conventional implicatures) and expressives such as 'idiot' and 'bastard'. Consider a couple of examples involving appositives:

67. "I am increasingly worried about my roommate. She seems to be growing paranoid. The other day, *she told me*⁸⁷ that we need to watch out for the mailman, a possible government spy" (Harris and Potts 2009a: 530)

68. "My friend Sal is absurdly optimistic. *He told me*⁸⁸ that the lottery ticket he bought yesterday, a sure winner, is the key to his financial independence." (*Ibidem*: 548)

The appositives 'a possible government spy' and 'a sure winner' would typically be attributed to the speaker even if they occur in the scope of a *verbum dicendi*; yet, in these contexts, they get attributed to the paranoid roommate and to the absurdly optimistic Sal, respectively. Harris and Potts interpret this shift as due to pragmatic

⁸⁷ Italics mine.

⁸⁸ Italics mine.

effects rather than to semantic binding by attitude predicates, because they find that the shift can also occur in the absence of semantic embedding (even if they found a less strong effect):

69. “I am increasingly worried about my roommate. She seems to be growing paranoid. The other day, she refused to talk with the mailman, a possible government spy.” (Harris and Potts 2009a: 530)

70. “My friend Sal is absurdly optimistic. All he could talk about at dinner was the lottery ticket he bought yesterday, a sure winner” (Ibidem: 548)

I agree with Harris and Potts that most of their examples, together with some examples from Amaral et al., (i) allow non-speaker-oriented readings for expressions that typically prefer speaker-oriented readings and (ii) allow a shift plausibly due to pragmatic effects. However, I do not take these examples to show that *all there is* about non-speaker oriented readings is pragmatically driven. They only show that certain pragmatic effects having to do with perspectives and playing the role of someone else are able to ascribe the content of expressions that typically prefer speaker-oriented readings to other doxastic agents. We shall come back to these cases at a greater length in Part III.

Let’s now consider how Potts could take into account the other perplexities that I voiced. Consider the reluctance of speakers to judge the truth value of slurring utterances: The presuppositional account can explain it in terms of non-catastrophic failure, whereas the CI account can claim that the reluctance only depends on a confusion between the technical notion of ‘true’, as distinct from ‘assertable’, and a folk notion of ‘true’, that somehow involves endorsement. It is less clear to me how to explain the more ‘social’ effects of slurs, such as complicity, within the CI-framework. In section 4.1, I mentioned a couple of solutions that Potts could adopt, even if I do not find them very convincing. The first one is to claim that people are in certain ways responsible for the beliefs endorsed by the people they deal with; in this sense, avoiding challenging the use of a slur would be endorsing the attitude associated with the slur, even if all the slur was doing was to inform about the mental

state of the slur-user. The second strategy would be to rely on a general notion of so-called presupposition of commonality, that is, as we saw, the expectation that the conversation participants share are “all alike under the relevant respects”. The notion is used by Lopez de Sa (2008) to talk about predicates of taste: “The relevant expression (...) triggers a presupposition of commonality to the effect that the participants of the conversation are all alike in the relevant respects”. As we said, if we admit an extension from the taste domain to the moral dimension, we would have the following account: A slur like ‘wop’ activates the conventional implicature that the speaker despises Italians *qua* Italians and in addition it also triggers the presupposition that the conversation participants feel the same way about Italians.

It should not be surprising after all that a presuppositional account and a conventional implicature account are very similar. As a matter of fact, it is a very debated issue among scholars how to draw a difference, if any, between lexically encoded presuppositions and conventional implicatures. There are two main lines, none of which is fully convincing. The first strategy has to do with the discourse status of presuppositions and conventional implicatures respectively (Potts 2005, Horn 2007): CIs typically introduce new information, whereas presuppositions are typically presumed to be already true. If things were so sharply and clearly divided, one could just check if the derogatory content of slurs is typically introduced as a new or an already-shared content. However, slurs are uttered both in contexts where a certain discriminatory attitude is part of the background and in contexts where it is entirely new (and it can be accommodated or not). As a matter of fact, on the one hand, presuppositions can be introduced as new information (it is the case of informative presuppositions, discussed in sections 4.1 and 5.2) and on the other, CIs can be evoked also when they are not new: Speakers just retrieve already known information in order to make it salient.

The second strategy to distinguish lexically encoded presuppositions and conventional implicatures is the possibility to felicitously neutralize the entailed content. Because the conventional implicature cannot interact with the at-issue content, the criterion is that if the not-at-issue content triggered by a certain lexical item can somehow be neutralized, then it is presuppositional; otherwise, it is a conventional implicature (see *i.a.* Potts 2005, 2015, McCready 2010). In chapter 5, I

argued that the suspendability of the triggered content should not be considered as a crucial test for presuppositions. My point was that many authors have come up with problematic examples showing that the presuppositions of certain triggers are very hard to neutralize and I mentioned some examples, including the gender features of pronouns (Heim and Kratzer 1998, Sudo 2012) and particles like ‘too’ (Chemla and Schlenker 2012). Moreover, some instances of the embeddability of slurs under *verba dicendi* seems to suggest that the pejorative content of slurs is, however, suspendable in certain cases, that is, it is *hard* to suspend, not *impossible* to suspend.

11.3 Conclusions

It is not surprising after all that there is not much consensus about how to distinguish presuppositions and conventional implicatures, as the two notions have a very complicated history (“close neighbor[s] in the linguistic literature”, as Abbott 2006 calls them): They have been sometimes conflated and sometimes one has been analyzed in terms of the other. Karttunen and Peters (1979), for example, used the label ‘conventional implicature’ to describe conventionally triggered presuppositions. Other scholars have similarly interpreted the label ‘conventional implicature’ as picking out a particular type of presupposition; in this framework, it would not be surprising to find lexical items with some kind of mixed behavior, somewhere between paradigmatic cases of presuppositions and paradigmatic cases of conventional implicatures (Lasnik 2007, Sauerland 2007, Schlenker 2007, 2009, 2010). Moreover, some lexical items feature in both conventional implicature and presupposition lists (it is the case of particles like ‘even’, ‘only’ and ‘too’ and, of course, slurs).⁸⁹

As we have seen, the distinction between CIs and presuppositions is in itself very controversial and I have the impression that the debate about pejoratives has been taken hostage by this more general debate. I do not aim at solving the main issue of how to characterize and distinguish presuppositions and conventional implicatures *in general*, I stick to the observation of slurs and go for the theoretical option that offers

⁸⁹ For a taxonomy of projective content that sets aside the notions of presupposition and conventional implicature, see Tonhauser et al. (2013).

better strategies to account for the phenomena, which, as I have argued, is the presuppositional account.

In the light of the above discussion, it does not come as a surprise that the CI analysis was far more popular in the literature on slurs than in the literature on thick terms. As we saw in chapter 6, the evaluative content of thick terms can usually be suspended in an easier way than the one of slurs and in a much easier way than conventional implicatures.

71. Back then, painting a naked person was regarded as lewd. Nowadays I find it hard to conceive that something might be considered bad for being sexually explicit beyond conventional boundaries.

72. ?? Back then, they regarded that ad, the first ad to represent a naked person, as inappropriate. Nowadays, we know for sure it was not the first ad to represent a naked person.

In (71) the evaluative content of ‘lewd’ is *not* ascribed to the speaker and therefore the second clause sounds felicitous; on the other hand, the CI content in (72) is strictly ascribed to the speaker even if it is embedded in the same ‘regarded as’ construction and therefore the second clause sounds infelicitous.

In sum, I discussed the ways in which a conventional implicature account of slurs is similar to a presuppositional one and, despite the controversy of how to characterize conventional implicatures vs. presuppositions, I hope I succeeded in showing that the presuppositional account offers a richer set of tools to understand the way in which slurs work and provides a deeper insight.

Part III. When the evaluation changes

In Part I and Part II I assessed the general question as to how slurs and thick terms are best analyzed and I compared my presuppositional account to alternative theories. In this part, I focus on a more specific issue. I propose a uniform account of two phenomena that have been treated as independent so far: Appropriation of slurs and Variability of thick terms. It is not surprising that Variability and Appropriation were treated as unrelated phenomena, given the fact that the investigation on slurs and thick terms in the first place has been mainly conducted in two different – even though related – fields (see section 1.3 for noteworthy exceptions): philosophy of language and linguistics for slurs and ethics and metaethics for thick terms. The proposal that I will present, designed to account for Appropriation and Variability within a presuppositional theory of hybrid evaluatives, is actually compatible with possibly any account of slurs and thick terms and should be considered as a contribution to the investigation on evaluatives that is independent from my presuppositional framework.

In chapter 12, I will focus on Appropriation, a particular use of slurs (12.1) and compare it to Variability, a particular use of thick terms (12.2). In chapter 13, I propose a uniform account of Appropriation and Variability, relying on the notion of “echoic use” of language, introduced by Relevance Theory (Sperber and Wilson, 1986, Carston, 1996, 2002, Wilson, 2006a, 2006b, 2007 and Wilson and Sperber, 2012) and employed by Bianchi (2014) to analyze the appropriated uses of slurs. I will argue that this echoic approach does not only explain the Appropriation of slurs, as Bianchi claims, but it also accounts for the Variability of thick terms. I show how the relevance-theoretic distinction between ‘merely attributive’ and ‘echoic’ uses of language sheds light on the crucial distinction between the two often conflated phenomena of ‘suspension’ and ‘reversal’ of the evaluative content of slurs and thick terms (13.1). I argue that the theoretically relevant distinction is not between Appropriation and Variability, rather between ‘suspension’ and ‘reversal’ of the evaluative content of slurs and thick terms. In chapter 14, I discuss the theoretical outcomes of my proposal for the theory of thick terms: If Variability cases are in fact derivative or parasitic uses of language, they do not count anymore as a counterexample to the claim that thick terms

convey or express⁹⁰ evaluation in a systematic way. This suggests once more that slurs and thick terms as similar phenomena. Finally, I conclude by suggesting that the relevance-theoretic notions of “attributive” and “echoic” uses of language seem to be related to similar notions such as “perspectival shift” (Harris and Potts 2009a, 2009b and Harris 2012) and “polyphony” (Ducrot 1984 and Recanati 2006). The phenomenon at stake – namely the possibility to attribute a certain content to someone else, while possibly expressing one’s attitude towards it – is a very general feature of human language that goes beyond evaluative terms and deserves deeper investigation.

12. The impermanence of evaluation

In this chapter, I present Appropriation and Variability, cases where slurs and thick terms typically conveying a certain evaluation fail to carry it or carry a reversed one.

12.1 Appropriation of slurs

As mentioned in section 4.1.1, with ‘Appropriation’ we refer to those cases where the members of a group can use among themselves the slur targeting their own group, in such a way that the slur is not offensive anymore in those contexts; on the contrary, appropriated slurs are used to express solidarity and underline intimacy. In such contexts, slurring does not constitute an act of genuine impoliteness and it resembles in a way the kind of sarcasm or mock impoliteness that is typically associated with intimacy (Leech 1993, Culpeper 1996, 2016). In light of the debate on hate speech among scholars, media and institutions, Appropriation deserves special attention, especially in the spirit of what Jacob Mey calls “emancipatory” or “anticipatory pragmatics” (Mey 1985, 1993, 1994, 2012), a discipline that “(...) proactively promotes use of language in non-oppressive ways; it foresees and prevents abusive language of all kinds by enabling the users, both on the domestic and the (inter) national scene” (Mey, 2012: 705).

In most cases, appropriated uses of slurs are typically available only for ‘in-groups’, i.e. the members of the group targeted by the epithet. The prototypical example is the

⁹⁰ In what follows, I will use ‘carry’, ‘communicate’, ‘convey’ and ‘express’ interchangeably. In this part I will stay neutral with respect to how slurs and thick terms and the associated derogatory content are in fact related, in order to show that my proposal is not dependent on any specific account of slurs and thick terms.

occurrence of the slur ‘nigger’ (or its variant ‘nigga’) in some African-American communities, especially in Rap cultures. For some terms the process of Appropriation reached a point where, under certain conditions, also out-groups could use the slur in a non-offensive way: It is the case of ‘queer’ used in certain academic contexts, where it is accepted – and definitely not homophobic – to talk about ‘Queer Studies’. For some other terms the process of Appropriation is fully completed and the standard use of the term does not carry derogatory contents anymore: It is the case of the term ‘gay’, that used to be derogatory until it underwent a process of Appropriation in the early Nineties and today it represents the standard non-offensive term for ‘homosexual’ (Brontsema 2004).⁹¹ As ‘gay’ shows, Appropriation is a process at the end of which epithets can lose their derogatory power for good; however, when we talk about ‘appropriated uses’ of slurs we refer to those cases where the process is ongoing, i.e. the cases where, in standard contexts, the term still carries a derogatory content.

Appropriation raises interesting issues such as whether the intention of speakers can interact with the disparaging content of slurs, whether the non-offensive occurrence of epithets should be considered the occurrence of a new lexical item etc. and it has been investigated both on theoretical (Tirrell 1999, Brontsema 2004, Croom 2011, 2013, 2014 and Bianchi 2014) and empirical grounds (Galinsky et al. 2003, Galinsky et al. 2013). Even though there is no unique unanimously accepted account of Appropriation, a fairly widespread view is that appropriated uses of slurs are somehow ‘derivative’ or ‘parasitic’ on the standard offensive uses.⁹² In other words, the mere existence of Appropriation is not *per se* evidence that the derogatory content associated with slurs is unstable, as the non-offensive uses are parasitic on the offensive ones.

⁹¹ Another interesting example is the term ‘deaf’: the Deaf community refused to be referred to as ‘hearing impaired’ people and they reclaimed the pride of the term ‘deaf’ associated with a specific culture and a specific language (Moore and Levitan, 2003). In this case, it is the in-group community that imposes a term to the out-groups.

⁹² Mišćević (2011:176) defends the idea that appropriated slurs are different lexical items from derogatory slurs. Mišćević and Perhat (2016: 140) give up the polysemy thesis and support an account that is close to Bianchi (2014)’s echoic account.

12.2 Variability of thick terms

Let us leave aside for a moment slurs and consider the literature on thick terms. Many authors (Hare 1952, Blackburn 1992, Väyrynen 2011, Eklund 2013) noticed that thick terms display a feature called Variability. Consider the following example:

1. The carnival was a lot of fun. But something was missing. It just wasn't lewd. I hope it'll be lewd next year. (Väyrynen 2013: 85)

'Lewd', that typically conveys a negative evaluation, is used in a positive sense in (1). The theoretical importance of Variability has to do with the so-called Variability Argument, according to which Variability shows that the relation between thick terms and the evaluation that they express is not systematic or semantically encoded and it must be fully pragmatic in nature. Väyrynen (2011: 3-4) formulates the argument as follows:

(V1) Variability Claim: Thick terms and concepts are contextually variable in evaluative valence.⁹³

(V2) Link Claim: If thick terms and concepts are contextually variable in evaluative valence, then the evaluations that they may be used to convey don't belong to their meanings.

(V3) Therefore, the evaluations that thick terms and concepts may be used to convey don't belong to their meanings; they are rather a feature of pragmatics.

Väyrynen (2011) considers a few ways in which the Variability Argument can be challenged. One option is to reject (V2) and explain the Variability of 'lewd' in utterances such as (1) by analyzing thick terms as context-sensitive expressions. The core idea is that thick terms do not necessarily encode only *one* kind of attitude, but

⁹³ Väyrynen (2013:115-116) accepts a version of (V1) according to which it is at least possible to conceive a *literal* occurrence of 'lewd' such that it does not convey a negative evaluation (either it fails to carry evaluations at all or it carries a positive one). In other words: One could admit that *certain* occurrences of (1) are ironic and nevertheless insist that (1) *can* have non-ironic felicitous occurrences. My proposal targets to this version of (V1) as well.

instead a *range* of possible attitudes that are contextually specified (Dancy 1995, 2004: 192-198). What is linguistically encoded is the potential to convey evaluative content; what is context-dependent, on the other hand, is the polarity of such an evaluation. In other words, according to Väyrynen, the context sensitivity analysis accepts (V1), but rejects (V2) and does not entail (V3), as evaluations (at least potential evaluations) are encoded in the linguistic meaning of thick terms. However, the strategy that I will present in this section consists in rejecting (V1), while staying neutral with respect to (V2): Therefore (V3) does not follow anymore.

To reject (V1) does not mean to deny basic linguistic evidence, such as the fact that ‘lewd’ conveys a positive evaluation in (1), nor that the evaluation usually conveyed by ‘lewd’ is negative. Rather, it consists in questioning the assumptions that utterances like (1) should be considered as genuine *literal* uses of ‘lewd’ and that the way in which ‘lewd’ conveys a positive evaluation in (1) is the same in which it usually conveys a negative one. Väyrynen (2011) mentions such a strategy, that he ascribes to the “semantic invariantist”. The semantic invariantist would consider (1) as a non-literal use of ‘lewd’, a “way of mocking the sorts of prudish evaluations which ‘lewd’ conveys as a matter of meaning” (Väyrynen 2011: 10). While this approach was not widespread among scholars in ethics and metaethics, Väyrynen (2011) remarks that Blackburn (1992, 1998) noted that (1)-like utterances sound as a way to mock prudish people and – I shall add – they would therefore count as non-literal instances of language.

In the next chapter I will present an account of Variability that does justice to the intuitions *à la* Blackburn by making reference to the relevance-theoretic notions of attributive and echoic uses of language. My proposal accounts for a very general mechanism of language, showing how Variability cases like (1) behave just in the same way as many other cases of attributive and echoic uses of language.

13. The echoic account

In the previous chapter, I presented two phenomena: Appropriation – concerning slurs – and Variability – concerning thick terms –. The two phenomena have a feature in common: They represent cases where slurs and thick terms that typically convey a certain evaluation, fail to carry it or carry a reversed one. However, there is a crucial

difference in the way in which the two phenomena were treated in the literature.⁹⁴ In the debate on thick terms, Variability gave rise to the Variability Argument, challenging the claim that the evaluation conveyed by thick terms is lexically encoded and supporting contrariwise the view that the relation between thick terms and the evaluation that they express must be fully pragmatic (Väyrynen 2013: 115-116 and 221-223). In the literature about slurs, on the other hand, appropriated uses have been mostly analyzed as derivative uses of language that do not challenge the claim that slurs convey evaluation in a systematic way. Appropriated uses of slurs are often viewed as non-literal uses of language that do not tell much about the lexical meaning of slurs. My proposal is to adopt the latter kind of approach to explain Variability as well. First, I will first present the account of Appropriation proposed by Bianchi (2014) and then I will show how the same approach suitably accounts for Variability as well.

Bianchi makes reference to the relevance-theoretic notion of “echoic use” (Sperber and Wilson 1986, Carston 1996, 2002, Wilson 2006a, 2006b, 2007 and Wilson and Sperber 2012). In order to present the notion of echo, we will need to introduce some useful distinctions, as well as some jargon. For Relevance Theory we can distinguish “descriptive” and “interpretive” uses of language:

An utterance is descriptively used when it is used to represent a possible or actual state of affairs; it is interpretively used when it is used to represent another representation (for instance, a possible or actual utterance or thought) that it resembles in content. (Wilson, 2006a: 1729).

Moreover, in some interpretive uses of language, the meta-represented content is attributed to someone else: These are “attributive” uses. Note that the meta-represented content (thought, utterance, expectation, etc.) does not necessarily have to have been *actually* uttered or thought by someone in particular, it is just *attributed* to someone. Consider the following examples (Wilson 2006: 1730, quoted in Bianchi 2014: 39):

2. a. The Dean spoke up. b. The university was in crisis.

⁹⁴ Richard (2008: 7) alludes to the idea that Appropriation and Variability could raise the same kinds of concerns.

3. a. The students were thoughtful. b. If they didn't act now, it might be too late.

The idea is that the thoughts expressed in (2b) and (3b) – that the university was in crisis and that if the students didn't act immediately it might be too late – is not something the utterers of (2) and (3) are asserting. Instead, the content of (2b) and (3b) is implicitly attributed to the Dean and the students respectively.

Echoic uses are a subclass of attributive uses that not only represent utterances, thoughts or expectations of some other individuals about certain states of affairs, but in addition they also communicate the attitude of the speaker with respect to such contents. In other words, through echoic uses speakers are able to communicate their attitude with respect to the evoked content. In Relevance Theory, echoic uses account for irony: In ironical utterances, speakers evoke some utterance, thought or expectation attributed to someone else and, at the same time, they express their dissociation⁹⁵ towards it and reject it “as ludicrously false (or blatantly inadequate in other ways)” (Wilson and Sperber 2012). Consider the following cases:

Case 1: Lucy and Sue are planning to have a pic-nic on the banks of the Seine. Lucy utters (4):

4. Let's go on Saturday! It's going to be a nice sunny day.

When they meet on Saturday, it is raining cats and dogs and Sue utters (5):

5. What a nice sunny day.

⁹⁵ About what should count as 'dissociation' in the relevant sense, see Wilson and Sperber (2012):

“Dissociative attitudes themselves vary quite widely, falling anywhere on a spectrum from amused tolerance through various shades of resignation or disappointment to contempt, disgust, outrage or scorn. The attitudes prototypical of verbal irony are generally seen as coming from the milder, or more controlled, part of the range. However, there is no cut off point between dissociative attitudes that are prototypically ironical and those that are not”.

Case 2: Lucy and Sue planned to have a pic-nic on the banks of the Seine. When they meet, it is raining cats and dogs and Sue utters (5).

In (5) Sue echoes the thought or the expectation that the weather should be good and at the same time, since the good-weather expectation was not met, she expresses her dissociative feelings towards it. The echoed content can be a thought, utterance or expectation that someone in particular actually expressed (Lucy, in Case 1), or it can be generally attributed to an indefinite someone else – a potential someone else, so to speak – without someone in particular in the context having expressed it (a general expectation of nice weather for a pic-nic, in Case 2).

Going back to slurs, Bianchi (2014) explains Appropriation by employing the framework proposed by Relevance Theory to account for irony. In the case of appropriated uses of slurs “in-groups echo derogatory uses in ways and contexts that make manifest the dissociation from the offensive contents” (Bianchi 2014: 40). For Bianchi, appropriated uses of slurs are in this sense ironical. My proposal is to extend such an analysis to Variability: Speakers can use a thick term that conventionally carries a certain evaluation in order to convey the opposite one. This means that, going back to the Variability Argument, I reject (V1). I therefore reject the idea that the evaluative content of thick terms is contextually determined (see Dancy 1995, 2004), as well as the idea that it entirely relies on pragmatic factors like conversational implications (see Väyrynen 2013). Instead, I favor the view according to which (1) echoes the lewd-user’s way to apply the term, it makes fun of it and it communicates the speaker’s dissociative attitude.⁹⁶ Note also that the echoic account can explain why it is easier to obtain reversed uses of *objectionable* thick terms: It is easier to dissociate from (and make fun of) a perspective that one clearly does not endorse. However, while the parameter of ‘objectability’ is relevant for reversed uses to obtain, ‘polarity’ is not. Take for instance this passage from the COCA (Davies 2008-), where ‘chaste’,

⁹⁶ It is worth noting that the echoic account of irony developed by Relevance Theory received considerable experimental investigation (Jorgensen et al. 1984, Gibbs 1986, Happé 1993, Keenan and Quigley 1999, Spotorno and Noveck 2014). This suggests that there might be experimental ways to establish whether speakers interpret (1)-like sentences as *literal* uses or as *echoic* uses of language. In particular, Spotorno and Bianchi (2015) put forward some experimental proposals to test whether appropriated uses of slurs are better understood as echoic uses of language. If the parallel I am drawing between Appropriation and Variability is on the right track, their designs might well be applied to the cases of Variability such as (1)-like utterances.

that typically conveys a positive evaluation, reversely conveys a negative one. The speaker makes fun of a prim perspective that he does not endorse and he calls a quick kiss on the cheek ‘chaste’ in order to complain about it:

6. ‘Not sure how long I’ll be gone. (...)’. Elaine gave him a quick kiss on the cheek. ‘That was a little chaste’. ‘Don’t look now, but we seem to be of interest to about fifty elderly women on the tour bus behind you’ ‘Should we give them something to stare at?’

To have echo, there need to be two distinct and opposed perspectives. If the speaker endorses the evaluative content of a term (or she is taken to endorse it) it is much more difficult to contextually reverse such an evaluation. Here is an example where the writer seems to *mock* a perspective that she does not share when using the term ‘chaste’:

7. The ultimate threat to the new model family, at least in discursive terms, was the unmarried mother. The ‘fallen woman’ stood in stark opposition to the chaste domestic angel of the ideal family, contained by marriage and economically dependent. (Simonton, Deborah (2006), (ed.) *The Routledge history of women in Europe since 1700*, London and New York, Routledge)

It is interesting that while the writer uses inverted commas to signal that she does not endorse the perspective encapsulated in the expression “fallen woman”, she does not need to use ‘visible’ inverted commas when she uses ‘chaste’. And she succeeds nevertheless in attributing such a prim perspective to the supporters of the “new model family”.

13.1 Merely attributive and echoic: a useful labelling

Scholars interested in Variability have tended to focus on valence-reversal cases – like (1) – where the speaker does display a dissociative attitude toward the evaluation conveyed by thick terms (see Blackburn 1992, 1998, Väyrynen 2013):

1. The carnival was a lot of fun. But something was missing. It just wasn't lewd. I hope it'll be lewd next year.

As I argued in the previous part of this chapter, such uses would count as echoic – or ironic – in relevance-theoretic terms. Nonetheless, the dissociative attitude and the ironical flavor are not a necessary requirement for the speaker not to endorse the evaluative content of slurs and thick terms. It will be useful here to rely on a distinction mentioned earlier: merely attributive vs. attributive and echoic uses. As we saw, not all attributive uses are echoic: For an attributive use to be echoic, the speaker must communicate her dissociative attitude towards the evoked content, i.e. the content that she attributes to someone else. The notion of merely “attributive use” can account for those cases where the evaluation conveyed by slurs and thick terms is suspended without being reversed. Consider the following examples involving thick terms. Here are some occurrences of the term ‘lewd’ in contemporary texts about Manet’s work and models:

8. As Abigail Solomon-Godeau (1986) has convincingly demonstrated, a rich tradition of pornographic photographs portraying lower-class prostitutes in a variety of lewd poses existed in 1865. (Bernheimer, Charles (1989), *Manet's Olympia: The Figuration of Scandal*, Duke University Press).

9. This meditation on the effects on time revealed to Manet that he was not wasting himself on transitory and lewd genre scenes. (Tinterow, Gary and Loyrette, Henri, *Origins of Impressionism*, Harry N. Abrams, New York).

One possibility is to infer that the authors of (8)-(9) endorse the negative evaluation conveyed by the term ‘lewd’. However, it looks like in (8)/(9) the negative attitude towards sexually explicit scenes is not necessarily endorsed by the speaker (writer, in this case), but it is rather attributed to someone else, probably Manet’s contemporaries, who commented on his work back in the 1860s. The hypothesis suggested by such cases is that when evaluative terms are involved, we often evoke someone’s perspective. Such a perspective usually coincides with the one of the speaker

(‘speaker-oriented’ uses), but it does not have to. Examples like (8)-(9) show that speakers can take perspectives that belong to some other (possibly virtual) actor, explicitly mentioned or even present in the context of utterance, or only implicitly evoked. If that is the case, the relation between (8)-(9)-like cases and the (1)-like cases can be satisfactorily captured by analyzing it in terms of merely attributive vs. attributive and echoic uses: In the latter use, in addition to evoking the perspective of someone else, the speaker also expresses her dissociation. Again, note that the availability of attributive uses is not related to a particular polarity. We can find the same patterns with thick terms conveying a positive evaluation. Take for instance the term ‘chaste’:

10. Thus, unmarried women are expected to be chaste and virginal, and once married, women are expected to sacrifice in favor of their children and husbands. (McNeill et al. 2001)

The positive evaluation of sexual abstinence conveyed by ‘chaste’ in (10) seems to be attributed to someone different from the writers, namely the community that the authors are describing. Note moreover that in some clearly attributive cases, it might be ambiguous whether the speaker is ‘just’ attributing a certain perspective to someone else, or whether she is *also* expressing her dissociation from it (and therefore echoing it). Consider the following cases from the COCA (Davies 2008-):

11. The Cherokee scholar Rayna Green points out as well that the Pocahontas legend constitutes part of a ‘perplex’ in which Native American women are objectified as chaste ‘Princesses’ or, conversely, as highly sexualized ‘squaws’.

12. There were so many rules we girls were supposed to adhere to, so much emphasis on propriety. Straight backs. Gloved hands. Unpainted (and unkissed) lips. Pressed skirts, modest words, downturned eyes, chaste thoughts. A lot of nonsense, in my view.

Let us now consider the case of slurs. It is debated whether the derogatory power of slurs can be suspended at all. I do not take into account the case of mention, as I focus on those cases where slurs are in fact used. Scholars who claim it is possible to suspend the derogatory power of slurs provide examples like the following. Consider:

13. I talked to my parents. It went bad. My father screamed that he would never allow me to marry a wop. (Adapted from Kratzer, 1999)

14. I am not prejudiced against Caucasians. But John, who is, thinks/claims that you are the worst honky he knows. (Schlenker 2003, 2007)

15. Institutions that treat Chinese as chinks are morally depraved.⁹⁷ (Hom 2008: 438)

In these cases, the speaker evokes the perspective of her father in (13), of John in (14) and of racist institutions in (15). By using a slur word, she seems to simply attribute a certain racist attitude to her father, to John and to racist institutions respectively, without necessarily being ironical about it. In a way, by employing the slur, the speaker temporarily takes another perspective. Harris and Potts (2009a, 2009b) introduced the notions of “perspectival orientation” and “perspectival shift”, developed further by Koev (2014) and closely related to the notion of attributive use. Harris and Potts are interested in non-speaker oriented readings of appositives and expressives. They consider cases involving non-slurring expressives like (16) and (17) (Harris and Potts 2009a: 538):

16. My classmate Sheila said that her history professor gave her a low grade. The jerk always favors long papers.

17. My buddy Connor said that his boss bought (such) expensive computers for the office. The imbecile wants to impress the CEO at the next board meeting.

⁹⁷ Hom (2008) talks about “pedagogical” uses of slurs. I am not entirely convinced by Hom’s examples of pedagogical cases, but assuming that such uses are in fact felicitous, they can be analyzed in terms of attributive uses of language.

According to them, the derogatory attitude conveyed by ‘the jerk’ in (16) is interestingly not ascribed to the speaker, but to Sheila, whose perspective is in fact taken by the speaker (and the same for ‘the imbecile’ attributed to Connar in (17)). Note that Koev (2014) claims that there must be an utterance whose speaker is the doxastic agent to which the perspective shifts. I disagree with Koev on this and I agree with Sperber and Wilson: There need not be an actual utterance with an actual speaker, because, as we saw, one can echo a thought or an expectation that is attributed to a real or an undefined individual. Even though the notion of perspectival shift presents some analogies with the one of attributive uses, Harris and Potts (2009b) tend to conflate attributive and echoic cases under the label of “perspectival shift”, mixing ironical and non-ironical cases; in this sense, the relevance theoretic taxonomy offers more fine-grained distinctions. Nevertheless, these examples support the claim that the evaluative content typically attributed to the speaker can be ascribed to some other agent *via* pragmatic mechanisms.

Note in passing that among all the different perspectives that a speaker can take, there is also her own past (or future) perspective, that counts as ‘other’. Consider:

18. I wasn’t a big opera fan when I was a child. I used to find opera quite boring. Then I grew up and the tedious arias suddenly engaged me deeply.

In this example, ‘tedious’ refers to the speaker’s past perspective and it counts as an attributive and echoic use. In the case of slurs, such a past-self perspective is harder to get because the issues involved in hate speech and taboo words are quite delicate. Nevertheless, an utterance like (19) gives a hint of what it would be for slurs to echo the speaker’s past-self perspective: In (19), the ‘other’ agent to whom the derogatory content triggered by ‘fags’ is ascribed is the speaker himself, but a past-version of him.

19. Before I finally accepted I am gay, I was very homophobic. When my family moved to San Francisco, I made a scene about living in a city of fags.

In a nutshell, I analyze Appropriation and Variability in terms of attributive and echoic uses of language: Attributive uses usually correspond to the suspension of the evaluative content of slurs and thick terms, while echoic uses typically correspond to the valence reversal, characterized by an ironical flavor.

13.2 Appropriation and Variability: residual differences

In the previous chapter, I briefly mentioned the fact that when slurs are involved, it is particularly difficult to obtain attributive readings, without the toxic attitude carried by slurs being ascribed to the speaker. This leads us to point out some crucial differences between the Variability of thick terms and the Appropriation of slurs. Those differences have to do with social factors rather than linguistic mechanisms, but should not be neglected. First of all, since slurs are taboo words with a strong social load, for appropriated uses to be admissible it must be completely unquestionable that the speaker is strongly and honestly opposed to the negative evaluation that slurs encode, to such an extent that appropriated uses are initially only available to in-groups; in the case of thick terms, it is less clear what an ‘in-group’ would be;⁹⁸ it might be harder to detect the echoic intentions of the speaker, since they might not be so evident (what could be common ground is whether the speaker is a thick-term-user or not, which could help interpreting an attributive use as such). In general, the clearer such a dissociative intentions will be in the context, the easier it will be for the reversed and ironical effect to obtain.

The second factor to consider is that Appropriation of slurs has to do with oppression, violence and fight and it therefore requires some awareness of the political significance of word-choice (see Tirrell 1999). In this sense, Appropriation is a fully aware process pursued by individuals and groups with a strong will to get rid of the toxic power of certain expressions while retaining the pride. In the case of thick terms, there is no political project, there typically is no pride movement: Evaluations can be reversed without any strong political valence, but just with standard irony.

That being said, the presence of some political awareness – for how interesting and worth noting – does not make the linguistic mechanism of Appropriation different

⁹⁸ Maybe someone who instantiates (or engages in behavior that instantiate) the descriptive properties that are the object of the evaluation.

from that of Variability and it has to do with political motivations rather than with linguistic mechanisms. In this sense, it does not make much sense to analyze along different lines Appropriation on the one hand and Variability on the other; rather, it is useful to distinguish attributive uses of evaluatives on the one hand from echoic uses on the other.

13.3 What's left of variability

The account of Appropriation and Variability that I have put forward can be seen, in a way, as deflationary, as it claims that there is no actual variation in the lexical meaning of slurs and thick terms in cases like (20) or (1).

20. I've been grinding outside, all day with my niggas. And I ain't going in, unless I'm with my nigga. (line from the rap song *My Nigga*, feat. Young Jeezy and Rich Homie Quan)

1. The carnival was a lot of fun. But something was missing. It just wasn't lewd. I hope it'll be lewd next year.

Nevertheless, this deflationary approach concerns the contextual variability and it does not exclude that the evaluations conventionally associated with certain expressions can change over time. That is very clear in the case of slurs. We can observe cases where a neutral term gets associated with a certain evaluative content to such an extent that the evaluation is finally built into the meaning of the term. Moreover, the opposite direction of change is also possible: The completed process of appropriation of slurs can turn expressions that used to convey a negative evaluation – like 'gay' – into neutral ones. It is less easy to find similar cases among thick terms.⁹⁹ Perhaps 'pride'¹⁰⁰ could be a good candidate. 'Pride' indicates a high opinion of one's own dignity and importance; today it doesn't have a negative valence (quite the

⁹⁹ Departing from slurs and thick terms, another case of valence change concerns the English adjective 'terrific', that appears in the 1660s with the meaning of 'frightening', coming from the Latin '*terrificus*'. Around 1888, we start finding cases where the connotation is inverted, meaning 'excellent'. Again, the evaluation expressed changed, but not contextually. See http://www.etymonline.com/index.php?allowed_in_frame=0&search=terrific&searchmode=none.

¹⁰⁰ I owe this example to Isidora Stojanovic (p.c.).

opposite), but it seems to have had a very negative valence in the past. Take this passage from Jane Austin's *Pride and Prejudice*.

21. "His pride," said Miss Lucas, "does not offend me so much as pride often does, because there is an excuse for it".

In general, the attributive-echoic analysis of Appropriation and Variability does not deny *tout court* the possibility that the evaluation conveyed by certain expressions can change over time; it only denies that literal occurrences of slurs and thick terms can convey different evaluations according to the context.

14. Echo and perspectives

In this part, I've shown how the echoic account of appropriated uses of slurs can explain the so-called Variability of thick terms. As we said, the investigation on slurs and thick terms has been mainly conducted in two different fields, philosophy of language and linguistics for slurs and ethics and metaethics for thick terms and scholars working on slurs often neglect thick terms and vice versa.¹⁰¹ Moreover, by rejecting (V1), my account provides an analysis of Variability that is compatible with the Conventionality Claim, according to which the evaluation conveyed by thick terms is lexically encoded. Nevertheless, as we said at the beginning, the echoic account is possibly compatible with any account of slurs and thick terms respectively, even though it is probably not a very economical explanation for pragmatic theories: If the evaluation is not encoded, it would be a bit complex – though not implausible – to claim that the reversed evaluation stems from the echo of an evaluation pragmatically and conversationally associated with a lexical item. However, Väyrynen himself, being the main defender of a pragmatic approach, acknowledges that (1)-like examples can be analyzed in terms of echo:

"Several different uses of a term get called inverted commas uses. One function of inverted commas is to echo or allude to the beliefs or attitudes of people other than the speaker herself (Hare 1952, 124)" (Väyrynen 2013: 85).

¹⁰¹ With some noteworthy exceptions, as we have seen in section 1.3.

Note also that if the evaluation of thick terms is analyzed as stable in their literal uses, this speaks in favor of the uniform account of slurs and thick terms that I discussed in Part I. The main obstacle to a unified treatment of the two kinds of expressions is in fact the idea that while slurs *systematically* convey a negative evaluation, the evaluation of thick terms is variable (in being *reversible* and *defeasible*, see chapter 6 and chapter 10).¹⁰² That being said, the echoic account of Variability and Appropriation is not committed to the idea that slurs and thick terms rest upon one and the same mechanism. Indeed, they could convey evaluations in very different ways and yet speakers could attribute and echo such evaluations via the same mechanisms. The ability of taking someone else's perspective while possibly expressing one's point of view regarding such a perspective is a very general feature of language that is not circumscribed to evaluatives. Evaluatives might be just a particularly interesting case for perspectival shift, given that evaluations typically require a point of view from which the evaluation is made (see among others McCready 2007). In this sense the investigation about Appropriation and Variability – or, more in general, about valence suspension and reversal for slurs and thick terms – might shed light on a far more general and pervasive mechanism of human language.

Finally, in these chapters I made reference to Relevance Theory and I employed some relevance-theoretic terminology. The main reason is that the notions of merely attributive vs. attributive and echoic uses of language capture particularly well the kinds of distinctions I am after. However, this choice does not commit my account to endorse Relevance Theory as a whole, as it only relies on a very specific portion of it. Moreover, several projects in the last years dealt with phenomena that are closely related to attributive and echoic uses of language. I have already briefly mentioned Harris and Potts' work on perspective shift (Harris and Potts 2009a, 2009b, Harris

¹⁰² “Supposing that ‘lewd’ introduces a negative evaluation of explicit sexual displays into the common ground would make (23) [for us (1)] sound highly odd. So it is very hard to see how the SP [Semantic Presupposition]View could explain tokens of (23) [here (1)] where lewd is used literally. The PP [Pragmatic Presupposition] View has in principle no problem allowing that such utterances of (23) [here (1)] can effect a contextual suspension of the normal negative implication of lewd-utterances” (Väyrynen 2013: 115-116).

2012) and I should also mention François Recanati's work on polyphony¹⁰³ (Recanati 2006). Despite the similarity of the notions involved, the experimental material used by Harris and Potts (2009a) includes at the same time cases that I would label as 'echoic' or 'ironic' together with cases that would be merely 'attributive'. The same goes for Recanati's approach: He re-analyzes Ducrot's notion of "polyphony" (Ducrot, 1984) in terms of locutionary and illocutionary contexts, but within the notion of "polyphony" he does not systematically distinguish what he calls "*écho*" (attributive uses) from "*ironie*" (echoic uses). However, even though the labels and the terminology introduced by Relevance Theory are more convenient for my purposes, it is worth noting that the account of Appropriation and Variability that I put forward seems at first sight compatible with both theoretical framework, Recanati's and Harris and Potts'.

Moreover, with respect to Bianchi's account of Appropriation, my proposal (i) extends the echoic approach to the case of Variability of thick terms, (ii) explains the non-echoic non-offensive uses of slurs by distinguishing cases of echo from cases of mere attribution, (iii) suggests a link between the relevance-theoretic notion of "attributive" and "echoic" uses of language and other notions tackling the same phenomenon.

In conclusion, my claim is that Variability and Appropriation do not constitute a case of meaning instability. They are just an instance of a more general phenomenon – shifting perspectives – that should be further investigated. But I shall leave that for future work.

¹⁰³ Capone (2014) talks about "polyphonic game" in indirect reports to talk about the ability of distinguishing the original speaker's voice and the indirect reporter's. In this paper, the notion of polyphony *à la* Ducrot/Recanati concerns cases that do not necessarily involve indirect report.

Conclusion

This work aims to develop a satisfactory account of slurs and thick terms, by adopting an interdisciplinary stance that mixes linguistics, metaethics and philosophy of language. The central claim of this thesis, developed in Part I, is that slurs and thick terms belong to the same class of terms that I call ‘hybrid evaluatives’ that rely on the same mechanisms: Both slurs and thick terms pick out a certain descriptive property and at the same time trigger an evaluation on that content. For example, ‘wop’ means ‘Italian’ but at the same time it triggers the presupposition that Italians are bad because of being Italians; ‘lewd’ means ‘sexually explicit beyond conventional boundaries’ and triggers the presupposition that things or individuals that are sexually explicit beyond conventional boundaries are bad because of being so and so on. I argued that the phenomenal differences between slurs and thick terms are superficial and do not constitute reasons to treat them differently. Part II was dedicated to the ‘negative’ contribution of my work: I presented the alternative theories that have been put forward to account for slurs and thick terms respectively, discussed their shortcomings and showed the reasons to prefer my presuppositional approach (Part II). In addition to the general question as to how slurs and thick terms are best analyzed (assessed in Part I and II), in Part III I focused on a more specific issue, namely how to treat Appropriation and Variability, particular uses of slurs and thick terms respectively, whose analysis is quite controversial. I proposed a uniform account based on the notions of ‘attributive’ and ‘echoic’ uses of language. The proposal, designed to account for Appropriation and Variability *within* my presuppositional theory of hybrid evaluatives, is actually compatible with possibly any account of slurs and thick terms and should be considered as a contribution to the investigation on evaluatives independently from my presuppositional framework.

The issues that I leave open concern metaethical questions as to how to define basic evaluative notions such as ‘good’ and ‘bad’, philosophical controversies such as what kind of concepts hybrid evaluatives express, together with linguistic issues such as how to best characterize the distinction between presupposition and conventional implicature and the like. However, I hope that this work, while focusing on two specific types of evaluative terms, contributes to the larger project of investigating the evaluative side of language.

Appendix

English slurs

Boche: old fashioned slur targeting German people.

Carpet-muncher: slur targeting homosexual women.

Chink: slur targeting Chinese people and Asian people in general.

Coon: slur targeting black people.

Curry-muncher: slur targeting Indian people.

Dyke: slur targeting homosexual women.

Fag: slur targeting homosexual men.

Faggot: slur targeting homosexual men.

Fairy: slur targeting homosexual men that engage in behavior that is taken to be stereotypically feminine.

Four-eyes: slur targeting people who wear glasses.

Frog: slur targeting French people.

Gook: slur targeting all native people in the regions occupied by the US army (Costa Ricans, Filipinos, Koreans, Nicaraguans, etc.).

Handicapped: slur targeting people with disabilities.

Honky: slur targeting Caucasian people.

JAP (Jewish American Princess): slur targeting wealthy young Jewish women.

Jap: slur targeting Japanese people.

Kike: slur targeting Jewish people.

Kraut: slur targeting German people.

Limey: slur targeting English people.

Negro: slur targeting black people.

Nigga: variation of the slur 'nigger', typically used in contexts of Appropriation.

Nigger: slur targeting black people.

Paki: slur targeting Pakistani people.

Queer: slur targeting homosexual people.

Rice queen: slur targeting homosexual men particularly attracted to East Asian men.

Slanty-eyed: slur targeting Asian people.

Spaghetti-eater: slur targeting Italian people.

Spade: slur targeting black people.

Stinkpotter: slur targeting the owners of motorized boats.

Twink: slur targeting homosexual men in their late teens or early twenties.

Wog (British English): slur targeting black people or Asian people.

Wog (Australian English): slur targeting Eastern and Southern European people or Asian people.

Wop: slur targeting Italian people.

Neutral terms *used as* slurs in English:

Firewood: targeting Jewish people.

Gorilla: targeting black people.

Lamp shade: targeting Jewish people.

English fictional slurs

Fangs: slur targeting vampires.

Furface: slur targeting werewolves.

Moondog: slur targeting werewolves.

Pointyear: slur targeting elves.

Toaster: slur targeting robots.

Italian slurs

Finocchio: slur targeting homosexual men.

Negro: slur targeting black people.

Negretto: slur targeting black people (especially children).

Neutral terms *used as* slurs in Italian:

Ciocolatino: targeting black people (especially children). It literally means 'chocolate praline'.

French slurs

Facho: slur targeting fascist people.

Nègre: slur targeting black people.

Summary

Keywords: slurs, thick terms, presuppositions, evaluatives, semantics, pragmatics

In the last few decades, philosophers of language and linguists have turned their attention to the evaluative and expressive dimensions of language. ‘Evaluative’ and ‘expressive’ are quite inclusive labels, ranging from expressive intensifiers like ‘damn’ to slurs like ‘wop’ and interjections like ‘shit’, from thick terms like ‘lewd’ to thin ones like ‘good’. In this dissertation I almost exclusively discuss two classes of terms for which I put forward a uniform account: slurs and thick terms.

The study of slurs and thick terms has been mostly conducted in two different – although related – fields: philosophy of language and linguistics for slurs and ethics and metaethics for thick terms. Despite the obvious relation between these disciplines, very few scholars adopted an interdisciplinary stance. On the one hand, the literature on thick terms focuses on issues related to philosophically central issues, such as the cognitivism/non-cognitivism dispute (i.e. whether evaluative judgements express cognitive or non-cognitive content, which affects the further issue of whether evaluative judgements are truth-apt), the fact/value distinction and so on. On the other hand, the debate on slurs focuses on the linguistic properties of epithets such as the projection power and it investigates the linguistic relation between descriptive and pejorative content.

Väyrynen (2009, 2011, 2012, 2013) is among the first scholars to systematically apply the tools of linguistics and philosophy of language to the study of thick terms, in order to investigate the relation between description and evaluation. Väyrynen’s series of papers together with his book play a crucial role in the discipline as an example of method. Not only do they make the study of thick terms accessible to linguists and philosophers of language, but they provide the basis for the *possibility* to ask whether slurs and thick terms rely on the same linguistic mechanisms or not.

In this thesis, I develop a uniform account of slurs and thick terms, classes of terms that I shall call ‘hybrid evaluatives’. My work aims to fill the gap between the research on thick terms conducted in metaethics on the one hand and the research on slurs conducted in linguistics and philosophy of language on the other, by showing

that the mechanism underlying slurs and thick terms is one and the same and that the phenomenal differences that we can observe depend on the peculiarities of their descriptive content. My investigation leaves aside the debate on thick concepts and pejorative concepts and focuses on thick terms and slurs instead: If one wants to infer claims about thick and pejorative concepts from my work on thick terms and slurs, one can do so only under the assumption that the meaning of thick terms and slurs amounts to what the corresponding concept expresses. I do not want to commit myself to this simplified view according to which there is a one-to-one relation between language and concepts: I will just stay neutral with respect to the relation between hybrid terms on the one hand and pejorative and thick concepts on the other hand and focus on language.

The dissertation consists of three parts. Part I presents a novel presuppositional analysis of hybrid evaluatives, a category including slurs and thick terms. After some preliminaries about how to define hybrid evaluatives and how to distinguish them from related notions (chapter 1), I focus on their presuppositional behavior with respect to projection and rejection (chapter 2). In chapter 3, I sketch a semantics for hybrid evaluatives in which I analyze their evaluative and their descriptive content respectively. I assess the issues of reference and extension and spell out what is exactly presupposed by hybrid evaluatives. I discuss how to interpret the presuppositional failure in the case of evaluative presuppositions and show how these topics relate to highly debated questions about values in ethics. In chapter 4, I discuss the effects of the *use* of hybrid evaluatives in conversation, while investigating their function: I show that the two questions as to how slurs and thick terms encode evaluation and the question as to what effects slurs and thick terms produce are deeply related. I discuss how the complex semantics of hybrid evaluatives brings about interesting issues concerning negotiations and disagreement. Chapter 5 is dedicated to the main objections that have been raised against presuppositional accounts of slurs and of thick terms respectively. I show how these objections do not constitute an obstacle to the presuppositional approach. In chapter 6 I discuss the residual differences between slurs and thick terms. I compare certain aspects of the semantics of thick terms and slurs respectively and their projective behavior and show how the features of the semantics of slurs and thick terms explain the

phenomenal divergences. In a nutshell, my analysis of hybrid evaluatives, as I presented it in Part I, is the following: both slurs and thick terms pick out a certain descriptive property and at the same time trigger an evaluation on that content. For example, ‘wop’ means ‘Italian’ but at the same time triggers the presupposition that Italians are bad because of being Italians; ‘lewd’ means ‘sexually explicit beyond conventional boundaries’ and triggers the presupposition that things or individuals that are sexually explicit beyond conventional boundaries are bad because of being so and so on. Slurs have a stronger projective power, fact that can be explained by appealing to some features of their descriptive content and to some extra-linguistic factors.

Part II is a critical review of some alternative theories that have been put forward to account for slurs and thick terms respectively. The two separate histories of these terms make it the case that the theories I consider often deal with only one class and neglect the other; as anticipated, the linguistic analysis of slurs is generally more detailed and developed than the one on thick terms. The first four chapters are dedicated to the “single-source” approaches, as Jeshion (2016) calls them: the truth-conditional approach (chapters 7-8) and the deflationary approach (chapters 9-10). In chapter 7 I discuss the truth-conditional analysis of slurs, by presenting the proposal of Hom (2008, 2010, 2012) and Hom and May (2013, 2014, forthcoming); in chapter 8, I discuss the truth-conditional analysis of thick terms formulated by Kyle (2013). The main problem with the truth-conditional approach is to account for the projection of the evaluative content. I show that the strategies employed to explain projection without appealing to non-truth-conditional meaning fail to account for the phenomena. In particular, I (i) show that Hom and May’s data in support of a non-pejorative readings of embedded slurs actually stem from metalinguistic readings and (ii) show that pragmatic mechanisms (both Hom and May’s offense and Kyle’s conversational implicatures) cannot explain projection. Therefore, the truth-conditional analysis of hybrid evaluatives fails to explain the basic data and should be abandoned. However, I endorse the following claim of the truth-conditional accounts of hybrid evaluatives: The evaluative content is linguistically encoded in the meanings of such terms and it does not arise only on the basis of pragmatic mechanisms. Chapters 9-10 are dedicated to the deflationary account of slurs and

thick terms, that analyze the evaluative content associated with hybrid evaluatives in terms of pragmatic implications. In chapter 9, dedicated to slurs, I present and discuss the so-called deflationary accounts of slurs (following the label that Anderson and Lepore 2013a gave to their own approach), according to which slurs have a fairly unexceptional semantics. I include in this category all those accounts according to which the derogatory content associated with slurs is not part of their encoded meaning (again, broadly intended: not just truth conditions, but also conventional implicatures, presuppositions and the like). I present three main approaches developed by Anderson and Lepore (2013a, 2013b), Bolinger (2015) and Nunberg (forthcoming), focusing on three notions respectively: taboo, contrastive preference and manner implicatures. I will discuss the specific problems met by each account, but in general I show that the deflationary accounts of slurs meet some difficulties in explaining the genesis of slurs, the complicity phenomenon (i.e. the fact that the pejorative content associated with slurs does not seem to concern the speaker *only*) and the contrast between the slurs and other ‘affiliatory’ terms. Chapter 10 is dedicated to the deflationary account of thick terms, developed by Pekka Väyrynen, according to which the evaluation associated with thick terms (*T-evaluation* in Väyrynen’s terms) conversationally arises as a pragmatic implication. These implications are not part of the asserted content, nor they are the main point of the utterance in their literal uses in normal contexts: in other words, they are typically backgrounded. The pragmatic account has the advantage of being more parsimonious than any other account of thick terms, but it fails to account for some instances of projection of the evaluative content (the so-called *defeasibility* data). In sum, in the first four chapters of Part II, I show that the single-source approaches fail to properly account for the behavior of hybrid evaluatives and I argue that the difficulties met by these proposals speak in favor of a hybrid approach. In chapter 11 I discuss an alternative hybrid approach, the conventional implicature view: It agrees with truth-conditional accounts that the derogatory content of epithets is linguistically encoded, but it holds with pragmatic theories that the pejorative content does not belong to the truth-conditional dimension of meaning. Such an approach, that analyzes the derogatory content of slurs in terms of conventional implicatures, was proposed by Christopher Potts (Potts 2005, 2007) and developed by Eric McCready (McCready

2010) and Daniel Gutzmann (Gutzmann 2011). It was furthermore supported *i.a.* by Copp (2009), Williamson (2009), Whiting (2013). While the previous approaches I presented have been independently developed for both slurs and thick terms, the conventional implicatures view was proposed for slurs only: The alleged defeasibility of the evaluative content of thick terms made scholars discard a conventional implicature approach. After presenting the account and showing how similar it is to the presuppositional approach (presupposition and conventional implicature being, in Abbott 2006's words "close neighbor[s] in the linguistic literature"), I point out the aspects for which the presuppositional view should be preferred.

In Part III I propose a uniform account of two phenomena that have been treated as independent so far: Appropriation of slurs and Variability of thick terms. My proposal, designed to account for Appropriation and Variability within my presuppositional theory of hybrid evaluatives, is actually compatible with possibly any account of slurs and thick terms. In chapter 12 I discuss Appropriation, a particular use of slurs, and Variability, a particular use of thick terms. In chapter 13, I propose a uniform account of Appropriation and Variability, relying on the notion of "echoic use" of language, introduced by Relevance Theory (Sperber and Wilson, 1986, Carston, 1996, 2002, Wilson, 2006a, 2006b, 2007 and Wilson and Sperber, 2012) and employed by Bianchi (2014) to analyze the appropriated uses of slurs. I will argue that such an echoic approach does not only explain the Appropriation of slurs, as Bianchi claims, but it also accounts for the Variability of thick terms. I show how the relevance-theoretic distinction between 'merely attributive' and 'echoic' uses of language sheds light on the crucial distinction between the two often conflated phenomena of 'suspension' and 'reversal' of the evaluative content of slurs and thick terms. I argue that the theoretically relevant distinction is not between Appropriation and Variability, rather between 'suspension' and 'reversal' of the evaluative content. In chapter 14, I discuss the theoretical outcomes of my proposal for the theory of thick terms: if Variability cases are in fact derivative or parasitic uses of language, they do not count anymore as a counterexample to the claim that thick terms convey or express evaluation in a systematic way. This suggests once more to consider slurs and thick terms as similar phenomena. Finally, I conclude by

suggesting that the relevance-theoretic notions of “attributive” and “echoic” uses of language seem to be related to similar notions such as “perspectival shift” (Harris and Potts, 2009a, 2009b and Harris, 2012) and “polyphony” (Ducrot, 1984 and Recanati, 2006). The phenomenon at stake – namely the possibility to attribute a certain content to someone else, while possibly expressing one’s attitude towards it – is a very general feature of human language that goes beyond evaluative terms and deserves deeper investigation.

The picture that emerges from this work is that hybrid evaluatives represent a device through which language implicitly conveys linguistically encoded evaluations. In my account, hybrid evaluatives rely on presuppositions, which are – in Chilton (2004)’s words – “at least one micro-mechanism in language use which contributes to the building of a consensual reality”. By employing these terms, we implicitly take for granted a certain moral perspective, a certain set of beliefs concerning what is good and what is bad. We implicitly apply a certain lens to the world and expect everyone else to do the same. Because the presupposed content is presented as not open to discussion, if it is not objected it has the potential to shape contexts. In this sense, using hybrid evaluatives is a powerful tool through which language not only encodes evaluation, but is also able to impose it. Talking about the stereotypes evoked by slurs, Nunberg (forthcoming) appeals to the notion of “shortcut”: “Stereotypes, negative and positive, are among the cognitive shortcuts we rely on to make sense of the world and to guide our responses to it”. I argue that this is true not just for stereotypes, but for hybrid evaluatives in general, as they are devices through which language can convey evaluations in a way that is both linguistically encoded and implicit.

Résumé

Mots clés : termes éthiques épais, termes d'offense, présuppositions, Appropriation, Variabilité, Termes d'évaluation, Sémantique, Pragmatique

Au cours des dernières décennies, les philosophes du langage et les linguistes se sont intéressés à la dimension expressive et évaluative du langage. « Évaluatif » et « expressif » sont des catégories très larges et incluent différentes espèces de termes, par exemple des termes d'offense comme « rituel », des interjections comme « merde ! », des termes éthiques épais comme « obscène » et des termes éthiques fins comme « bon ». Dans ce travail de thèse je vais me focaliser de façon quasiment exclusive sur deux espèces d'expressions, pour lesquelles je propose une analyse uniforme en termes de présuppositions : les deux espèces sont les termes d'offense et les termes éthiques épais. Je prends comme point de départ pour mon enquête des champs d'études différents, mais connectés, la philosophie du langage et la linguistique pour les termes d'offense, l'éthique et la meta-éthique pour les termes éthiques épais. Bien que ces champs de recherches soient connectés, très peu de chercheurs ont adopté une perspective réellement interdisciplinaire. D'un côté, la littérature sur les termes éthiques épais s'est focalisée sur des questions centrales pour le débat philosophique, comme l'opposition entre cognitivisme et non-cognitivisme (c'est-à-dire la question de savoir si les jugements évaluatifs expriment des contenus cognitifs ou non-cognitifs, débat qui soulève d'autres problématiques concernant la possibilité pour un jugement évaluatif d'avoir une valeur de vérité), la distinction entre faits et valeurs, et ainsi de suite. De l'autre côté, le débat sur les termes d'offense s'est focalisé sur les propriétés linguistiques de ces expressions, par exemple leur pouvoir de projection, et sur la relation entre le contenu descriptif et le contenu péjoratif.

Väyrynen (2009, 2011, 2012, 2013) a été parmi les premiers à avoir appliqué de façon systématique les outils de la philosophie du langage et de la linguistique à l'étude des termes éthiques épais afin de clarifier la relation entre description et évaluation. Les contributions de Väyrynen jouent un rôle central dans le domaine en s'imposant comme une référence sur le plan méthodologique. La méthodologie de ces travaux rend l'étude des termes éthiques épais accessible aux linguistes et aux

philosophes du langage, mais elle offre aussi les bases pour se poser la question du fonctionnement des termes d'offense et des termes éthiques épais comme étant fondé sur les mêmes mécanismes linguistiques.

Dans ma thèse, je développe une analyse uniforme des termes d'offense et des termes éthiques épais, classes de termes que j'appelle les « évaluatifs hybrides ». Le but de mon travail est de combler (ou de contribuer à combler) la lacune entre la recherche sur les termes éthiques épais menée principalement en meta-éthique, et la recherche sur les termes d'offense, menée principalement en linguistique et en philosophie du langage. Mon but est de démontrer que le mécanisme à la base des termes d'offense et des termes éthiques épais est le même et que les différences superficielles entre ces catégories dépendent de certaines caractéristiques du contenu descriptif de ces termes.

Ma recherche fait passer au second plan le débat sur les concepts éthiques épais et sur les concepts d'offense et se concentre sur le langage, c'est-à-dire sur les termes éthiques épais et sur les termes d'offense. Si l'on voulait passer directement des expressions linguistiques aux concepts, ou en d'autres mots, si l'on voulait appliquer aux concepts d'offense et aux concepts éthiques épais les mêmes réflexions que je propose ici pour les termes d'offense et les termes éthiques épais, il faudrait supposer que la signification de ces expressions ne soit rien d'autre que ce qu'expriment les concepts correspondants. Je ne veux pas m'engager dans une telle hypothèse qui traite de façon trop simplifiée la relation complexe entre entités linguistiques et entités conceptuelles comme une correspondance un à un. Par conséquent je me concentre exclusivement sur le langage, en restant neutre sur la relation entre concepts évaluatifs et expressions évaluatives.

La structure de la thèse est la suivante. Dans la première partie je présente une nouvelle analyse présuppositionnelle des évaluatifs hybrides, catégorie qui inclut les termes d'offense et les termes éthiques épais. J'introduis d'abord quelques notions préliminaires, je définis les évaluatifs hybrides et je les distingue d'autres notions voisines (chapitre 1), ensuite je me concentre sur le comportement présuppositionnel des évaluatifs hybrides, en discutant les phénomènes de la projection et du rejet (chapitre 2). Dans le chapitre 3 j'esquisse une sémantique pour les évaluatifs hybrides en analysant séparément leur contenu descriptif de leur contenu évaluatif. Je

prends en compte des questions comme celles de la référence et de l'extension et je cherche à élaborer la meilleure formulation possible pour comprendre le contenu présupposé des évaluatifs hybrides. Je discute de comment interpréter l'échec présuppositionnel dans le cas des présuppositions évaluatives (plutôt que descriptives) et je montre comment la question de cet échec est profondément liée aux questionnements et aux problématiques qui concernent la théorie des valeurs en éthique et en meta-éthique. Dans le chapitre 4 j'observe les effets générés par l'usage des évaluatifs hybrides dans les conversations, en explorant leur fonction : je montre que la question linguistique de la façon dont les termes d'offense et les termes éthiques épais *codifient* un contenu évaluatif est profondément liée à la question non linguistique de quels sont les *effets* produits par les évaluatifs hybrides et quelle fonction ils exécutent.

En outre, j'avance des considérations sur comment la sémantique complexe et riche des évaluatifs hybrides détermine des questions intéressantes qui concernent les concepts tels que la négociation et le désaccord. Le chapitre 5 est dédié aux objections qui ont été soulevées contre l'analyse présuppositionnelle des termes d'offense et des termes éthiques épais. Mon but est de montrer que ces objections ne constituent pas de véritables obstacles pour l'approche présuppositionnelle. Dans le chapitre 6 je prends en considération les différences restantes entre les termes d'offense et les termes éthiques épais : je compare certains aspects de leur sémantique avec leur comportement projectif et je montre que les caractéristiques des évaluatifs que j'ai mises en lumière peuvent expliquer les différences superficielles. En d'autres termes, mon analyse des évaluatifs hybrides, présentée dans la Partie I, est la suivante : à la fois les termes d'offense et les termes éthiques épais identifient des propriétés descriptives, mais en même temps présupposent un contenu évaluatif *qui concerne* ce contenu descriptif. Par exemple, « rital » veut dire « italien », mais en même temps présuppose que les italiens sont méprisables en tant qu'italiens, « obscène » signifie quelque chose comme « sexuellement explicite », mais en même temps présuppose que les choses et les personnes sexuellement explicites sont méprisables en tant que telles, et ainsi de suite. Les termes d'offense possèdent un pouvoir projectif plus grand que les termes éthiques épais. Ceci est

expliqué par des propriétés du contenu descriptif des termes d'offense et par des facteurs extralinguistiques.

La Partie II est une présentation critique de certaines théories alternatives proposées pour expliquer les termes d'offense et les termes éthiques épais. Les deux histoires séparées de ces analyses ont eu comme conséquence le fait que les théories en considération se sont concentrées sur une seule classe d'expressions et ont ignoré l'autre. Comme j'ai déjà mentionné, l'étude des termes d'offense par les linguistes et les philosophes du langage est en général plus développée et approfondie que l'analyse linguistique des termes éthiques épais.

Les premiers chapitres de la Partie II sont dédiés aux approches « monosource » (selon la définition de Jeshion 2016) : les théories vériconditionnelles d'un côté (chapitres 7-8) et celles déflationnistes de l'autre (chapitres 9-10). Dans le chapitre 7 je présente et discute l'analyse vériconditionnelle des termes d'offense, en me focalisant sur la théorie proposée par Hom (2008, 2010, 2012) et Hom et May (2013, 2014, à paraître) ; dans le chapitre 8 je discute l'analyse vériconditionnelle des termes éthiques épais formulée par Kyle (2013). Les difficultés principales des approches vériconditionnelles sont dues au comportement projectif des évaluatifs hybrides. Mon but est de montrer que les stratégies qui tentent d'expliquer la projection sans faire appel aux composants *non* vériconditionnels de la signification ne parviennent pas à expliquer les phénomènes en jeu. En particulier, (i) je démontre que les données utilisées par Hom et May pour soutenir l'idée que les termes enchâssés ne sont pas péjoratifs dérivent en réalité des usages métalinguistiques des termes et (ii) je montre que les mécanismes pragmatiques invoqués par Hom et May (comme la notion d'« insulte ») et par Kyle (comme la notion d'« implicature conversationnelle ») ne parviennent pas à rendre compte du comportement projectif.

Etant donné que l'analyse vériconditionnelle n'arrive pas à expliquer les aspects fondamentaux des évaluatifs hybrides, elle doit être abandonnée. Cependant, je retiens de l'analyse vériconditionnelle l'idée que le contenu évaluatif des termes d'offense et des termes éthiques épais est codifié linguistiquement et non pas généré dans la conversation à partir des mécanismes pragmatiques. Les chapitres 9 et 10 sont en effet dédiés aux analyses pragmatiques et déflationnistes des termes d'offense et des termes éthiques épais. Ces théories rendent compte du contenu

évaluatif associé à ces expressions en termes d'implications pragmatiques. Dans le chapitre 9, dédié aux termes d'offense, j'introduis et je discute les analyses déflationnistes (en adoptant la définition que Anderson et Lepore 2013a appliquent à leur propre théorie). Selon ces approches, les termes d'offense ont une sémantique assez courante, et la signification de ces expressions n'inclut pas un contenu évaluatif. J'inclus dans la catégorie des théories déflationnistes toutes ces théories qui soutiennent que le contenu péjoratif des termes d'offense n'appartient pas à la signification linguistiquement codifiée (signification est à comprendre au sens large du terme : elle inclut les conditions de vérité, mais aussi les implicatures conversationnelles, les présuppositions, etc.). J'introduis trois théories principales, développées par Anderson et Lepore (2013a, 2013b), Bolinger (2015) et Nunberg (à paraître), qui se focalisent sur trois notions clés : tabou, préférence contrastive (« contrastive preference ») et implicatures conversationnelles de manière. Je discute les problèmes spécifiques associés avec chacune de ces théories et je démontre qu'en général les théories déflationnistes des termes d'offense n'expliquent pas de façon satisfaisante l'origine des termes, le phénomène de la complicité (c'est-à-dire le fait que le contenu péjoratif associé aux termes d'offense n'est pas attribué seulement à celui ou à celle qui prononce le terme) et le contraste entre les termes d'offense et les autres mots de code ou mots 'associants'. Le chapitre 10 est dédié à la théorie déflationniste des termes éthiques épais, introduite et développée par Pekka Väyrynen, selon laquelle le contenu évaluatif associé aux termes épais (« T-évaluation » selon la définition de Väyrynen) émerge dans la conversation comme implication (et *non pas* implicature) pragmatique. Ces implications ne font pas partie de ce qui est affirmé, et ne sont pas la raison principale de l'énonciation dans les usages littéraux dans les contextes standards, mais ce sont normalement des contenus de « deuxième plan ». La théorie pragmatique a l'avantage d'être plus parcimonieuse que tout autre théorie concernant les termes épais, mais elle n'explique pas de façon satisfaisante certains cas de comportement projectif du contenu évaluatif (ce que l'on appelle les exemples de neutralisabilité (« neutralizability »). En bref, dans les quatre premiers chapitres de la seconde partie je montre que les approches « mono-source » (vériconditionnelles et déflationnistes) ne permettent pas de rendre compte de façon satisfaisante du comportement des évaluatifs hybrides et que les difficultés

rencontrées par ces approches suggèrent qu’une approche hybride est préférable. Le chapitre 11 est donc dédié à l’analyse d’une approche hybride alternative à la mienne, c’est-à-dire une analyse des termes en utilisant la notion d’implicature conventionnelle. Une telle théorie partage avec les approches vériconditionnelles l’idée que le contenu péjoratif des termes est linguistiquement codifié, mais soutient aussi – tout comme les théories déflationnistes – que le contenu péjoratif n’appartient pas à la signification vériconditionnelle de ces expressions. Cette théorie a été proposée par Christopher Potts (Potts 2005, 2007) et développée ensuite par McCreedy (2010) et Gutzmann (2011). Elle a été soutenue, entre autres, par Copp (2009), Williamson (2009) et Whiting (2013). Tandis que les approches discutées dans les chapitres précédents ont été développées *indépendamment* mais *parallèlement* pour rendre compte des termes d’offense ou des termes épais, la stratégie des implicatures conventionnelles a été proposée seulement pour les termes d’offense : la possibilité de neutraliser dans certains contextes le contenu évaluatif des termes éthiques épais a poussé les philosophes et les linguistes à écarter une approche en termes d’implicatures conventionnelles pour les termes éthiques épais. Après avoir présenté la théorie et avoir montré les ressemblances entre cette théorie et une approche présuppositionnelle (ressemblance peu étonnante si l’on considère que les implicatures conventionnelles et les présuppositions sont des « notions voisines dans la littérature linguistique »¹⁰⁴, comme suggère Abbott 2006), je mets en lumière les raisons pour lesquelles une théorie présuppositionnelle est préférable.

Dans la Partie III, je suggère une analyse uniforme de deux phénomènes qui ont été considérés jusqu’au présent comme indépendants : l’Appropriation des termes d’offense et la Variabilité des termes épais. La suggestion, développée au départ pour expliquer l’Appropriation et la Variabilité dans le contexte d’une approche présuppositionnelle, est en réalité potentiellement compatible avec n’importe quelle théorie des termes d’offense et des termes épais.

Dans le chapitre 12 j’introduis le phénomène de l’Appropriation, un usage particulier des termes d’offense, et le phénomène de la Variabilité, un usage particulier des termes éthiques épais. Dans le chapitre 13 je propose une analyse

¹⁰⁴ Abbott (2006:2), ma traduction.

uniforme de l'Appropriation et de la Variabilité, en me référant à l'usage échoïque du langage, notion introduite par les théoriciens de la pertinence (Sperber et Wilson, 1986, Carston, 1996, 2002, Wilson, 2006a, 2006b, 2007 et Wilson et Sperber, 2012) et adoptée par Bianchi (2014) pour analyser les usages d'appropriation des termes d'offense. Mon but est de montrer que l'approche échoïque est non seulement capable d'expliquer le phénomène de l'Appropriation des termes d'offense (comme soutient Bianchi), mais aussi capable de rendre compte du phénomène de la Variabilité des termes éthiques épais. En outre, je montre que la distinction introduite par la théorie de la pertinence entre usages purement attributifs et usages échoïques du langage clarifie aussi la distinction centrale entre deux phénomènes trop souvent rapprochés : la suspension et le renversement du contenu évaluatif des termes d'offense et des termes éthiques épais. Je soutiens qu'il n'y a pas une distinction importante entre Appropriation et Variabilité, mais plutôt une distinction entre suspension et renversement du contenu évaluatif. Dans le chapitre 14, je développe les conséquences théoriques de ma suggestion pour la théorie des termes éthiques épais : si les exemples de Variabilité sont effectivement analysables en tant qu'usages dérivatifs et secondaires du langage, ceux-ci ne sont pas à considérer comme contrexemples à la thèse selon laquelle les termes éthiques épais expriment de façon systématique une évaluation. Ce point fournit les bases pour un rapprochement entre termes d'offense et termes éthiques épais comme phénomènes similaires. Enfin, je conclus en suggérant que les notions introduites par la théorie de la pertinence, comme l'usage attributif et échoïque du langage ont des affinités intéressantes avec des notions similaires, comme la notion de 'changement de perspective' (*perspectival shift*), introduite par Harris et Potts (2009a, 2009b) et Harris (2012), et la notion de polyphonie, introduite par Ducrot (1984) et Recanati (2006). Le phénomène en question – la possibilité d'attribuer un contenu à quelqu'un d'autre qu'au locuteur et en même temps le fait de pouvoir exprimer sa propre attitude envers le contenu lui-même – est une caractéristique des langages naturels, plus large que le cas spécifique des termes évaluatifs et qui mériterait un traitement à part.

Pour conclure, la thèse générale de mon travail est que les évaluatifs hybrides représentent un outil à travers lequel le langage transmet implicitement des évaluations linguistiquement codifiées. Dans mon analyse les évaluatifs hybrides

impliquent les présuppositions, qui sont « au moins un des micro-mécanismes de l'usage du langage qui contribuent à la création d'une réalité consensuelle », pour utiliser des mots de Chilton (2004).¹⁰⁵ En utilisant ces expressions, nous adhérons implicitement à une certaine perspective morale, à un ensemble de croyances qui concernent le bien et le mal, le juste et l'injuste. Nous appliquons implicitement au monde un filtre, une vision, et nous nous attendons à ce que les autres fassent pareil. Du moment que les choses qui sont présupposées se présentent comme données et non pas modifiables, un tel contenu a le pouvoir de façonner les contenus, à moins que les locuteurs présents ne s'y opposent explicitement. En ce sens, l'usage des évaluatifs hybrides est un outil puissant, qui pas seulement encode un contenu évaluatif, mais est aussi capable, dans une certaine mesure, de l'imposer. En parlant des stéréotypes évoqués par les termes d'offense, Nunberg (à paraître) utilise la notion de « raccourci » : « Les stéréotypes, négatifs ou positifs, sont des raccourcis cognitifs que nous utilisons pour donner un sens au monde et pour guider nos réponses au monde »¹⁰⁶. Dans mon travail je soutiens l'idée que les considérations de Nunberg s'appliquent non seulement aux stéréotypes, mais aussi aux évaluatifs hybrides en général, du moment qu'ils sont des dispositifs employés par le langage pour transmettre des évaluations d'une manière à la fois linguistiquement codifiée et implicite.

¹⁰⁵Chilton (2004: 64), ma traduction.

¹⁰⁶Nunberg (à paraître), ma traduction.

References

Abbott, Barbara (2006), Where have some of the presuppositions gone, in Birner, Betty and Ward, Gregory (eds.), *Drawing the boundaries of meaning: Neo-Gricean studies in pragmatics and semantics in honor of Laurence R. Horn*, Philadelphia, John Benjamins, pp. 1–20.

Abrusán, Márta (2016), Presupposition Cancellation: Explaining the Soft-Hard Distinction, *Natural Language Semantics*, 24 (2), pp. 165–202.

Abusch, Dorit (2002), Lexical alternatives as a source of pragmatic presuppositions, in Jackson, Brendan (ed.), *Proceedings of SALT 12*, Ithaca (NY), CLC Publications, pp. 1–20.

Abusch, Dorit (2010), Presupposition triggering from alternatives, *Journal of Semantics*, 27 (1), pp. 37–80.

Amaral, Patricia, Roberts, Craige and Smith, Allyn (2007), Review of The Logic of Conventional Implicatures by Chris Potts, *Linguistics and Philosophy*, 30, pp. 707–749.

Anderson, Luvell, Haslanger, Sally and Langton, Rae (2012), Language and Race, in Fara, Delia and Russell, Gillian (eds.), *Routledge Companion to the Philosophy of Language*, London, Routledge, pp. 753–767.

Anderson, Luvell and Lepore, Ernest (2013a), Slurring words, *Nous*, 47 (1), pp. 25–48.

Anderson, Luvell and Lepore, Ernest (2013b), What did you call me? Slurs as prohibited words, *Analytic Philosophy*, 54 (3), pp. 350–363.

Beaver, David I. and Geurts, Bart (2014), Presupposition, in Zalta, Edward N. (ed.), *The Stanford Encyclopedia of Philosophy*, URL=<<http://plato.stanford.edu/archives/win2014/entries/presupposition/>>.

Bergström, Lars (2002), Putnam on the Fact-Value Dichotomy, *Croatian Journal of Philosophy*, 2, pp. 1–13.

Bianchi, Claudia (2014), Slurs and appropriation: An echoic account, *Journal of Pragmatics*, 66, pp. 35–44.

Bierwisch, Manfred (1989), The semantics of gradation, in Bierwisch, Manfred and Lang, Ewald (eds.), *Dimensional Adjectives*, Berlin, Springer-Verlag, pp. 71–261.

Blackburn, Simon (1992), Through Thick and Thin, *Proceedings of the Aristotelian Society*, Supplementary Volume 66, pp. 284–99.

Blackburn, Simon (1998), *Ruling Passions*, Oxford, Clarendon Press.

Bolinger Jorgensen, Renée (2015), The Pragmatics of Slurs, *Noûs*, 50 (3).

Brentano, Franz (1969) [1889], *The Origin of Our Knowledge of Right and Wrong*, London, Routledge.

Brink, David Owen (1989), *Moral realism and the foundations of ethics*, Cambridge, Cambridge University Press.

Brontsema, Robin (2004), A queer revolution: reconceptualizing the debate over linguistic reclamation, *Colorado Research in Linguistics*, 17, pp. 1–17.

Brower, Bruce W. (1988), Virtue Concepts and Ethical Realism, *The Journal of Philosophy*, 85, pp. 675–693.

Burgess, Alexis, and David Plunkett (2013a), Conceptual Ethics I, *Philosophy Compass*, 8, pp. 1091–1101.

Burgess, Alexis, and David Plunkett (2013b), Conceptual Ethics II, *Philosophy Compass*, 8, pp. 1102–1110.

Bylinina, Lisa (2014), *The Grammar of Standards Judge-dependence, Purpose-relativity, and Comparison Classes in Degree Constructions*, Utrecht, LOT.

Camp, Elizabeth (2013), Slurring Perspectives, *Analytic Philosophy*, 54 (3), pp. 330–349.

Campbell, Richmond (2015), Moral Epistemology, in Zalta, Edward N. (ed.), *The Stanford Encyclopedia of Philosophy*, URL = <http://plato.stanford.edu/archives/win2015/entries/moral-epistemology/>.

Capone, Alessandro (2014), Indirect reports, slurs, and the polyphonic speaker, *Reti, saperi, linguaggi*, 1 (2), pp. 299–318.

Carston, Robin (1996), Metalinguistic negation and echoic use, *Journal of Pragmatics*, 25, pp. 309–330.

Carston, Robin (2002), *Thoughts and Utterances: The Pragmatics of Explicit Communication*, Oxford, Blackwell.

Cepollaro, Bianca (2015), In Defense Of A Presuppositional Account Of Slurs, *Language Sciences*, 52, pp. 36–45.

Cepollaro, Bianca (2016), Building evaluation into Language, *Phenomenology and Mind*, 11, pp. 158–168.

Cepollaro, Bianca (forthcoming), The shortcut of discrimination, *Rivista di Estetica*, forthcoming.

Cepollaro, Bianca and Stojanovic, Isidora (2016), Hybrid Evaluatives, *Grazer Philosophische Studien*, 93 (3), pp. 458–488.

Cepollaro, Bianca and Thommen, Tristan (ms), What's wrong with truth-conditional account of slurs, submitted, under review.

Chemla, Emmanuel and Schlenker, Philippe (2012), Incremental vs. Symmetric Accounts of Presupposition Projection: An experimental Approach, *Natural language semantics*, 20 (2), pp. 177–226.

Chilton, Paul (2004), *Analysing Political Discourse: Theory and Practice*, London and New York, Routledge.

Copp, David (2001), Realist-expressivism: A neglected option for moral realism, *Social Philosophy and Policy*, 18, pp. 1–43.

Copp, David (2009), Realist-expressivism and conventional implicature, *Oxford Studies in Meta-Ethics*, 4, pp. 167–202.

Croom, Adam (2011), Slurs, *Language Sciences*, 33, pp. 343–358.

Croom, Adam (2013), How to do things with slurs: Studies in the way of derogatory words, *Language & Communication*, 33, pp. 177–204.

Croom, Adam (2014), Spanish slurs and stereotypes for Mexican-Americans in the USA: A context-sensitive account of derogation and appropriation, *Sociocultural Pragmatics*, 2, pp. 1–35.

Cruse, D. Alan (1986), *Lexical Semantics*, Cambridge, Cambridge University Press.

Culpeper, Jonathan L. (1996), Towards an anatomy of impoliteness, *Journal of Pragmatics*, 25, pp. 349–367.

Culpeper, Jonathan L. (2011), *Impoliteness: Using Language to Cause Offense*, Cambridge, Cambridge University Press.

Culpeper, Jonathan L. (2016), Impoliteness strategies, in Capone, Alessandro and Mey, Jacob L. (eds.), *Interdisciplinary Studies in Pragmatics, Culture and Society*, Springer, pp. 421–445.

Čupković, Gordana (2015), Diachronic variations of slurs and levels of derogation: On some regional, ethnic and racial slurs in Croatian, *Language Sciences*, 52, pp. 215–230.

Dancy, Jonathan (1995), In Defence of Thick Concepts, *Midwest Studies in Philosophy*, 20, pp. 263–279.

Dancy, Jonathan (2004), *Ethics without Principles*, Oxford, Clarendon Press.

Davies, Mark (2008-), *The Corpus of Contemporary American English: 520 million words, 1990-present*, available online at <http://corpus.byu.edu/coca/>

DiFranco, Ralph (2014), Pejorative Language, in *Internet Encyclopedia of Philosophy*, <http://www.iep.utm.edu/pejorati/>.

DiFranco, Ralph (2015), Do Racists Speak Truly? On the Truth-Conditional Content of Slurs, *Thought: A Journal of Philosophy*, 4 (1), pp. 28–37.

Ducrot, Oswald (1984), Esquisse d'une théorie polyphonique de l'énonciation, in

Ducrot, Oswald, *Le Dire et le Dit*, Paris, Minuit, pp. 171–233.

Dummett, Michael (1973), *Frege's Philosophy of Language*, Oxford, Clarendon Press.

Eklund, Matti (2011), What are thick terms, *Canadian Journal of Philosophy*, 41 (1), pp. 25–49.

Eklund, Matti (2013), Evaluative Language and Evaluative Reality, in Kirchin, Simon (ed.), *Thick Concepts*, Oxford, Oxford University Press, pp. 161–181.

Elstein, D. Y., and Hurka, T. (2009), From thick to thin: Two moral reduction plans, *Canadian Journal of Philosophy*, 39 (4), pp. 515–535.

Ewing, Alfred Cyril (1939), A Suggested Non-Naturalistic Analysis of Good, *Mind*, 48, pp. 1–22.

Ewing, Alfred Cyril (1947), *The Definition of Good*, London, MacMillan.

von Fintel, Kai (2004), Would you believe it? The King of France is back!, in Bezuidenhout, Anne and Reimer, Marga (eds.), *Descriptions and Beyond: An Interdisciplinary Collection of Essays on Definite and Indefinite Descriptions and other Related Phenomena*, Oxford, Oxford University Press, pp. 328–354.

von Fintel, Kai (2008), What is presupposition accommodation, again?, *Philosophical Perspectives*, 22 (1), pp. 137–170.

Fricker, Elizabeth (1987), The Epistemology of Testimony, *Proceedings of the Aristotelian Society*, Supplementary Volume 61, pp. 57–83.

Galinsky, Adam D., Hugenberg, Kurt, Groom, Carla and Bodenhausen, Galen V. (2003), The reappropriation of stigmatizing labels: Implications for social identity, In

Polzer, Jeffrey (ed.), *Identity Issues in Groups*, Bingley, Emerald Group Publishing Limited, pp. 221–256.

Galinsky Adam D., Wang, Cinthya S., Whitson, Jennifer A., Anicich, Eric M., Hugenberg, Kurt and Bodenhausen, Galen V. (2013), The reappropriation of stigmatizing labels: The reciprocal relationship between power and self-labeling, *Psychological Science*, 24, pp. 2020–2029.

Gallie, Walter Bryce (1956), Essentially Contested Concepts, *Proceedings of the Aristotelian Society*, 56, pp. 167–98.

García-Carpintero, Manuel (2015), Contexts as Shared Commitments, *Frontiers in Psychology*, 6, pp. 1–13.

Gauker, Christopher (2008), Against accommodation: Heim, van der Sandt, and the presupposition projection problem, *Philosophical Perspectives*, 22 (1), pp. 171–205.

Gazdar, Gerald (1979), *Pragmatics: Implicature, Presupposition, and Logical Form*, New York, Academic Press.

Geurts, Bart (1998), The Mechanisms of Denial, *Language*, 74, 274–307.

Gibbard, Allan (1992), Thick Concepts and Warrant For Feelings, *Proceedings of the Aristotelian Society*, Supplementary Volume 61, pp. 267–283.

Gibbard, Allan (2003a), Reasons Thin and Thick, *The Journal of Philosophy*, 100 (6), pp. 288–304.

Gibbard, Allan (2003b), *Thinking How to Live*, Cambridge (MA), Harvard University Press.

Gibbs, Raymond W. (1986), On the Psycholinguistics of Sarcasm, *Journal of Experimental Psychology: General*, 115 (1), pp. 2–15.

Glanzberg, Micheal (2007), Context, Content, and Relativism, *Philosophical Studies*, 136, pp. 1–29.

Greenberg, Joseph (1966), *Language Universals*, The Hague, Mouton.

Grice, H. Paul (1989), *Studies in the Way of Words*, Cambridge (MA), Harvard University Press.

Gutzmann, Daniel (2011), Expressive Modifiers & Mixed Expressives, in Bonami, Olivier and Cabredo Hofherr, Patricia (eds.), *Empirical Issues in Syntax and Semantics*, 8, pp. 123–141.

Gutzmann, Daniel (2015), *Use-conditional meaning: Studies in multidimensional semantics*, Oxford, Oxford University Press.

Happé, Francesca G. E. (1993), Communicative competence and theory of mind in autism: A test of relevance theory, *Cognition*, 48 (2), pp. 101–119.

Harcourt, Edward and Thomas, Alan (2013), Thick Concepts, Analysis, and Reductionism, in Kirchin, Simon (ed.), *Thick Concepts*, Oxford, Oxford University Press, pp. 20–43.

Hare, Richard M. (1952), *The Language of Morals*, Oxford, Oxford University Press.

Hare, Richard M. (1963), *Freedom and Reason*, Oxford, Oxford University Press.

Hare, Richard M. (1981), *Moral Thinking*, Oxford, Clarendon Press.

Harris, Jesse A. (2012), *Processing Perspectives*, doctoral dissertation, Amherst, University of Massachusetts.

Harris, Jesse A. and Potts, Christopher (2009a), Perspective-shifting with appositives and expressives, *Linguistics and Philosophy*, 32 (6), pp. 523–552.

Harris, Jesse A. and Potts, Christopher (2009b), Predicting perspectival orientation for appositives, in *Proceedings from the Annual Meeting of the Chicago Linguistic Society*, 45 (1), pp. 207–221.

Haslanger, Sally (2013), *Social Meaning and Philosophical Method*, APA Presidential Address, handout.

Haspelmath, Martin (2006), Against markedness (and what to replace it with), *Journal of Linguistics*, 42 (1), pp. 25–70.

Hay, Ryan J. (2013), Hybrid expressivism and the analogy between pejoratives and moral language, *European Journal of Philosophy*, 21 (3), pp. 450–474.

Hedger, Joseph A. (2012), The semantics of racial slurs: using Kaplan's framework to provide a theory of the meaning of derogatory epithets, *Linguistic and Philosophical Investigations*, 11, pp. 74–84.

Heim, Irene (1992), Presupposition projection and the semantics of attitude verbs, *Journal of semantics*, 9 (3), pp. 183–221.

Heim, Irene (2000), Degree Operators and Scope, in Jackson, Brendon and Matthews, Tanya (eds.), *Proceedings of SALT 10*, Ithaca (NY), CLC Publications, Cornell University.

Hom, Christopher (2008), The semantics of racial epithets, *Journal of Philosophy*, 105, pp. 416–440.

Hom, Christopher (2010), Pejoratives, *Philosophy Compass*, 5 (2), pp. 164–185.

Hom, Christopher (2012), A puzzle about pejoratives, *Philosophical Studies*, 159, pp. 383–405.

Hom, Christopher and May, Robert (2013), Moral and Semantic Innocence, *Analytic Philosophy*, 54 (3), pp. 293–313.

Hom, Christopher and May, Robert (2014), The inconsistency of the identity thesis, *Protosociology*, 31, pp. 113–120.

Hom, Christopher and May, Robert (forthcoming), Pejoratives as Fiction, in Sosa, David (ed.) (forthcoming), *Bad words*, Oxford, Oxford University Press.

Horn, Lawrence (1985), Metalinguistic Negation and Pragmatic Ambiguity, *Language*, 61 (1), 121–174.

Horn, Lawrence (1989), *A Natural History of Negation*, Chicago, University of Chicago Press.

Hornsby, Jennifer (2001), Meaning and Uselessness: How to think about derogatory words, *Midwest Studies in Philosophy: Figurative Language*, 25, pp. 128–141.

Jacobson, Daniel (2011), Fitting Attitude Theories of Value, in Zalta, Edward N. (ed.), *The Stanford Encyclopedia of Philosophy*,
URL=<<http://plato.stanford.edu/archives/spr2011/entries/fitting-attitude-theories/>>

Jeshion, Robin (2013a), Expressivism and the Offensiveness of Slurs, *Philosophical Perspectives*, 27 (1), pp. 231–259.

Jeshion, Robin (2013b), Slurs and Stereotypes, *Analytic Philosophy*, 54 (3), pp. 314–329.

Jeshion, Robin (2014), *Slurs and the Expression of Contempt*, Handout for Names, Demonstratives and Expressives Conference, Gargnano.

Jeshion, Robin (2016a), *The Social Dimension of Slurs*, Handout for Normativity, Emotions and Social Life – SRSSP, Milan.

Jeshion, Robin (2016b), Slur Creation, Bigotry Formation: the Power of Expressivism, *Phenomenology and Mind*, 11, pp. 130–139.

Jeshion, Robin (forthcoming), Slurs, Dehumanization, and the Expression of Contempt, to appear in Sosa, David, *Bad Word*, Oxford, Oxford University Press.

Jorgensen, Julia, Miller, Jorge A. and Sperber, Dan (1984), Test of the mention theory of irony, *Journal of Experimental Psychology: General*, 113 (1), pp. 112–120.

Kaplan, David (1999), *The meaning of Ouch and Oops*, manuscript.

Karttunen, Lauri (1971), Some Observations on Factivity, *Papers in Linguistics*, 5, pp. 55–69.

Karttunen, Lauri (1973), Presuppositions of compound sentences, *Linguistic inquiry*, 4 (2), pp. 169–193.

Karttunen, Lauri (1974), Presupposition and linguistic context, *Theoretical linguistics*, 1 (1–3), pp. 81–194.

Karttunen, Lauri, and Peters, Stanley (1979), Conventional Implicatures in Montague Grammar, in Oh, Choon-Kyu and Dinneen, David A. (eds.), *Syntax and Semantics 11: Presupposition*, New York, Academic Press, pp. 1–56.

Keenan, Thomas R. and Quigley, Kathleen (1999), Do young children use echoic information in their comprehension of sarcastic speech? A test of echoic mention theory, *British Journal of Developmental Psychology*, 17 (1), pp. 83–96.

Kennedy, Christopher (2007), Vagueness and grammar: The semantics of relative and absolute gradable adjectives, *Linguistics & Philosophy*, 30, pp. 1–45.

Kennedy, Christopher and McNally, Louise (2005), Scale structure, degree modification and the semantics of gradable predicates, *Language*, 81, pp. 345–381.

Klein, Ewan (1980), A semantics for positive and comparative adjectives, *Linguistics and Philosophy*, 4, pp. 1–45.

Koev, Todor (2014), Two puzzles about appositives: Projection and perspective shift, in Etxeberria, Urtzi, Fălăuş, Anamaria, Irurtzun, Aritz, Leferman, Bryan, (eds.), *Proceedings of Sinn und Bedeutung 18*, pp. 217–234.

Kratzer, Angelika (1999), *Beyond Ouch and Oops: How Descriptive and Expressive Meaning Interact*, Handout for Cornell Conference on Theories of Context Dependency, Amherst.

Kyle, Brent G. (2013), How Are Thick Terms Evaluative?, *Philosophers' Imprint*, 13 (1), pp. 1–20.

Langton, Rae (2012), Beyond Belief: Pragmatics in Hate Speech and Pornography, in Maitra, Ishani and McGowan, Mary Kate (eds.), *What Speech Does*, Oxford, Oxford University Press, pp. 72–93.

Lasnik, Peter (2007), Expressives, perspective and presupposition, *Theoretical Linguistics*, 33 (2), pp. 223–230.

Leech, Geoffrey N. (1983), *Principles of Pragmatics*, London, Longman.

Lenehan, Rose (2015), *Beyond slurs: communicating evaluative perspectives*, master dissertation, Boston, Massachusetts Institute of Technology.

Levinson, Stephen C. (2000), *Presumptive Meaning: The theory of generalized conversational implicature*, Cambridge (MA), MIT Press.

Lewis, David (1989), Dispositional Theories of Value, *Proceedings of the Aristotelian Society*, Supplementary Volume 63, pp. 113–137.

López de Sa, Dan (2008), Presuppositions of commonality: An indexical relativist account of disagreement, in García-Carpintero, Manuel, and Kölbel, Max (2008), *Relative truth*, Oxford, Oxford University Press, pp. 297–310.

Macià, Josep (2002), Presuposición y significado expresivo, *Theoria: Revista de Teoria, Historia y Fundamentos de la Ciencia*, 3 (45), pp. 499–513.

Macià, Josep (2006), *Context, Presupposition and Expressive Meaning*, Handout for Milan Meeting 2006, Milan.

Macià, Josep (2011), *A defense of the presuppositional view of expressive meaning*, Handout for ECAP-7, Milan.

Mason, Michelle (2003), Contempt as a Moral Attitude, *Ethics*, 113, pp. 234–272.

McCready, Eric (2007), Context Shifting in Questions and Elsewhere, in Puig-Waldmüller, Estela (ed.), *Proceedings of Sinn und Bedeutung 11*, pp. 433–447.

McCready, Eric (2010), Varieties of conventional implicature, *Semantics and Pragmatics*, 3 (8), pp. 1–57.

McNally, Louise and Stojanovic, Isidora (2017), Aesthetic Adjectives, in Young, James (ed.), *The Semantics of Aesthetic Judgment*, Oxford, Oxford University Press.

McNeill, Brian W., Prieto, Loreto R., Niemann, Yolanda F., Pizarro, Marc, Vera, Elizabeth M. and Gómez, Sylvia P. (2001), Current directions in Chicana/o psychology, *The Counseling Psychologist*, 29 (1), pp. 5–17.

Meier, Cécile (2003), The Meaning of too, enough and so... that, *Natural Language Semantics*, 11 (1), pp. 69–107.

Mey Jacob, L. (1985), *Whose Language? A Study in Linguistic Pragmatics*, Amsterdam and Philadelphia, John Benjamins.

Mey Jacob, L. (1994), How to do good things with words: a social pragmatics for survival, *Pragmatics*, 4 (2), pp. 239–263.

Mey Jacob, L. (1993), *Pragmatics: An Introduction*, Oxford and Malden (MA), Blackwell Publishers.

Mey Jacob, L. (2012), Anticipatory pragmatics, *Journal of Pragmatics*, 44 (5), pp. 705–708.

Miščević, Nenad (2006), Moral concepts: From thickness to response-dependence, *Acta analytica*, 21 (1), pp. 3–32.

Miščević, Nenad (2011), Slurs & Thick Concepts - Is the New Expressivism Tenable?, *Croatian Journal of Philosophy*, 11 (32), pp. 159–182.

Miščević, Nenad and Perhat, Julija (2016), *A word which bears a sword*, Zagreb, KruZak.

Moore, Mathew S. and Levitan, Linda (2003), *For hearing people only*, Rochester (NY), MSM Productions.

Nunberg, Geoffrey (forthcoming), The Social Life of Slurs, in Fogal, Daniel, Harris, Daniel and Moss, Matt (eds.), *New Work on Speech Act*, Oxford, Oxford University Press.

O'Dea, Conor J., Miller, Stuart S., Andres, Emma B., Ray, Madelyn H., Till, Derrik F. and Saucier, Donald A. (2015), Out of bounds: factors affecting the perceived offensiveness of racial slurs, *Language Sciences*, 52, pp. 155–164.

Panzeri, Francesca and Carrus, Simone (2016), Slurs and Negation, *Phenomenology and Mind*, 11.

Plunkett, David and Sundell, Tim (2013), Disagreement and the semantics of normative and evaluative terms, *Philosophers' Imprint*, 13 (23), pp. 1–37.

Popa-Wyatt, Michaela (2016), Not all Slurs are Equal, *Phenomenology and Mind*, 11.

Potts, Christopher (2005), *The logic of conventional implicatures*, Oxford, Oxford University Press.

Potts, Christopher (2007a), Conventional implicatures, a distinguished class of meanings, in Gillian Ramchand and Charles Reiss (eds.), *The Oxford Handbook of Linguistic Interfaces Studies in Theoretical Linguistics*, Oxford, Oxford University Press, pp. 475–501.

Potts, Christopher (2007b), The centrality of expressive indexes. Reply to commentaries, *Theoretical Linguistics*, 33 (2), pp. 255–268.

Potts, Christopher (2007c), The expressive dimension, *Theoretical Linguistics*, 33 (2), pp. 165–197.

Potts, Christopher (2012), The pragmatics of conventional implicature and expressive content, in Maienborn, Claudia, von Stechow, Klaus and Portner, Paul (eds.), *Semantics: An International Handbook of Natural Language Meaning*, 3, Berlin, Mouton de Gruyter, pp. 2516–2536.

Potts, Christopher (2015), Presupposition and Implicature, in Lappin, Shalom and Fox, Chris (eds.), *The Handbook of Contemporary Semantic Theory*, second edition, Oxford, Wiley-Blackwell, pp. 168–202.

Predelli, Stefano (2010), From the Expressive to the Derogatory: On the Semantic Role for Non-Truth-Conditional Meaning, in Sawyer, Sara (ed.), *New Waves in Philosophy of Language*, Palgrave-MacMillan, pp. 164–185.

Predelli, Stefano (2013), *Meaning Without Truth*, Oxford, Oxford University Press.

Rabinowicz, Wlodek, and Rønnow-Rasmussen, Toni (2004), The Strike of the Demon: On Fitting Pro-attitudes and Value, *Ethics*, 114 (3), pp. 391–423.

Rappaport, Jesse (ms), *Communicating with Slurs*.

Rappaport, Leon (2005), *Punchlines: The Case for Racial, Ethnic, and Gender Humor*, Westport, Praeger Publishers.

Recanati, François (2006), D'un contexte à l'autre, in Behr, Irmtraud, Marillier, Jean-François and Dalmas, Martine (eds.), *Text und Sinn. Studien zur Textsyntax und Deixis im Deutschen und Französischen. Festschrift für Marcel Vuillaume*, Tübingen, Stauffenburg Verlag.

Recanati, François (forthcoming), Local pragmatics: reply to Mandy Simons, *Inquiry*.

Richard, Mark (2008), *When Truth Gives Out*, Oxford, Oxford University Press.

Richard, Mark (2011), Reply to Lynch, Mišćević, and Stojanovic, *Croatian Journal of Philosophy*, 32, pp. 197–208.

Roberts, Craig (1996), Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics, in Yoon, J. and Kathol, A. (eds.), *Ohio State University Working Papers in Linguistics*, Columbus, The Ohio State University Department of Linguistics, pp. 91–136.

Roberts, Debbie (2013), It's Evaluation, Only Thicker, in Kirchin, Simon (ed.), *Thick Concepts*, Oxford, Oxford University Press, pp. 78–96.

Saka, Paul (2007), *How to Think About Meaning*, Springer.

Sassoon, Galit W. (2013), A Typology of Multidimensional Adjectives, *Journal of Semantics*, 30, pp. 335–380.

Saucier, Donald A., Till, Derrik F., Miller, Stuart S., O'Dea, Connor J., and Andres, Emma B. (2015), Slurs against masculinity: masculine honor beliefs and men's reactions to slurs, *Language Sciences*, 52, pp. 108–120.

Sauerland, Uli (2007), Beyond unpluggability, *Theoretical Linguistics*, (33) 2, pp. 231–236.

Sbisà, Marina (1999), Ideology and the persuasive use of presupposition, in Verschueren, Jef, *Language and ideology. Selected papers from the 6th International Pragmatics Conference*, Antwerp: International Pragmatics Association, pp. 492–509.

Sbisà, Marina (2007), *Detto non detto: le forme della comunicazione implicita*, Roma-Bari, Laterza.

Scanlon, Thomas M. (1998), *What We Owe to Each Other*, Cambridge (MA), Harvard University Press.

Schlenker, Philippe (2003), A plea for monsters, *Linguistics & Philosophy*, 26, pp. 29–120.

Schlenker, Philippe (2007), Expressive presuppositions, *Theoretical Linguistics*, 33 (2), pp. 237–245.

Sidgwick, Henry (1884), *The Methods of Ethics*, third edition, London, Macmillan.

Simons, Mandy (2005), Presupposition and relevance, in Szabó Zoltán (ed.), *Semantics vs. Pragmatics*, Oxford, Oxford University Press, pp. 329–355.

Simons, Mandy, Tonhauser, Judith, Beaver, David and Roberts, Craig (2010), What projects and why, in Li, Nan and Lutz, David (eds.) (2010), *Semantics and linguistic theory*, Ithaca (NY), CLC Publications, pp. 309–327.

Sperber, Dan and Wilson, Deidre (1986), *Relevance: Communication and cognition*, Cambridge (MA), Harvard University Press.

Spotorno, Nicola and Bianchi, Claudia (2015), A plea for an experimental approach on slurs, *Language Sciences*, 52, pp. 241–250.

Spotorno, Nicola and Noveck, Ira A. (2014), When is irony effortful?, *Journal of Experimental Psychology: General*, 143 (4), pp. 1649–1665.

Sreenivasan, Gopal (2001), Understanding alien morals, *Philosophy and Phenomenological Research*, 62 (1), pp. 1–32.

Stanley, Jason (2015), *How propaganda works*, Princeton, Princeton University Press.

Stojanovic, Isidora (2016a), Evaluative Adjectives and Evaluative Uses of Ordinary Adjectives, in Bekki, Daisuke and McCready, Eric (eds.), *Proceedings of LENLS12: Language Engineering and Natural Language Semantics*.

Stojanovic, Isidora (2016b), Expressing aesthetic judgments in context, *Inquiry*, 59 (6), pp. 663–685.

Sudo, Yasutada (2012), *On the semantics of phi features on pronouns*, doctoral dissertation, Boston, Massachusetts Institute of Technology.

Tenchini, Maria Paola, and Frigerio, Aldo (2016), A multi-Act Perspective on Slurs, in Finkbeiner, Rita, Meibauer, Jörg and Wiese, Heike, *Pejoration*, Amsterdam, John Benjamins, pp. 167–186.

Terkourafi, Marina (2003), Generalised and Particularised Implicatures of Politeness, in Kuhnlein, Peter, Rieser, Hannes and Zeevat, Henk (eds.), *Perspectives on Dialogue in the New Millennium*, Amsterdam, John Benjamins, pp. 151–166.

Terkourafi, Marina (2005a), Beyond the Micro-level in Politeness Research, *Journal of Politeness Research: Language, Behavior, Culture*, 1 (2), pp. 237–262.

Terkourafi, Marina (2005b), Pragmatic Correlates of Frequency Use: the Case for a Notion of ‘Minimal Context’, in Marmaridou, Sophia, Nikiforidou, Kiki and Antonopoulou, Eleni (eds.), *Reviewing Linguistic Thought: Converging Trends for the 21st Century*, Berlin, Mouton de Gruyter, pp. 209–233.

Terkourafi, Marina (2005c), An Argument for a Frame-based Approach to Politeness: Evidence from the Use of the Imperative in Cypriot Greek, in Lakoff, Robin and Ide, Sachiko (eds.), *Broadening the Horizon of Linguistic Politeness*, Amsterdam, John Benjamins, pp. 99–116.

Thomson, Judith J. (1992), On Some Ways in Which A Thing Can be Good, *Social Philosophy and Policy*, 9 (2), pp. 96–117.

Thomson, Judith J. (2008), *Normativity*, Chicago, Open Court.

Tirrell, Lynne (1999), Derogatory Terms: Racism, Sexism, and the Inferential Role Theory of Meaning, in Hendricks, Christina and Oliver, Kelly (eds), *Language and Liberation: Feminism, Philosophy, and Language*, Albany (NY): SUNY Press.

Thommen, Tristan (2014), *Response-Dependence and the Semantics of Slurs*, Master Dissertation, Paris, École Normale Supérieure.

Tonhauser, Judith (2012), Diagnosing (not-) at-issue content, in *Proceedings of Semantics of Under-represented Languages of the Americas (SULA) 6*, Amherst, GLSA, pp. 239–254.

Tonhauser, Judith, Beaver, David, Roberts, Craig and Simons, Mandy (2013), Toward a taxonomy of projective content, *Language*, 89 (1), pp. 66–109.

Väyrynen, Pekka (2006), Resisting the Buck-Passing Account of Value, in Shafer-Landau, Russ (ed.), *Oxford Studies in Metaethics*, Oxford, Oxford University Press, pp. 295–324.

Väyrynen, Pekka (2009), Objectionable Thick Concepts in Denials, *Philosophical Perspectives*, 23, pp. 439–469.

Väyrynen, Pekka (2011), Thick Concepts and Variability, *Philosophers' Imprint*, 11, pp. 1–17.

Väyrynen, Pekka (2013), *The lewd, the rude and the nasty: A study of thick concepts in ethics*, Oxford, Oxford University Press.

Väyrynen, Pekka (2014), Essential Contestability and Evaluation, *Australasian Journal of Philosophy*, 92 (3), pp. 471–488.

Väyrynen, Pekka (2016a), *Evaluatives and Pejoratives*, Handout for Linguistics Seminars-Scuola Normale Superiore, Pisa.

Väyrynen, Pekka (2016b), Thick Ethical Concepts, in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, URL=<http://plato.stanford.edu/archives/win2016/entries/thick-ethical-concepts/>.

Wang, Linton, Reese, Brian and McCready, Eric (2005), The projection problem of nominal appositives, *Snippets*, 10, pp. 13–14.

Whiting, Daniel (2013), It's Not What You Said, It's the Way You Said It: Slurs and Conventional Implicatures, *Analytic Philosophy*, 54 (3), pp. 364–377.

Williams, Bernard (1985), *Ethics and the Limits of Philosophy*, Cambridge (MA), Harvard University Press.

Williamson, Timothy (2009), Reference, Inference, and the Semantics of Pejoratives, in Almog, Joseph and Leonardi, Paolo (eds.), *The Philosophy of David Kaplan*, New York, Oxford University Press, pp. 137–158.

Wilson, Deidre (2006a), The pragmatics of verbal irony: echo or pretence?, *Lingua*, 116, pp. 1722–1743.

Wilson, Deidre (2006b), Attributive uses of concepts, *Issues in Pragmatics*, (PLINM301), ms.

Wilson, Deidre (2007), Where does irony come from?, *Pragmatic Theory* (PLIN2002) 2007-08, ms.

Wilson, Deidre and Sperber, Dan (2012), *Meaning and Relevance*, Cambridge, Cambridge University Press.

Yablo, Stephen (2006), Non-catastrophic presupposition failure, in Thomson, Judith and Byrne, Alex, *Content and modality: Themes from the philosophy of Robert Stalnaker*, Oxford, Clarendon Press, pp. 164–190.

Résumé

Dans cette thèse, je développe une analyse présuppositionnelle des termes d'offense (ou péjoratives), d'une part, et des termes étiques épais, d'autre part. Je soutiens que les termes d'offense et les termes étiques épais – même si ils ont été étudiés par différentes disciplines – appartiennent à la même classe de termes que j'appelle 'évaluatifs hybrides'. Mon travail vise à combler l'écart entre la recherche sur les termes étiques épais d'une part et la recherche sur les termes d'offense d'autre part. Je montre que le mécanisme sous-jacent les termes d'offense et les termes étiques épais est le même : les différences qui peuvent être observées sont superficielles et elles dépendent des particularités du contenu descriptif des évaluatifs hybrides.

Dans la première partie de ma thèse, j'introduis un critère pour déterminer ce qu'on peut définir 'évaluatif hybride' et je présente mon analyse des évaluatifs hybrides en termes de présuppositions. La deuxième partie est un examen critique des théories qui ont été développées pour expliquer le comportement linguistique respectivement des termes d'offense et des termes étiques épais. Je présente les théories vériconditionnelles, les théories déflationnistes, l'approche en termes des implicatures conventionnelles et je mets en évidence les aspects à l'égard desquels l'analyse présuppositionnelle est préférable. Dans la troisième partie, je propose une analyse uniforme de deux utilisations particulières des termes d'offense et des termes étiques épais qui ont été traitées comme indépendantes jusqu'à présent : il s'agit de l'Appropriation des termes d'offense et de la Variabilité des termes étiques épais.

Mots Clés

Termes étiques épais, termes d'offense, présuppositions, Appropriation, Variabilité, termes d'évaluation, sémantique, pragmatique.

Abstract

In this thesis I develop a uniform account of slurs and thick terms in terms of presuppositions. I argue that slurs and thick terms – even though they were studied by different disciplines – belong to the same class of 'hybrid evaluatives'. My work aims to fill the gap between the research on thick terms on the one hand and the research on slurs on the other, by showing that the mechanism underlying slurs and thick terms is one and the same and that the phenomenal differences that one can observe depend on the peculiarities of their descriptive content. In Part I, I define what counts as 'hybrid evaluative' and present my presuppositional account. Part II is a critical review of some alternative theories that have been put forward to account for slurs and thick terms respectively: I focus on truth-conditional theories, on deflationary theories and on the conventional implicature approach and show in which respects the presuppositional account is superior. In Part III, I propose a uniform account of two particular uses of slurs and thick terms that have been usually treated as independent so far, namely Appropriation of Slurs and Variability of thick terms. I put forward an analysis in terms of 'echo', that distinguishes from 'merely attributive' uses of language and 'echoic' ones. Such a uniform account is largely independent from the specific analysis of slurs and thick terms that I favour, but it supports once more the idea that slurs and thick terms should be analysed along the same lines.

Keywords

Thick terms, pejoratives, presupposition, Appropriation, Variability, evaluative terms, semantics, pragmatics.