The social meaning in swearing variation

Matthew Hunt

Submitted in partial fulfilment of the requirements
of the degree of Doctor of Philosophy

January 2022

School of Languages, Linguistics and Film
Queen Mary University of London

Statement of originality

I, Matthew Christopher Hunt, confirm that the research included within this thesis is my own work or that where it has been carried out in collaboration with, or supported by others, that this is duly acknowledged below and my contribution indicated. Previously published material is also acknowledged below.

I attest that I have exercised reasonable care to ensure that the work is original, and does not to the best of my knowledge break any UK law, infringe any third party's copyright or other Intellectual Property Right, or contain any confidential material.

I accept that the College has the right to use plagiarism detection software to check the electronic version of the thesis.

I confirm that this thesis has not been previously submitted for the award of a degree by this or any other university.

The copyright of this thesis rests with the author and no quotation from it or information derived from it may be published without the prior written consent of the author.

Signature:

Date: 13/01/2022

Abstract

This thesis is about linguistic variation in swearing and its consequences for how speakers are socially evaluated. Abundant research has established that, beyond its perception as rude or impolite, swearing is hugely socially meaningful in a variety of ways (Stapleton, 2010; Beers Fägersten, 2012). Swearing has been shown to index solidarity (Daly et al., 2004), intimacy (Stapleton, 2003), differing forms of masculinity (De Klerk, 1997) and femininity (S. E. Hughes, 1992), honesty (Feldman et al., 2017), believability (Rassin & Heijden, 2005) and lack of intelligence (DeFrank & Kahlbaugh, 2019), among other traits. The activation of these social meanings also depends on language-external factors such as speaker gender (Howell & Giuliano, 2011), ethnicity (Jacobi, 2014) and social status (T. Jay & Janschewitz, 2008). What has not been established is whether this also depends on language-internal factors such as pronunciation, word formation or sentence structure.

This thesis investigates the effect of variation from three different domains of language - phonetics, morphology and semantics/pragmatics - on social evaluation of a speaker. To do so, the thesis takes an experimental approach using the variationist sociolinguistic framework. For variation in each domain, two experiments were used to test for different levels of awareness, following Squires's (2016) approach for grammatical variation (see also Schmidt, 1990). One experiment tested whether people *perceived* the variation, while a second tested whether people *noticed* the variation in the process of social evaluation; the concepts of *perceiving* and *noticing* roughly map to the Labovian concepts of the sociolinguistic *indicator* and *marker* respectively (Labov, 1972).

At the level of phonetics, variation in the realisation of variable (ING) in swearwords (e.g., *fucking* vs *fuckin*) was first tested using a variant categorization task, revealing that listeners have an implicit bias towards the velar [ɪŋ] variant when hearing swearwords, compared to neutral words and non-words. An auditory matched-guise task then revealed that this same bias affects how listeners extract social information from (ING) tokens attached to swearwords in relation to social meanings typically associated with the variable (Schleef et al., 2017). This result suggests that, rather than pronunciation affecting how swearwords are socially evaluated, swearwords can affect how other phonetic sources of social meaning are evaluated.

At the level of morphology, the linguistic constraints on swearing infixation (e.g., *fan-fucking-tastic*) were first tested using an acceptability judgment task, revealing that native speakers of English have an implicit knowledge of the prosodic constraints on swearing infixation (McCarthy, 1982). A visual matched-guise task then tested whether infixation affects the social perception of *fucking*. The results suggest that, compared to non-infixed swearing constructions, infixed swearing constructions are perceived as ruder, funnier and more sarcastic; this was modulated by the well-formedness of the infixed construction. The results provide an example of morphology, specifically expressive morphology (Zwicky & Pullum, 1987), that can act as a source of social meaning.

At the level of semantics and pragmatics, variation in object animacy (e.g., *the fucking paint* vs *the fucking kids*) and adjective gradability (e.g., *fucking long* vs *fucking wooden*) in swearing modification was first tested using a self-paced reading task. These variables are described using insights from formal semantics (Potts, 2005, 2007) and pragmatics (Grice, 1975). Results of the self-paced reading task then revealed a significant slowdown for inanimate, compared to animate, objects (but not for non-gradable, compared to gradable, adjectives). A matched-guise task using just the animacy variable then revealed that people using swearwords to modify animate objects (e.g., *kids*, *chef*, *daughter*) are perceived as ruder than people using swearwords to modify inanimate objects (e.g., *paint*, *heat*, *carpet*). The findings further previous work on socio-semantics (Beltrama, 2020), as well as suggesting that reactions to swearwords are modulated by the hearer's perception of the target of the heightened emotion.

The findings from each chapter are discussed in relation to Eckert's (2008) conceptualisation of the *indexical field*. This thesis pushes swearing research forward by shedding light on the as-yet-unexamined area of intra-word variation in swearing as a significant factor in how it is perceived. With regard to sociolinguistics more generally, the thesis provides a methodological advance by applying Squires's approach for grammatical variation to new domains of language. The consequences of these findings for sociolinguistic cognition are also discussed.

Dedication

I'd like to dedicate this thesis to a guy who swore at me once when I was working as a doorman in Southampton. After being unceremoniously ejected from our establishment, this chap shouted at me 'you're a fucking prick you cunt!'. This seemingly contradictory but perfectly coherent statement is what inspired me to write my undergraduate thesis on swearing, which was eventually followed by my Master's thesis on swearing, which has now been followed by my PhD thesis on swearing. I hope you got home safe, and thank you for your contribution.

Acknowledgements

Whenever I read people's theses, for some reason I always have a quick read of their acknowledgements section. Obviously the information it contains isn't actually relevant to whatever I'm trying to find, but its nice to remind myself that, behind the 300 odd pages that I'm reading, there's a real person responsible for putting them together. A person with family, friends, co-workers, supervisors and a whole range of other people that supported them along the way. So for the small handful of people that might read this thesis, here's a quick rundown of the people to whom I owe a great deal of thanks.

My supervisory team: Colleen Cotter, Linnaea Stockall and Hazel Pearson. It would be remiss of me to not to include a rugby analogy at some point, so here goes. Thank you to Colleen, my DOR, who has overseen my project since it's inception in early 2017 as part of my funding proposal; you've been there every step of the way, critiquing, questioning and, most importantly, supporting me in every aspect of my PhD journey. Thank you to Linnaea, my forward's coach, who has been the driving force behind my experiments, pushing me to always be rigorous in their design and to leave no room for doubt in the claims I draw from them. Thank you to Hazel, my back's coach, who has been my rock when it comes to the semantics elements of my thesis in particular; I've lost count of the number of times you've had to tolerate my attempts at the discipline and help turn those attempts into something meaningful.

The staff at Queen Mary. Thank you to all the members of staff that have provided me feedback, attended my practice talks or taught me on modules over the last 4 years, including Adam Chong, Kathleen McCarthy, Devyani Sharma, Erez Levon, Jenny Cheshire, Jonathan Kasstan, Hagit Borer and Luisa Martí. In particular, I'd like to single out Adam for all the support you've given me, on both an academic and personal level. Before I arrived at Queen Mary, I received a strong grounding in linguistics from the staff at the University of Southampton. In particular, thank you to Patrick Stevenson for everything you taught me and for giving me sound and honest advice on pursuing postgraduate study.

My collaborators on other projects. The only way I have kept sane over the last 4 years is by spending some time away from my thesis on other projects, so I'd to thank the people that made those projects possible. Thank you to Eric Acton for working with me on our 'the wife' project. I'm very proud to have my first academic publication with

your name alongside mine and hopefully it won't be the last. Thank you to the illustrious Sue Denim, with whom I worked on a project about linguistic profiling against Uyghurs in the job market; thank you for taking a risk with me and also just for being an immense person. I've also benefited from interactions with numerous other linguists along the way who have given me sound advice on my project, including Kristy Beers Fägersten, Karyn Stapleton and Andrea Beltrama.

The other PhD students at Queen Mary. The LingLab has often been a place of comfort for me during my PhD and that is largely due its inhabitants. Thank you to Louis, Sarah, Elvis, Stamantina, Christian, Nate, Maame, Ynda, Zoe, Daniella, Elisa, Scott, Tom Rausch, Pietro, Chiara, Margaret, Lujain, Tom Meadows, Pedro, Adrian, Songyan, Liam, Celia, Jessica, Daniella and Zoe. In particular, I'd like to thank my "twin" Louis for being there with me through the whole journey and for always being a willing sounding board to my questions.

To my various cohabitants since the start of my PhD. Thank you to Rose and Sophie for putting up with me for two years in Elephant and Castle. Thank you to Pietro for putting up with me (on and off) for 1 and a half years in Canning Town, particularly the first 6 months of lockdown 2020. And thank you to Rosie, the person with whom I hope to be cohabiting for many years to come; you have supported me more than you can possibly know.

Finally, thank you to my parents. To my mother for inspiring me to continue to pursue my academic interests and to my father for helping me to stay grounded in the real world.

Before I conclude these acknowledgements, and in a slight departure from what these sections are normally about, there is one last thing I want to acknowledge: this was hard. I spent the first year and a half of my PhD in relative normality, before the rest of it was interrupted by the COVID-19 pandemic. Like everyone else in the world, my life moved online. This presented a lot of challenges. While my data collection likely would have occurred online regardless, losing the chance to have daily, in-person interactions with students and staff at Queen Mary made the process very difficult. I'm noting this here for no other reason than to remind myself, whenever I come back to this thesis, that I got through it and, if I can get through this, then whatever is to come shouldn't seem so hard.

Contents

List of Tables **12**

List of Figures **14**

1 Introduction **16**
1.1 What do we know about swearing? . . . 16
1.2 A division of meaning(s) - semantics . . . 20
1.3 It's all in the word - sociolinguistics . . . 22
1.4 Combining disciplines . . . 25
1.5 Summary . . . 26
1.6 Thesis overview . . . 26

2 Studying swearing **31**
2.1 Introduction . . . 31
2.2 The features and parameters of swearing . . . 33
2.3 Swearing in linguistics . . . 35
2.3.1 Phonetics . . . 35
2.3.2 Morpho-syntax . . . 36
2.3.3 Semantics . . . 38
2.3.3.1 A homogenous set of meanings . . . 39
2.3.3.2 A heterogeneous set of uses . . . 46
2.3.4 Pragmatics . . . 48
2.3.5 Psycholinguistics . . . 51
2.4 Swearing in sociolinguistics . . . 56
2.4.1 Production data . . . 57
2.4.1.1 Corpus linguistics . . . 57
2.4.1.2 Swearing in interaction . . . 59
2.4.2 Perception data . . . 62
2.4.2.1 Language-external factors . . . 65
2.4.3 Social meaning . . . 66

2.4.3.1 The indexical field . . . 67
2.4.3.2 Swearing and the indexical field . . . 69
2.5 The knowledge gap . . . 71

3 New methods for studying swearing 74
3.1 Introduction . . . 74
3.2 Perceiving vs Noticing in Sociolinguistics . . . 75
3.2.1 Perceiving vs Noticing: Swearing variation . . . 78
3.3 Measuring what people perceive . . . 80
3.3.1 The variant categorisation task . . . 81
3.3.2 The acceptability judgment task . . . 83
3.3.3 The self-paced reading task . . . 85
3.4 Measuring what people notice . . . 88
3.4.1 The matched guise task . . . 88
3.5 Participant Recruitment . . . 94
3.6 Statistical analysis . . . 95
3.7 Summary . . . 95

4 Swearing and the socio-phonetic variable 96
4.1 Introduction . . . 96
4.2 Socio-phonetic variation . . . 97
4.3 The variable: (ING) . . . 98
4.4 Experiment I: Perceiving phonetic variation in swearing . . . 101
4.4.1 Methods . . . 101
4.4.1.1 Experimental Hypothesis . . . 101
4.4.1.2 Stimuli . . . 105
4.4.1.3 Procedure . . . 111
4.4.1.4 Participants . . . 113
4.4.2 Results . . . 113
4.4.2.1 Reaction time data . . . 116
4.4.3 Interim discussion . . . 116
4.5 Experiment II: Noticing socio-phonetic variation in swearing . . . 119
4.5.1 Methods . . . 120
4.5.1.1 Experimental Hypothesis . . . 120
4.5.1.2 Stimuli . . . 122
4.5.1.3 Procedure . . . 124
4.5.1.4 Participants . . . 126
4.5.2 Results . . . 126
4.5.3 Interim discussion . . . 129

4.6 Chapter discussion and conclusions . . . 131

5 Swearing and the socio-morphological variable 134
5.1 Introduction . . . 134
5.2 Socio-morphogical variation . . . 135
5.3 The variable: infixation . . . 138
5.4 Experiment I: Perceiving morphological variation in swearing . . . 141
5.4.1 Methods . . . 142
5.4.1.1 Experimental Hypothesis . . . 142
5.4.1.2 Stimuli . . . 143
5.4.1.3 Procedure . . . 143
5.4.1.4 Participants . . . 144
5.4.2 Results . . . 144
5.4.3 Interim discussion . . . 146
5.5 Experiment II: Noticing morphological variation in swearing . . . 147
5.5.1 Methods . . . 150
5.5.1.1 Experimental Hypotheses . . . 150
5.5.1.2 Stimuli . . . 151
5.5.1.3 Selecting the scales . . . 152
5.5.1.4 Procedure . . . 153
5.5.1.5 Participants . . . 155
5.5.2 Results . . . 155
5.5.3 Interim discussion . . . 156
5.6 Chapter discussion and conclusions . . . 160

6 Swearing and the socio-semantic variable 162
6.1 Introduction . . . 162
6.2 Socio-semantic variation . . . 163
6.3 The variables . . . 165
6.3.1 Variable I: Animacy . . . 167
6.3.1.1 What is animacy? . . . 167
6.3.1.2 The culprit hypothesis . . . 168
6.3.2 Variable II: Gradability . . . 171
6.3.2.1 What is gradability? . . . 172
6.3.2.2 The scale hypothesis . . . 174
6.3.2.3 Scales in intensification . . . 176
6.4 Experiment I: Perceiving semantic variation swearing . . . 180
6.4.1 Methods . . . 181
6.4.1.1 Experimental hypotheses . . . 181

6.4.1.2 Stimuli . . . 185
6.4.1.3 Procedure . . . 189
6.4.1.4 Participants . . . 190
6.4.2 Results . . . 191
6.4.2.1 Animacy results . . . 192
6.4.2.2 Gradability results . . . 196
6.4.3 Interim discussion . . . 199
6.5 Experiment II: Noticing socio-semantic variation in swearing . . . 201
6.5.1 Methods . . . 202
6.5.1.1 Experimental hypothesis . . . 202
6.5.1.2 Stimuli . . . 203
6.5.1.3 Selecting the scales . . . 204
6.5.1.4 Procedure . . . 206
6.5.1.5 Participants . . . 207
6.5.2 Results . . . 207
6.5.3 Interim discussion . . . 209
6.6 Chapter discussion and conclusions . . . 212

7 Conclusion: Swearing and the sociolinguistic variable 215
7.1 Introduction . . . 215
7.2 The indexical field: a recap . . . 216
7.3 The social meaning in swearing: contextualizing previous findings . . . 220
7.4 The social meaning in swearing: contextualizing linguistic variation . . . 224
7.4.1 The indexical field of swearing infixation . . . 225
7.4.2 The indexical field of animate and inanimate swearing modification 227
7.4.3 The indexical field of swearing and variable (ING) . . . 229
7.5 Consequences for swearing research . . . 231
7.6 Limitations . . . 234
7.7 Future directions . . . 235
7.8 The final word . . . 237

Appendix 238

A 239
A.1 Infixing Experiment I: Trisyllabic adjectives . . . 240
A.2 Infixing Experiment II - Sentences and adjectives . . . 243
A.3 Infixing Experiment II - Example stimuli list . . . 244

B **246**

B.1 Animacy/Gradability Experiment I - Animacy: Test Nouns 247

B.2 Animacy/Gradability Experiment I - Animacy stimuli 248

B.3 Animacy/Gradability Experiment I - Gradability: Test adjectives 249

B.4 Animacy/Gradability Experiment I - Gradability stimuli 251

B.5 Animacy/Gradability Experiment II - Animacy Stimuli: Animate sentences 252

B.6 Animacy/Gradability Experiment II - Animacy Stimuli: Inanimate sentences 253

References **254**

List of Tables

4.1 Experiment I: List of items . . . 107
4.2 Experiment I: Norming model summary - Response ∼ Continuum Step + Item Type + (1|Participant) + (1|word) . . . 110
4.3 Experiment I: Example trial sheet . . . 111
4.4 Experiment I: Model summary - Response ∼ Continuum Step + Item Type + Pre-step + Participant Gender + (1|Participant) + (1|word) . . . 114
4.5 Experiment I: Model summary - RT ∼ Item Type + Continuum Step + (1|Participant) + (1|word) . . . 116
4.6 Experiment II: Test recording key . . . 120
4.7 Experiment II: Model summaries for *articulate*, *rich*, *working class* and *educated* - Response ∼ Condition + 'Swearing = LAP' + 'Social Class' + (1|Participant) + (1|Speaker) . . . 127
4.8 Experiment II: Model summaries for *posh*, *hard-working* and *casual* - Response ∼ Condition + (1|Participant) + (1|Speaker) . . . 128

5.1 Rule for expletive infixation from Siegel (1974) and Aronoff (1976). . . . 138
5.2 Rule for expletive infixation from McCarthy (1982). . . . 140
5.3 Experiment I: Model summary - Response ∼ Stress * Form + Valency + Orthographic Neighbourhood Density + First Vowel Match + (1|Participant) + (1|Word) . . . 146
5.4 Experiment II: Experimental design . . . 150
5.5 Experiment II: Model summary - Response ∼ Infixed + Formedness + Participant Age + Adjective Log Freq + Participant Income + Infix Freq + (1|Participant) + (1|Adjective) . . . 156

6.1 Experiment I: Animacy stimuli - 3x2 research design . . . 186
6.2 Experiment I: Gradability stimuli - 2x2 research design . . . 186
6.3 Experiment I: Nominal alpha values calculated in *GroupSeq* using Power Family: $alpha * t^{\phi}$ . . . 191
6.4 Experiment I: Animacy stimuli - sentence labels . . . 192

6.5 Experiment I: Model summary: Modifier, Object, +1 - Reaction Time $\sim$ Modifier + Noun Type + Participant Age + Swearing = appropriate + (1|Participant) + (1|Word) 193
6.6 Experiment I: Model summary: +2, +3, +4 - Reaction Time $\sim$ Modifier * Noun Type + Participant Age + Swearing = appropriate + (1|Participant) + (1|Word) 193
6.7 Experiment I: Model summary: +1, +2, +3, +4 - Summed Reaction Time $\sim$ Modifier + Noun Type + (1|Participant) + (1|Object) 194
6.8 Experiment I: Model summary: +1, +4 - Reaction Time $\sim$ Modifier + Animacy Category + Participant Age + Swearing = appropriate + (1|Participant) + (1|Object) 195
6.9 Experiment I: Gradability stimuli - sentence labels 197
6.10 Experiment I: Model summary: Modifier, Adjective and +1 - Reaction Time $\sim$ Modifier + Adjective Type + (1|Participant) + (1|Word) 197
6.11 Experiment I: Model summary: +2, +3, +4 - Reaction Time $\sim$ Modifier + Adjective Type + (1|Participant) + (1|Word) 197
6.12 Experiment I: Model summary: +1, +2, +3, +4 - Summed Reaction Time $\sim$ Modifier + Adjective Type + (1|Participant) + (1|Adjective) 198
6.13 Experiment II: Experimental design 202
6.14 Experiment II: Model summary for *rude* and *angry*: Response $\sim$ Condition * Swear (1|Participant) + (1|$Noun$) 208
6.15 Experiment II: Model summary for *annoyed* and *aggressive*: Response Condition + Swear (1|Participant) + (1|$Noun$) 209

List of Figures

1.1 Plot of valency and arousal using data from Janschewitz (2008) 24

4.1 TANDEM-STRAIGHT: Distance Matrix for two sounds 109
4.2 Experiment I: Norming experiment line-graph: Response ∼ Continuum Step 110
4.3 Experiment I: Experimental procedure 113
4.4 Experiment I: Mean probability of '-ing' selection ∼ Continuum Step * Item Type 115
4.5 Experiment I: Mean probability of '-ing' selection ∼ Continuum Step * Item Type - colour blind-friendly version 115
4.6 Experiment II: Mean responses for the scales *articulate*, *rich*, *working class* and *educated* by condition. *All-ing* = 100% velar, *All-in* = 100% alveolar, *Swear-ing* = swearwords are velar, *Swear-in* = neutral words are velar. 128
4.7 Experiment II: Mean responses for the scales *casual*, *posh* and *hard-working* by condition. *All-ing* = 100% velar, *All-in* = 100% alveolar, *Swear-ing* = swearwords are velar, *Swear-in* = neutral words are velar . . 129

5.1 Metrical structure tree for *fantastic* 139
5.2 Metrical structure tree for *fabulous* 139
5.3 Well-formed metrical structure for fan-fucking-tastic 140
5.4 Well-formed metrical structure for fabu-fucking-lous 140
5.5 Ill-formed metrical structure for fanta-fucking-stic 141
5.6 Ill-formed metrical structure for fab-fucking-ulous 141
5.7 Experiment I: Boxplot of mean responses by Stress and Form 145
5.8 Experiment II: Words associated with 'That's fucking fantastic' and 'That's fan-fucking-tastic' 153
5.9 Experiment II: Boxplot - Scale response Infixation * Form 155
5.10 Experiment II: Principal Component Analysis: plotted using *autoplot* in the *ggfortify* package 158

6.1 Experiment I: Animacy results for *fucking* 194
6.2 Experiment I: Animacy results for *damn* 194
6.3 Experiment I: Animacy results for *bloody* 194
6.4 Experiment I: Gradability results 198
6.5 Experiment II: Scale selection adjective frequencies 205
6.6 Experiment II: Mean response for *angry* and *rude* by Condition and Swear 208
6.7 Experiment II: Mean response for *aggressive* and *annoyed* by Condition and Swear 209

7.1 Indexical field of (ING. Black = meanings for the velar variant, gray = meanings for the apical variant). From Eckert (2008). 216
7.2 Indexical field of /t/ release. Boxes = social types, black = permanent qualities, gray = stances. From Eckert (2008). 217
7.3 Indexical field for *favelado*. Lower case = permanent qualities; upper case = social types; dotted line separates positive and negative qualities and social types. From Beaton (2015). 218
7.4 Indexical field of 'swearing': Boxes = social types, black = permanent qualities, grey = stances. Based on the findings of Christie (2013). 219
7.5 Indexical field of 'swearing' in the workplace: Boxes = social types, black = permanent qualities, grey = stances/interpersonal functions. Based on the findings of Daly (2004), Baruch (2007) and McLeod (2011). 222
7.6 Indexical field of 'swearing' and gender: Boxes = social types, black = permanent qualities, grey = stances/interpersonal functions. Based on various findings on swearing and gender. 223
7.7 Indexical field of variable (ING) in Manchester: Boxes = social types, black = permanent qualities, grey = stances/interpersonal functions. From Schleef and Flynn (2015). 224
7.8 Indexical field of *fucking infixation*. 225
7.9 Indexical fields of animate and inanimate swearing modification for *fucking* and *damn*. 227
7.10 The intersecting indexical fields of alveolar [ɪn], swearing and velar [ɪŋ] . 230
7.11 The *predicted* indexical fields of alveolar [ɪn] and swearing 231

Chapter 1

Introduction

1.1 What do we know about swearing?

Swearing is a type of language with properties that interest linguists from a range of disciplines, including, but not limited to, sociolinguistics, semantics and psycholinguistics. As a research area in its own right, improving our understanding of 'swearing' ultimately increases our understanding of human behaviour and language. A swearword is typically categorized as such because it has been deemed inappropriate, or in some cases strictly prohibited, for civil conversation due to some association with potential for harm to occur. In fact, the prohibition of certain types of language has been ever-present throughout history. The themes from which swearwords are drawn may change (Montagu, 1967; G. Hughes, 1998), but the existence of swearwords persists, insofar as 'swearwords' denote some subset of vocabulary items that are commonly considered 'taboo' and/or 'offensive' by a particular speech community (a more precise definition will follow in Chapter 2).

Swearwords also present an interesting paradox for linguists. Despite its apparent potential for causing harm and offense, swearing is ubiquitous across the majority of languages and cultures (T. Jay, 2009b). More paradoxically still, of the subset of vocabulary items frequently circumscribed as 'swearwords' in English, the items often considered the most *offensive* are also the ones that are uttered the most frequently (Beers Fägersten, 2012). Unless we assume that speakers are constantly walking around deliberately offending one another with their words, swearing must be more complicated than simply taboo or offensive.

The majority of instances of swearing are socially motivated, rather than being spontaneous and uncontrolled (Beers Fägersten, 2012). We would therefore expect a variety of social meanings to be available to speakers through swearing. The term 'social meaning' will be fully explicated and operationalized in Chapter 2, but for now, it will suffice to say that social meaning pertains to information that is encoded in language

that expresses some aspect of the speaker's identity, which can include their membership to a particular macro-social category (e.g. gender, ethnicity, social class) or their orientation towards locally relevant norms of communication (Eckert, 2000, 2008; Zhang, 2008; Mendoza-Denton, 2008, among others). Sociolinguistics is resplendent with examples of non-standard language varieties that allow speakers to do their own context-specific "meaning-making" (Eckert, 2008, p. 465) and we would expect the same to be true for swearing, particularly because swearing arguably epitomises the notion of non-standardness.

In the literature, this is shown to be the case. Swearing is associated with a multiplicity of meanings outside of offensiveness. Swearing can be used to mark solidarity and positive politeness (Daly et al., 2004), toughness (S. E. Hughes, 1992), intimacy or trust (Stapleton, 2003) and humour (Stapleton, 2010). It can be used to evoke a range of emotions, both positive and negative (Wang, Chen, Thirunarayan, & Sheth, 2014; Lutzky & Kehoe, 2016) and to index differing forms of both masculinity (De Klerk, 1997) and femininity (S. E. Hughes, 1992; Stapleton, 2003). It is often used by speakers to converge towards localised norms of communicating (S. E. Hughes, 1992; Drummond, 2020). Swearwords show interesting spoken usage trends across macro-social categories such as age, gender and social class (McEnery, 2004; McEnery & Xiao, 2004; Love, 2017, 2021); similar trends have also been found in in online contexts (Thelwall, 2008; Wang et al., 2014; Gauthier & Guille, 2017).

Swearing is also shown to play a role in social evaluation. In some studies, swearing has been shown to negatively affect how a speaker is perceived in terms of overall impression, trustworthiness and intelligence (DeFrank & Kahlbaugh, 2019; Johnson & Lewis, 2010), as well as professional capability (Paradise, Cohl, & Zweig, 1980) and effectiveness (Howell & Giuliano, 2011). In other studies, swearing has shown a positive effect on how a speaker is socially evaluated. For example, swearing can improve the general impression of a politician, although it can also decrease the persuasiveness of their message (Cavazza & Guidetti, 2014). Elsewhere, swearing has been shown to increase a speaker's perceived believability (Rassin & Heijden, 2005) and workplace effectiveness (Johnson, 2012), as well as improving a hearer's attitude towards a topic of conversation (Scherer & Sagarin, 2006).

How swearing is perceived also varies as a function of contextual factors. Females typically rate swearwords are more offensive than males (Beers Fägersten, 2012), with swearing more broadly associated with male speech (Coates, 1986; Martin, 1997). When evaluating pairs of speakers, listeners typically react more favorably to swearing when those speakers are of the same gender (DeFrank & Kahlbaugh, 2019), while swearing by male sports coaches is deemed less appropriate when aimed at a female team than at a male one (Howell & Giuliano, 2011). The ethnicity of the swearer also contributes to

how they are perceived, with black swearers evaluated less favourably than white swearers when uttering the same swearwords (Jacobi, 2014). A listener's native language influences their perception of the emotional force of swearing (Dewaele, 2004), as well as their specific dialect of a language (Dewaele, 2015). Finally, the formality of the setting and the status of the listener relative to the speaker also influences how swearing is perceived (T. Jay & Janschewitz, 2008).

The studies that have focused on perceptions of swearing are lacking in scope, however. In each, *swearing* is treated as a homogeneous set of lexical items. The majority of the papers on perceptions of swearing (see Section 2.4.2 of Chapter 2) state that *swearing* makes people seem more or less X, where X is a social trait; the authors talk only of the presence or absence of swearing as a variable to be manipulated. *Swearing*, in these studies, is always well-defined as a subset of language, but is typically a vague descriptor when one actually examines what the researchers are measuring. Typically, they measure changes in perceptions of a social trait as a function of the use of one or a couple of swearwords. The conclusions they then draw about *swearing* are only based on the results of a few items. For example, they may be based on the use of one or two swearwords in the context of one conversation, with the presence or absence of swearing included as a between-subjects variable.

Even in studies that compare different swearwords to one another, the researchers assume homogeneity in the way each swearword will be produced. Historically, the majority of studies used word-rating tasks, in which participants were presented with a list of words and asked to rate, using a scale, how offensive each word would be to them (see Beers Fägersten, 2012 for a review). Those words are typically treated as unvarying entities, however. The same is true in more recent social perception work (e.g., Jacobi, 2014). What these studies do not consider is how the social meanings attributed to swearwords may vary as a function of the way in which those words are used. This is despite the fact that, as well as representing a complex and varying social phenomenon, there also exists a large amount of linguistic variation in swearing.

In many ways, swearwords in English behave just like other words in the language. Consider, for example, the uses of *fucking* and *damn* in (1-a). In the semantics literature, swearwords like this are referred to as *expressive adjectives* (EAs). Syntactically, they appear to behave like other attributive adjectives, or descriptive adjectives (DAs), for example in (1-b). In languages that mark gender agreement, EAs also behave like DAs, as in (2) (Gutzmann, 2019).

(1) a. The damn/fucking dog ate too quickly
    b. The brown/aggressive dog ate too quickly

(2) a. Gestern hat die jung**e** Hund die ganze Nacht gebellt

b. Gestern hat die verdammt**e** Hund die ganze Nacht gebellt

In other ways however, EAs differ significantly from DAs. EAs are restricted to the positive form and cannot be used in comparative form (3-a). Furthermore, they cannot be degree modified (3-b), appear in predicative position (3-c) or be the target of adverbial modification (Gutzmann, 2019).

(3) a. *The damner dog
b. *The very damn dog
c. *The dog is damn
d. *The probably damn dog

EAs also differ semantically from DAs. Unlike DAs, EAs license non-local readings whereby they can target a constituent larger than the nominal in which they occur (Potts, 2005). For example, the uses of *fucking* and *damn* in (4-a) can target *dog*, *my dog* or *I lost my dog in the park*; in the latter case, the speaker expresses their heightened emotion about the whole state of affairs described by the sentence. The same is not true for DAs. For example, there is no way in which *brown* or *crazy* in (4-b) target anything other than *dog*.

(4) a. I lost my fucking/damn dog in the park
b. I lost my brown/crazy dog in the park

Swearwords have previously received attention in the domain of semantics (Kaplan, 1999; Potts, 2005, 2007), with many semanticists suggesting that swearwords are a form of expressive or not-at-issue content. That is, they contribute a meaning that is separate from the sentence's truth conditions. Further accounts of this type have suggested that some expressives can convey both at-issue and not-at-issue meanings (McCready, 2010; Gutzmann, 2011, 2015). The distinction could have consequences for how swearwords are perceived. For example, literal uses of swearwords (e.g. *he took a shit* or *they fucked each other*) are typically considered more offensive than metaphorical uses (e.g. *the film was shit* or *they really fucked up*) (Beers Fägersten, 2012). Variation in the type of meaning contributed by a swearword can therefore affect how the person using it is perceived.

Relatedly, phonetic variation in swearing can contribute to its pragmatic meaning. For example, subtle changes in the /ʌ/ vowel in *fuck* can mark the difference between different intended meanings, such as intensification, confusion, dissatisfaction or suspicion (Gold & McIntyre, 2016). In English, swearwords in general are also more likely than neutral words to contain harsher, plosive consonants (Yardy, 2010). Furthermore, just like any other set of words, swearwords will be pronounced by an array of different speakers in ways which are bound to change how they are perceived; work in sociophonetics has

shown this to be true for neutral words (see e.g., Campbell-Kibler, 2005), so it follows that the same could be true for swearwords.

Swearwords can also vary in their morphology, such as in infixed constructions e.g., *fan-fucking-tastic*. While infixation is common in other languages for inflection or derivational purposes, expletive infixation is the only common use of infixation in English (although see Elfner & Kimper, 2008 on the use of *diddly*-infixation in The Simpsons). Thought to be constrained by prosody and universal well-formedness conditions (McCarthy, 1982), swearing infixation is a rare case of expressive morphology in English (Zwicky & Pullum, 1987). Determiner phrases containing the word *fuck* can also be inserted into set phrases including wh-questions (e.g. *what the fuck is that?*). Expletive insertion can be used as a conversational device for sequential organising and sanctioning (Hoey et al., 2020). Another example of expressive morphology in swearing is compounding, as in the English constructions *fuckton* or *fuckload* or German constructions such as *Arschgesicht* (arse+face) (Meibauer, 2013).

In sum, there is a large amount of linguistic variation in swearing. This variation occurs at multiple levels of structure, including phonetics, morpho-phonology and syntax and semantics. Given that perceptions of swearing are shown to vary as a function of language-external factors, including speaker- and context-specific factors, we might also expect variation as a function of language-internal factors. Are all pronunciations of a swearword perceived the same? All word formations? All sentence structures? These are questions I will consider in this thesis. If the social meanings attributed to swearing depend on who and in what context the swearing occurs, it would be logical to expect those meanings to also depend on how they occur.

1.2 A division of meaning(s) - semantics

There are two main reasons to expect linguistic variation in swearing to influence social evaluation. The first relies on the notion of content being either at-issue or not-at-issue (Roberts, 1996; Simons, Tonhauser, Beaver, & Roberts, 2010). While this concept will be further developed in Chapter 6, it is relevant to the thesis as a whole with respect to how listeners distribute their attention to incoming information. At-issue (or descriptive) meaning, on the one hand, comprises all and any meaning contained within the truth conditions of a sentence. For a sentence such as (5-a), we can know what the world has to be like for the statement to be true; this is illustrated using the truth conditions in (5-b)[1]

(5) a. The dog is hungry

[1]This is obviously an overly simplified semantics for (5-a). For the purposes of distinguishing between at-issue and not-at-issue content however, it is sufficient.

b. "the dog is hungry" is true in a world w, iff $w \in [\![\text{the dog is hungry}]\!]^t$

The situation is more complicated for the sentence in (6-a). While the meaning of the rest of the sentence remains the same, it is harder to pin down the exact meaning contribution of *fucking*. Potts (2007) suggests that the meaning contributed by *fucking* is, among other things, immediate, perspective dependent and descriptively ineffable. That is, *fucking* is "inflict[ed]" on the common ground and reflects the perspective of the speaker in a way that cannot be simply described. The meaning of *fucking* is also argued to behave independent of the rest of the sentence. For example, while the at-issue content can be negated, you cannot felicitously negate the expressive meaning conveyed by *fucking*, as in (6-b). That is, the negation doesn't reverse the expression of heightened emotion on behalf of the speaker.

(6) a. The dog is fucking hungry
b. The dog isn't fucking hungry (# And I feel neutral about the situation)

To deal with these two types of meaning, many semanticists typically think of meanings being either at-issue or not-at-issue. At-issue content is the main point of the utterance, while not-at-issue content is secondary. Following Gutzmann and Turgay (2019, p. 1), I take secondary content to be content that "is not the main point of the utterance, but instead provides side and/or background information, which is less prominent and less active than the utterance's main content". Other examples of not-at-issue content include slurs, appositives and dog-whistles.

Another aspect of language which is, arguably, a form of not-at-issue content is *social meaning*. Many parallels can be drawn between social meaning and other forms of not-at-issue meaning. Consider the example in (7).

(7) John is playin' football

The sentence in (7) has the at-issue meaning that 'John is playing football'. In addition however, it has a source of social meaning in its phonetic form, namely the use of alveolar [ɪn] in the word *playin'*. A variety of studies in sociolinguistics have shown this pronunciation to be socially meaningful across multiple varieties of English in both production and perception (Labov, 1972; Houston, 1985; Campbell-Kibler, 2005; Schleef et al., 2017, among others). Much like expressives (including swearwords), the social meaning of alveolar [ɪn] operates as not-at-issue content. It is independent of the truth conditions of (7) (although see Acton, 2020). It immediately inflicts something onto the common ground of the conversation; you can't negate the social meaning, for example. It conveys something about the speaker's perspective e.g. their background, social class, chosen speech style etc. Finally it is descriptively ineffable, insofar as it is very difficult

to provide a propositional paraphrase for the meaning of alveolar [ɪn] (see Blakemore, 2011).

Both expressive meanings (e.g. those expressed by swearwords) and social meanings are not-at-issue in the sense that they are secondary to the main content of what is being expressed. The main pieces of information for a listener to take away from (6-a) and (7) are contained in their at-issue contents. Secondary to that, they may also take away some information about the speaker based on their use of *fucking* or alveolar [ɪn]. For this reason, if two forms of not-at-issue meaning are present, for example in a form like *fuckin*, we might expect one to influence perception of the other.

Work in sociolinguistic perception has already demonstrated that, when multiple sources of extra-linguistic meaning are available, the result is not just the sum of its parts. That is, if two socially meaningful variables are present in the speech signal, each with their own distinct set of potential social meanings, the speaker isn't associated with the sum total of those meanings. Rather, different variables can combine to produce unique sociolinguistic profiles (Campbell-Kibler, 2011). In some cases, this has been argued to reflect the different levels of social salience that variables display (Levon, 2014). Furthermore, some variables may be more or less prominent depending on the overall speech style in which they are used (Campbell-Kibler, 2007; Pharao, Maegaard, Møller, & Kristiansen, 2014).

In cases where listener attention is already focused on the at-issue content of an utterance, there are limited attentional resources available to process not-at-issue content (see Campbell-Kibler, 2016 on cognitive economy in sociolinguistics). We know, on the basis of the studies referenced above, that swearwords influence social evaluation. We also know, on the basis of work in sociolinguistics (see next section), that sociolinguistic variation influences social evaluation. As yet to be explored is whether the two interact to produce different social evaluations than they would in isolation.

1.3 It's all in the word - sociolinguistics

The second reason to expect language-internal variation in swearwords to influence social evaluation relates to similar findings for other word-specific factors. A consistent finding in work on sociolinguistic perception is that people's social evaluations of a speaker can vary as a function of subtle changes in the speech signal; this idea has been central to the variationist tradition, where speech communities can be grouped together not just via shared behavioural norms, but also shared evaluative norms (Labov, 1972; Campbell-Kibler, 2010b). A range of experimental studies on different language varieties suggest that listeners extract social information from the speech signal at all levels of structure, including phonetics (Campbell-Kibler, 2005; Levon, 2014; Schleef et al., 2017), phonol-

ogy (Tyler, 2015; Levon & Ye, 2019), syntax (Bender, 2005; Levon & Buchstaller, 2015) and semantics/pragmatics (Beltrama & Staum Casasanto, 2017; Hunt & Acton, 2022).

There is also a strong body of evidence suggesting that the extraction of social information from language depends on other social cues. If a listener is primed to expect to hear a speaker with particular social characteristics, this expectation can influence how they interpret aspects of the speech signal. This has been shown in speech communities with emerging sound changes, including vowel mergers (e.g. *pin-pen*). In experiments, listeners are tasked with classifying potentially ambiguous sounds as either the old sound or the new sound. When provided information that the speaker belongs to a group which are stereotypically associated with the sound change - but with no actual change in the speech signal - listeners are more likely to classify those sounds as the new sound. This has been shown for both age- (Hay, Warren, & Drager, 2006; Koops, Gentry, & Pantos, 2008) and region-based (Niedzielski, 1999; D'Onofrio, 2015, 2018) stereotypes. Sociolinguistic perception has also been shown be sensitive to information pertaining to gender (Strand, 1999) and ethnicity (Casasanto, 2008; D'Onofrio, 2019).

As well as varying as a result of social factors, the perception of sociolinguistic variables also depends on language-internal factors. Production of sociolinguistic variables can vary due to lexical frequency (Hay, Jannedy, & Mendoza-Denton, 1999) and grammatical category (Kendall, 2013). The same two factors have also been shown to condition listener expectation in speech perception (Vaughn & Kendall, 2018); speakers have stored knowledge about the likelihood of a linguistic variant occurring as a function of word-internal properties. This knowledge can also show up in social evaluation. For example, in a study of copula absence in African American Vernacular English (e.g. *she teaching me piano*), Bender (2005) shows that listener's social evaluations of the variable depend on grammatical factors. While copula absence was associated with a particular set of social meanings overall (when compared to copula presence), among certain speakers, evaluations also depended on how marked copula absence was in a particular grammatical environment (e.g. pre-verb vs pre-noun). A similar result was found for social evaluations of the intensifier *totally*, depending on the gradability properties of the following adjective (Beltrama & Staum Casasanto, 2017), and for variable (ING), depending on the part-of-speech properties of the particular word (Vaughn, 2021). Finally, word frequencies vary across different social groups (by e.g. region, age, gender etc). Listeners show an awareness of this in word recognition tasks, in which words uttered by a member of a congruent social group are recognised faster (Walker & Hay, 2011; Kim, 2016). In sum, listeners have stored knowledge of how particular words (and groups of words) relate to both linguistic and social representations.

Given their unique properties, we might expect listeners to store similar knowledge for swearwords. As well as being characterised by *offensiveness*, swearwords are rated

high in *arousal* and *tabooness* and low in *valency* (i.e. how 'happy' the word makes someone feel) (Janschewitz, 2008). Using the norming data from Janschewitz (2008) for the valency, arousal and tabooness of 460 words, I have plotted the 50 words rated as most taboo by their participants according to how they were rated for arousal and valency in blue (see Figure 1.1), with the other 410 words plotted in red. As the plot shows, there is fairly clear separation between the most taboo and the remaining words with respect to valency and arousal, with only a few overlapping data points. While neutral words are fairly spread out, taboo words are mostly localised to the low valency and high arousal space (with a few exceptions).

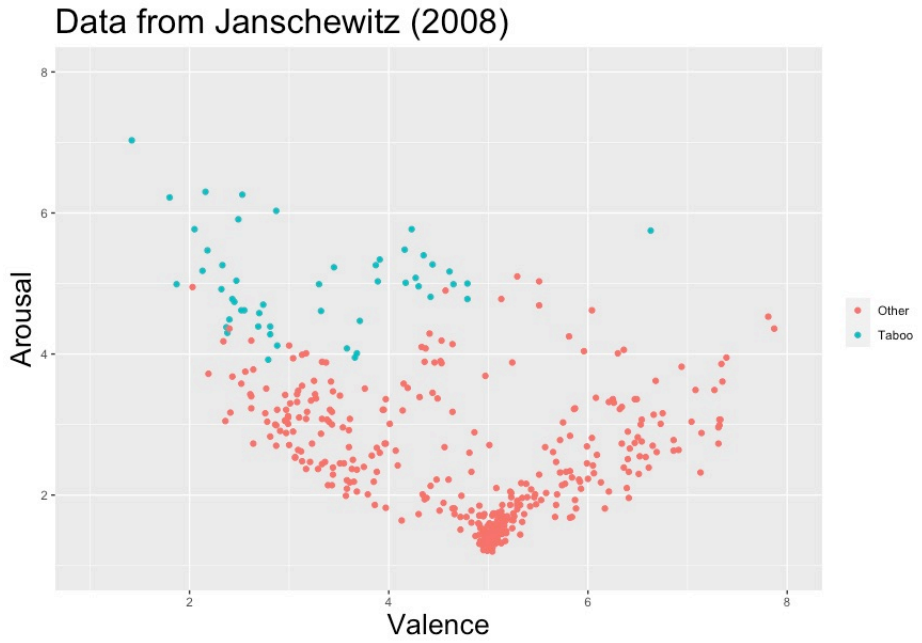

Figure 1.1: Plot of valency and arousal using data from Janschewitz (2008)

The role of valency, arousal and tabooness has been studied extensively within psycholinguistics; this work has typically focused on how participant behaviour in simple linguistic or non-linguistic tasks varies as a function of these factors, with particular focus on swearwords. For example, swearwords have been shown to improve lexical decision reaction times and free recall. Improved performance in these two tasks is driven by different factors however; lexical decision times improved as a result of lexical access and tabooness, while free recall was improved by tabooness, valency and arousal (Madan, Shafer, Chan, & Singhal, 2017). In other tasks, the presence of swearwords can inhibit performance. For example, taboo language causes slower reaction times in modified Stroop (MacKay et al., 2004; Guillet & Arndt, 2009; Eilola & Havelka, 2011), attentional blink (Mathewson, Arnell, & Mansfield, 2008) and picture-word inference tasks (Dhooge & Hartsuiker, 2011). The processing of swearwords can also be affected by social factors such as speaker gender (Tuft, McLennan, & Krestar, 2018).

In all of these tasks, participants were required to process a stimulus and provide a response; their behaviour changed as a result of factors specific to swearwords, namely tabooness, valency and arousal. The same process may occur in social evaluation tasks. In such tasks, listeners hear or read a stimulus. They then provide a response to that stimulus. Given that listeners show an awareness of internal linguistic factors when evaluating language for social characteristics (Bender, 2005; Beltrama & Staum Casasanto, 2017), as well as having awareness of word-specific factors that relate to social meanings (Walker & Hay, 2011; Kim, 2016), it follows that they might do the same for swearword-specific factors like *tabooness*, *valency* and *arousal*.

1.4 Combining disciplines

The points covered in Sections 1.2 and 1.3, while distinct, refer to the same basic phenomenon: some meanings are secondary to the main content of an utterance. In the semantics literature, these secondary meanings (or not-at-issue meanings) are often discussed in relation to attitude, for example in swearwords or slurs. In the sociolinguistics literature, secondary meanings are discussed in relation to a broad range of concepts including social categories, stereotypes, personae and ideological orientations, among others. While the two disciplines may study aspects of secondary meaning using different tools, they can be used in combination to further our understanding of *meaning*, broadly construed.

Interest in combining the tools of these two disciplines, as well as those from formal pragmatics, has picked up new momentum in recent years (Burnett, 2017b; Eckert, 2019; Beltrama, 2020; Acton, 2020). Eckert (2019) suggests that, for a broader theory of social practice, meanings can be characterised along a continuum of simultaneously decreasing reference and increasing performance, with purely semantic meanings at the highly referential end and purely social meanings at the highly performative end. While the work of sociolinguists has often involved excluding truth-conditional meanings, focusing instead on the social meanings indexed by abstract linguistic units such as sound segments, Eckert argues for an integrated theory of meaning that balances the referential and performative properties of linguistic forms.

Similarly, Beltrama (2020) suggests that sociolinguists should pay more attention to the role of semantics in socially meaningful variation, for example with respect to the relation between semantic and pragmatic properties and the social salience of a form. Burnett (2017b) provides a framework for modelling sociolinguistic variation using formal pragmatic principles, suggesting that such an approach may help link together the more qualitative and quantitative approaches to sociolinguistics. Finally, Acton (2020) argues that the traditions of pragmatics and sociolinguistics are "mutually enriching" in

potentially unexpected ways. It is possible for linguistic forms that are commonly considered to be purely performative, such as sociophonetic variables, to influence a sentence's truth-conditions. It is also possible for forms commonly considered purely semantic, such as determiners, to have social value (Acton & Potts, 2014; Acton, 2019).

1.5 Summary

The salient points to take from this introduction are as follows. Swearwords are associated with a variety of social meanings in perception. This perception is influenced by language-external factors. No previous work has tested the influence of language-internal factors however. We would expect language internal factors to influence the social perception of swearing for two main reasons. Firstly, because swearwords and social meanings fall into the same not-at-issue meaning category, thus they are competing for a listener's attention. Secondly, because other word-specific factors have been shown to influence the perception of social meanings. Combining tools from sociolinguistics, pragmatics and semantics can improve our understanding of *meaning*, broadly construed.

1.6 Thesis overview

The remainder of this thesis continues as follows. Chapter 2 is a comprehensive review of research on swearing across a wide range of sub-disciplines including phonetics, morpo-syntax, semantics, pragmatics and psycholinguistics, as well as an in-depth review of research from sociolinguistics, including work on both production and perception. Chapter 2 will also review the definitions of swearing used in the literature and will provide my own criteria for swearing to be used in this thesis. Chapter 2 will also introduce the *indexical field* (Eckert, 2008), with a view to operationalising it to analyse both previous findings on the social meaning of swearing and my own findings on linguistic variation in swearing and the effect this has on social evaluations of the speaker.

Chapter 3 introduces the methods to be used in this thesis. To investigate variation from three different areas of language - namely phonetics, morphology and semantics/pragmatics - I employed an approach suggested by Squires (2016), inspired by Schmidt (1990), for investigating linguistic knowledge at different levels of awareness. This approach separates out what people *perceive* from what they *notice* and from what they *understand*. Focussing on these first two processes, I employ two experiments in each of my three content chapters (Chapters 4, 5 and 6). The first assesses whether people *perceive* linguistic variation in swearing, that is, whether they have some internal representation of the variation below the level of consciousness. To do this, I employ a variant

categorization task, an acceptability judgment task and a self-paced reading task respectively; each of these techniques is reviewed in Chapter 3. The second assesses whether people *notice* linguistic variation in swearing, that is, whether their implicit knowledge of the variation affects the conscious process of social evaluation. To do this, I employ either an auditory or visual matched guise task in each chapter; this technique is reviewed in Chapter 3.

Chapter 4 focuses at the level of sound and how swearing may interact with a phonetic variable to influence social evaluation, namely variable (ING). Variable (ING) relates to variation in the realisation of *-ing*, for which both the velar [ɪŋ] and alveolar [ɪn]variants are available options. Variable (ING) has been studied widely in sociolinguistics (Labov, 1966; Trudgill, 1974; Tagliamonte, 2004; Hazen, 2008, among others). The alveolar [ɪn] variant is commonly associated with certain characteristics in perception, including a working-class identity, an informal speech style and lower intelligence (Campbell-Kibler, 2005; Labov, 2006; Schleef et al., 2017). Swearing has been shown to share these associations (S. E. Hughes, 1992; T. Jay & Janschewitz, 2008; Stapleton, 2010; DeFrank & Kahlbaugh, 2019).

The first experiment in Chapter 4 measures speakers' implicit association of swearing with alveolar [ɪn]. This was done using a variant categorization task in which participants heard a mixture of phonetically-matched swearwords, neutral words and non-words, all ending with *-ing*. This suffix was artificially manipulated in MATLAB to create a 7-step nasal continuum from maximally velar to maximally alveolar, with acoustically ambiguous tokens in the middle. Participants categorized each item as ending in either '*-ing*' or '*-in*'. The prediction was that, due to their shared associations with particular speakers and speech styles, swearwords would prime listeners to select '*-in*' when the actual (ING) token was ambiguous. The results instead suggest that listeners were more likely to select '*-ing*' for swearwords compared to neutral words. I present two possible explanations for this effect. The first is that swearing primed listeners to hear the velar [ɪŋ] variant due to its association with verbal emphasis (Stapleton, 2010) and, therefore, careful speech. The second possible explanation was that, due to the increased attentional resources taken up by the swearword stems, participants' attention was drawn away from the (ING) token, meaning they defaulted to the most likely option overall i.e. the '*-ing*' form.

The second experiment in Chapter 4 measures whether this effect holds in full sentences and whether this, in turn, influences social evaluation. This was done using a n auditory matched-guise task, in which participants listened to multiple speakers telling a story and then rated those speakers on a selection of Likert scales. The story contained 10 words ending with *-ing*, including 5 swearwords and 5 neutral words. Four speakers were recorded reading the story with both velar and alveolar realisations of those words. The recordings were then artificially manipulated to either include exclusively velar to-

kens (*All-ing*), exclusively alveolar tokens (*All-in*) or 50% of each, with either all the swearwords (*Swear-ing*) or all the neutral words (*Swear-in*) as velar realisations. The prediction was that the alveolar [ɪn] tokens on swearwords would be heard velar [ɪŋ] tokens, but those on neutral words would not. Responses to the *Swear-in* guises were therefore predicted to approximate those for the *All-ing* guise; the same was not predicted to occur for the *Swear-ing* guise, despite both having equal proportions of each variant. The results suggested this to be the case on a subset of scales typically associated with variable (ING) in the UK (Schleef et al., 2017), namely *articulate*, *rich*, *working-class* and *educated*. This suggests that swearwords can interact with a phonetic variable to influence social evaluation in way that sets them apart from neutral words.

Chapter 5 focuses at the level of words and how different word formations of a swearing construction influence social evaluation. The variable of interest in this chapter is infixation. Expletive infixation has been studied extensively by phonologists (Siegel, 1974; Aronoff, 1976; McCawley, 1978). The most prominent account is McCarthy (1982), who claims that the rules of expletive infixation in English fall out from the universal conditions of prosodic well-formedness. While expletive infixation immediately prior to a metrical foot (e.g. *fan-fucking-tastic*) is well-formed, an infix which is interior to a foot (e.g. *fanta-fucking-stic*) is ill-formed. A full account of McCarthy's generalisation is provided in Chapter 5.

The first experiment of Chapter 5 tests whether people have an implicit knowledge of the prosodic constraints on swearing infixation, as well as testing a range of other linguistic factors that might be expected to influence the acceptability of an infixed swearing construction. This was done using an acceptability judgement task. Participants were presented with 201 different trisyllabic adjectives, including both stress-intial and stress-medial words, infixed with the swearword *fucking*. Participants were presented with either the well- or ill-formed version of each adjective, providing a judgment of the *acceptability* of that construction in English (either Yes or No). Adjectives were also coded for a number of other factors that could influence their acceptability. The results suggest that people do have implicit knowledge of the prosodic constraints on swearing infixation. In addition, the results suggest that the acceptability of a swearing infix is also influenced by stress position, valency, orthographic neighbourhood density and partial vowel harmony.

The second experiment of Chapter 5 tests the effect that sentences containing swearing infixation have on social evaluation compared to those containing identical non-infixed constructions (e.g. *fan-fucking-tastic* vs *fucking fantastic*). In addition, the experiment tests whether the phonological constraints tested in the first experiment also influence social evaluation, that is, whether it matters if a person's infixation is well-formed or not. This was done using a visual matched-guise task, in which participants read two-line dialogues containing well- or ill-formed infixed or non-infixed swearing constructions

and rated the 'speaker' on Likert scales. The results suggest that swearing infixation is perceived as funnier, ruder and more sarcastic compared to equivalent non-infixed expressions; these effects were partially influenced by well-formedness, suggesting that stored knowledge of the linguistic constraints of swearing infixation is tracked in social evaluation.

Chapter 6 focuses at the level of sentences and how swearing intensification can variably influence social evaluation, depending on whether that intensification is perceived to be local or non-local. A swearword such as *fucking* can appear prejacent to both adjectives, nouns and verbs. The syntactic behaviour of swearing intensifiers does not always match up with their semantics however. The literature on expressives (Potts, 2005, 2007; McCready, 2010; Gutzmann, 2011) suggests that, unlike descriptive adjectives, expressive adjectives are not limited to operating over their syntactic sister; rather, they can also be interpreted non-locally, expressing heightened emotion towards the state of affairs described by the sentence. I present two factors that may influence whether the local or non-local reading is the most likely reading, namely *animacy* and *gradability*. I suggest that the dominance of one reading over another stems from pragmatic reasoning on behalf of the listener; I provide detailed accounts of both factors in Chapter 6.

The first experiment of Chapter 6 tests whether speakers have some stored knowledge of the difference between local and non-local swearing intensification by measuring differences in processing. This was done using a self-paced reading task. Participants read sentences word-by-word through a moving window, controlling the transition from one word to the next using their keyboard. The task contained two sets of test sentences. In the first set, sentences contained a noun, modified by either *fucking*, *damn* or *bloody*. The noun was either animate (*dog*, *chef*) or inanimate (*shoe*, *stove*). In the second set, sentences contained an adjective, modified either by *fucking* or *probably* (a non-swear baseline). The adjective was either gradable (*tall*, *casual*) or non-gradable (*wooden*, *weekly*). The prediction was that sentences in which the non-local reading was dominant (i.e. non-gradable adjectives and inanimate objects) would be read faster than those in which the local reading was dominant. The results suggest that this was the case for animacy, where a slowdown was detected in the reading of words following inanimate objects, but not for gradability, where no significant effects on reading speed were found.

The second experiment of Chapter 6 tests whether the processing effect found for animacy in the first experiment also shows up in social evaluation, that is, whether it matters if a person is swearing to express their attitude towards a state of affairs or towards an individual. This was done using a visual matched-guise task in which participants read a series of two-sentence statements containing stimuli from the set of animacy sentences use in the first experiment. The results suggest that the use of either *fuck* or *damn* is perceived to be ruder if it precedes an animate object rather than an inanimate object; this

in turn suggests that swearwords are ruder if they are interpreted locally rather than non-locally. Similarly, for *damn* only, swearing about animate objects was perceived to make the speaker seem angrier than if they swore about an inanimate object. I take these results to show that people modulate their reactions to swearing depending on who or what the target of the swearing is perceived to be.

In the final chapter of the thesis, Chapter 7, I discuss the findings from Chapters 4-6. I revisit the concept of the *indexical field* introduced in Chapter 2, using this to bring together previous findings in swearing research to paint a more detailed picture of the social meanings of swearing; in this way I link together the micro-level stances and interpersonal functions, the social traits and characteristics and the more concrete social types and personae associated with swearing. Building on this, I then present original indexical fields for my own findings. I discuss the consequences of my findings for swearing research before discussing both the limitations of my work and the potential for future research in this area.

Chapter 2

Studying swearing

2.1 Introduction

Findings from research on swearing are frequently reported on by major news outlets. Examples in 2021 include the news in The Observer that swearing means you are more intelligent and helps you cope with pain (LaMotte, 2021), and that people are swearing more frequently, according to the BBC (Coughlan, 2021), but are increasingly choosing different words to do so, according to The Guardian (Davies, 2021). The findings reported in the news articles detail prominent work by K. Jay and Jay (2015), Stephens, Atkins, and Kingston (2009), Love (2017) and the British Film Classification Board. The prominence of findings from swearing research in the news media suggests that the public are interested in swearing. The work of Stephens et al. (2009) even led to Richard Stephens winning an Ig Nobel prize, an alternative to the Nobel praise for work that makes people "laugh first and think later". It would be easy to think, therefore, that while swearing clearly makes for attention-grabbing headlines, the motivations for doing academic research on swearing are not so clear. Linguists themselves may have similar doubts about swearing research, particularly in relation to social meaning, given how disparate the findings are across different disciplines.

As I will show in this chapter, however, research on swearing has much to teach us about human behaviour. Firstly, research on swearing contributes to larger questions about language. As a ubiquitous phenomenon cross-linguistically (T. Jay, 2009b), swearing could be considered a linguistic universal, insofar as every language has words that are prohibited due to societal taboo. If linguistics is concerned with the study of what is universal in natural language therefore, then swearwords must be included. Research in a number of domains of linguistics has already shown swearwords to be highly grammatically variable (Bergen, 2016). In particular, a large body of work in formal semantics (Potts, 2007; McCready, 2010; Gutzmann, 2015) has suggested that swearwords belong to a distinct subset of language known as *expressives*, a group that also includes slurs and

insults. Models of swearing in this tradition highlight a number of key components of swearing that distinguish them from neutral words, such as their perspective dependence and projective properties, which in-turn explain much about the role of swearing, and the conveying of emotion more generally, in communication.

Secondly, research on swearing in the domain of psycholinguistics contributes to answering questions about language processing and the allocation of attentional resources. Just as swearwords appear to be grammatically distinct from neutral words, they also differ from neutral words in a number of other properties relevant to how humans process incoming linguistic stimuli. In norming studies, English swearwords are rated as highly arousing and low in valency (Warriner, Kuperman, & Brysbaert, 2013), as well as highly taboo (Janschewitz, 2008). Based on numerous experiments testing the effects of these properties on both linguistic and non-linguistics tasks, swearwords have been argued to take up more attentional resources in processing (MacKay et al., 2004; Guillet & Arndt, 2009; Eilola & Havelka, 2011; Madan et al., 2017). Studying swearwords in the context of cognition therefore reveals much about how the brain prioritises different types of linguistic stimuli.

Finally, research on swearing as a social phenomenon contributes to answering questions about how people draw links between linguistic forms, on the one hand, and the people that use them, on the other. The majority of swearing is *social swearing*, rather than uncontrolled and emotional swearing (Beers Fägersten, 2012). In cases of social swearing therefore, a speaker has the choice between swearing or not swearing, presumably opting for the former because it fulfils the particular conversational or social goal they wish to achieve. Swearing behaviour also patterns as a function of social variables such as age, gender and social class (Love, 2017). When hearing swearing, listeners often form different impressions of the speaker than had they not sworn (Johnson & Lewis, 2010; Cavazza & Guidetti, 2014; DeFrank & Kahlbaugh, 2019, *inter alia*). Studying swearing in the context of social evaluation therefore reveals much about people's attitudes towards different types of language use and the factors that modulate these attitudes.

In summary, studying swearing is a worthwhile academic pursuit because of what it can tell us about how people use language, how people process language and how people associate language with the people that use it. Swearwords behave like a distinct subset of language in each of these processes. If we are to have comprehensive accounts of these processes therefore, those accounts must also be true for swearwords. As I will demonstrate in the final section of this chapter, however, previous work has stopped short of uniting our understandings of these three processes in relation to swearwords. Before highlighting this knowledge gap however, I will first review the previous findings in these areas that motivate studying swearing.

2.2 The features and parameters of swearing

As Limbrick (1991, p. 79) suggests, "swearing resists concrete definition; exactly what constitutes a swear word is generally determined by social codes". A number of linguists have attempted to provide definitions for swearing despite this difficulty, however. Expanding on criteria provided by Andersson and Trudgill (2007), Ljung (2010, p. 4) suggests the following four criteria for swearing:

1. Swearing is the use of utterances containing taboo words.
2. The taboo words are used with non-literal meaning.
3. Many utterances that constitute swearing are subject to severe lexical, phrasal and syntactic constraints, which suggest that most swearing qualifies as formulaic language.
4. Swearing is emotive language: its main function is to reflect, or seem to reflect, the speaker's feelings and attitudes.

The first criterion is universal in swearing research. Swearwords refer to something that is societally taboo, the mention of which is typically avoided due to the nature of the topic. By swearing, a speaker is breaking a taboo. The most common taboo topics for swearwords are scatology (e.g., *shit*, *piss*), sex including genitalia (e.g., *fuck*, *cunt*, *wank*, *prick*) and religion (e.g., *damn*, *bloody*) (Stapleton, 2010). Although there are additional categories of taboo used for swearing, such as diseases in Dutch, linguists generally agree that a word needs to break a taboo to qualify as a swearword.

Ljung (2010) argues for the second criterion by claiming that swearwords behave as a natural class in non-literal uses, despite being semantically unrelated. Consider the word *fuck*. In its literal use (e.g., *I fucked this guy*), it has numerous synonyms including *screw*, *bonk*, *shag* and *frig*. In a non-literal use such as the exclamation *fuck!*, the same words are not suitable replacements. The same is true in the non-literal use *fuck you!*, with the exception of *screw*. Replacement words are available from the pool of other swearwords, however. *Damn*, for example, could be used in both non-literal constructions (i.e., *damn!* and *damn you!*), despite being semantically unrelated to the taboo topic of sex. Ljung makes similar claims about *prick*, which can be replaced by *cock*, *dick*, *dork*, *pecker* or *pisser* in literal uses (e.g., *he hurt his prick*), but not in non-literal uses (e.g., *the dumb prick didn't listen*). Ljung acknowledges that literal uses of e.g., *fuck* are certainly taboo; they clearly still pertain to the taboo topic if sex. But, according to his criteria, only non-literal uses count as swearwords.

Several other linguists, including McEnery (2004), Love (2017) and Drummond (2020), include literal uses of swearing in their works, however. Love (2017) makes the case that

swearwords which do not have polysemous non-taboo uses (e.g., *fuck*, *cunt*, *shit*, contrasted with e.g., *bloody*) are shown to be more psychologically arousing (Janschewitz, 2008) and memorable (T. Jay, Caldwell-Harris, & King, 2008) than non-taboo words when encountered out of semantic context, that is, when there is no evidence to say that the use is either literal or non-literal. Drummond (2020) similarly argues for including literal uses of swearwords as part of a "common sense approach", whereby using the literal version of a swearword at a family meal would still likely be counted as swearing.

I take a perspective that aligns with that of McEnery (2004), Love (2017) and Drummond (2020). Firstly, the constraints posited by Ljung (2010) are far from categorical. In the *prick* example, while *dumb prick* is certainly more common than e.g., *dumb dick* or *dumb cock*, it isn't the case that one cannot replace *prick* with his suggested replacements. It also isn't the case that all other, semantically unrelated swearwords would be suitable replacements in non-literal uses. One would be unlikely to say *shit you!*, for example. The argument posed by Ljung (2010) also belies the fact that languages change; there is no obvious linguistic reason why e.g., *dumb pecker* couldn't become widely used.

Secondly, I agree with Drummond (2020) that, taking a common-sense approach, the literal uses of taboo words like *fuck* and *shit* would be considered as swearing by the average person. If a child swore at the dinner table and was scolded by a parent for swearing, it is unlikely that they could argue their case by claiming that literal uses don't count as swearing. Ultimately, all uses of swearwords are dysphemistic; a person opts to use a swearword where a non-taboo term could have been used. This is true for both literal and non-literal uses.

The third criterion concerns the tendency for swearwords to appear in set phrases in which the swearword cannot be easily replaced with a neutral word. As Ljung (2010) explains, many occurrences of swearing involve set idiomatic phrases such as 'For fuck's sake!' or 'the shit hit the fan'. Many swearwords can appear outside of such phrases however, behaving in many ways like neutral words; for example, swearwords like *damn* or *fucking* can be used in attributive position in determiner phrases like *the damn/fucking dog*. Finally, the fourth criterion reflects a fact central to semantic analyses of swearing, namely that swearwords convey meaning from the perspective of the person uttering them (Potts, 2007; cf. J. A. Harris & Potts, 2009). The emotive nature of swearwords is also reflected in their consistent rating as some of the most highly arousing words in English (Janschewitz, 2008; Warriner et al., 2013).

A considerable number of pages could be, and have been elsewhere, devoted to the task of defining swearing. While there is merit to doing so, it does not serve this thesis a great deal of good. This thesis concerns linguistic variation within swearwords, not across swearwords or between instances of either swearing or not swearing. Comparing the 'offensiveness' of different swearwords has received considerable attention in the liter-

ature already (see Beers Fägersten, 2012 for a review) and numerous studies (see Section 2.4) have already tested the binary distinction of swearing vs not swearing in perception. While a precise definition of swearing is vital to such studies, this thesis intends to deal with only those swearwords that are unambiguously so.

Furthermore, while the criteria provided by Ljung (2010) provide some insight into what makes a swearword, it is really individual speech communities that decide what counts as a swearword. Although there might be some merit to including or excluding swearwords from a study on the basis of such criteria, a similar common-sense approach to that used by Drummond (2020) would suggest that they are not necessary to decide whether words like *fuck*, *shit* or *cunt* are swearwords in British English; they are unambiguously so.

Instead of using Ljung's (2010) criteria, therefore, I will instead rely on words that are prohibited in the speech variety of interest, namely British English. The Office of Communications (Ofcom), the government's regulatory and competition authority for broadcasting, published a 2016 guide to offensive language in broadcast media (Ofcom (Office of Communications), 2016). This guide was informed by research by Ipsos Mori with participants across the UK. Importantly, the guide distinguishes between discriminatory and non-discriminatory language. The category of discriminatory language includes offensive language that targets particular groups, including racial and homophobic slurs. The category of non-discriminatory language includes words commonly considered to be swearwords. I will use the list of non-discriminatory language as the sole inclusion criterion for swearwords in this thesis. The list reflects swearing usage in British English specifically and, as such, is a more reliable indicator of the "social codes" (Limbrick, 1991) that determine swearing.

2.3 Swearing in linguistics

2.3.1 Phonetics

Swearing has received minimal attention in the study of phonetics, with only a few studies of note. Yardy's (2010) unpublished undergraduate dissertation concerns the sound symbolism of swearwords. Using wordlists from Christmas carols, heavy metal music and lullabies as comparisons, Yardy found that swearwords were significantly more likely to contain harsher consonants and significantly less likely to contain more sonorant consonants.

A study by Reilly et al. (2020) shows that the phonetic features of a word contribute to its probability of being rated as taboo and its plausibility in novel swearing compounds. Firstly, in a tabooness rating task, in addition to several other factors including topic and

semantic concreteness, words were significantly more likely to be rated as taboo if they contained obstruents. In a second experiment using novel swearing compounds (e.g., *shit-gibbon*), Reilly et al. showed that acceptability ratings improved if the non-taboo part of the compound contained an obstruent; ratings also improved for compounds with fewer syllables. Similar work by Tessier and Becker (2018) showed swearing compounding to be sensitive to vowel harmony, but not consonant harmony. That is, swearing compounds are more "satisfying" if they have matching vowels (e.g., *fuck-puffin*) but not if they have matching consonants (e.g., *fuck-frisbee*). A further finding from Aryani, Conrad, Schmidtke, and Jacobs (2018) suggests that words with short vowels, voiceless consonants and hissing sibilants are more likely to be rated as highly arousing and negatively valent, the typical emotional profile of a swearword.

Finally, Gold and McIntyre (2016) show that the acoustic-phonetic properties of swearword productions can influence its implicated meaning. Taking an acoustic-pragmatic approach, Gold and McIntyre analyse all occurrences of *fuck* from a single scene of the TV programme *The Wire*, in which *fuck* is uttered 29 times. After categorizing each occurrence into one of five functional categories - *disbelief*, *insult*, *functional*, *surprise/realisation* and *idiomatic* - the authors extracted the durations of each /ʌ/ vowel. Results suggested that the longest durations of /ʌ/ are found when *fuck* was used in disbelief or surprise, while the shortest were found for insults. It is worth noting however that Gold and McIntyre's small token count restricts their statistical analysis.

While work on the phonetics of swearwords is minimal, the above studies do indicate a degree of phonetic patterning within swearing (Yardy, 2010). The findings of Reilly et al. (2020), Tessier and Becker (2018) and Aryani et al. (2018) suggest that speakers have an awareness of which phonetic features are more or less likely to occur in swearwords. Furthermore, the findings of Gold and McIntyre (2016) suggest that the word-specific acoustic phonetic properties are connected to that word's intended meaning in context. This would suggest that phonetic variation within a swearword is pragmatically, and therefore potentially socially, meaningful.

2.3.2 Morpho-syntax

Work on the morphology and syntax of swearing has been predominantly descriptive. Linguists in these disciplines have focused on showing how swearwords can be used in innovative ways in words and sentences and on examining the ways in which they are constrained by the underlying grammar of the language. Reviewing the observations made in morpho-syntax therefore serves the purpose of further illustrating the significant linguistic variability of swearing.

There are two notable cases of what has been termed 'expressive morphology' (Zwicky

& Pullum, 1987) that involve swearwords, namely infixation and compounding. Both involve a consistent trend in a language to derive words using all or parts of a swearword. Expressive morphology, it is argued, is "associated with an expressive, playful, poetic or simply ostentatious effect of some kind" (Zwicky & Pullum, 1987, p. 335). The nature of this expressive effect is imprecise, as I will further discuss in the following section on semantics in relation to expressive content. Importantly however, the adding of expressive morphology adds something to the pragmatics of the word's usage.

Swearing infixation involves placing a swearword in the middle of another word, such as in *fan-bloody-tastic* or *Ala-fucking-bama*. While other languages make use of infixation for inflectional purposes, the only forms of infixation in English are expressive, with swearing infixation the most common example. The constraints on swearing infixation have been discussed by several linguists including Siegel (1974) and Aronoff (1976), with the most prominent contribution having been made by McCarthy (1982). While his account will be discussed in greater detail in Chapter 5, McCarthy's (1982) central claim is that swearing infixation is restricted to the position immediately before a metrical foot; this generalization accounts for the well-formedness of *fan-fucking-tastic* and the ill-formedness of *fanta-fucking-stic*, as well as the well-formedness of both *un-fucking-believable* and *unbe-fucking-lievable*.

The other notable case of expressive morphology involving swearwords is expressive compounding (Meibauer, 2013). This process involves the combining of an evaluative morpheme with a neutral one. Meibauer (2013) cites a number of examples of expressive compounding in German, including *Arschgesicht* ('arse face') and *Reformscheisse* ('reform shit'). Just as in other cases of swearing, the same expressive morpheme can have pejorative (*Arschgesicht*) and meliorative (*Arschgut* - 'arse good') connotations. As Meibauer (2013) argues, new coinages of words with expressive morphemes are easily possible, citing the relatively recent addition of *Rattenscharf* ('rat good'). Similar constructions are also possible in English, such as *fuckload* and *fuckton*, as well as dialect-specific examples such as discourse *-ass* expressions in African American English (e.g., *ugly-ass* or *sweet-ass*, the second of which is comparable to the German *Arschgut*) (Irwin, 2015). Phrases like *shitgibbon*, which are distinct from the taboo measure phrases like *fuckload*, have their own rules of grammar, with compounds involving animate entities (e.g., *fuckpuffin*), receptacles (e.g., *fucksack*) and body-parts (e.g., *shithead*) preferred (Reilly et al., 2020). The examples of both infixation and compounding suggest that swearwords can be used in interesting ways in word formation to convey expressive meaning.

Hoeksema (2019) discusses numerous examples of the innovative uses of swearing in sentence structure and, importantly, the constraints placed on this innovation by the grammar. There are a number of fixed constructions that require an evaluative taboo

term. For example, in evaluative vocatives like 'you're a piece of X', X must be a taboo term like *shit* and cannot be a non-taboo term like *dentist* (Corver, 2008). Swearwords are available in English in a number of other grammatical uses, including emphatic *wh*-questions (*what the hell...*), negative polarity items (*I don't give a fuck*), complex degree expressions (*stupid as fuck*) and degree resultatives (*scared shitless*) (Hoeksema, 2019). Across a sentence, swearwords are 'grammatically versatile', appearing as both a primary and secondary predicate, objects, modifiers and intensifiers (Napoli & Hoeksema, 2009). While a full account of the syntax of swearwords would be a useful endeavour, this thesis does not deal directly with syntactic variation. The examples cited here therefore only serve the purpose of showing that swearwords have been studied in this domain and to point the reader to suitable accounts.

2.3.3 Semantics

For all the interesting ways in which swearwords can be used in a sentence, the most important question concerns what they actually mean. That is, what contribution, if any, does a swearword make to the truth-conditional meaning of a sentence. The literature on expressives, the subset of language that includes swearwords, is very large, encompassing views from both semantics and pragmatics. In this section I will review a subset of this literature. The purpose of doing so is two-fold. Firstly, observations from this literature show important linguistic distinctions between expressive (or not-at-issue) content (such as swearwords), on the one hand, and descriptive (or at-issue) content, on the other. These distinctions further our understanding of what swearwords, as a homogenous group, actually do in conversation that neutral terms do not. Secondly, and slightly conversely, observations from this literature point to a degree of heterogeneity within the category of swearing, showing that not all occurrences of the same swearword make the same meaning contribution (be it semantically or pragmatically motivated), which highlights their potential significance for the study of swearing variation.

In reviewing the literature, however, I take no particular position on the optimal way to analyse expressives. This is a much larger question that many have attempted to answer (Karttunen & Peters, 1979; Kratzer, 1999; Potts, 2005; McCready, 2010, among others) and which, given the focus of this thesis on the social perception of swearing variation, I would fail to do justice in attempting to solve. More specifically, I take no position on whether expressives are best analysed using a unidimensional (Schlenker, 2007, 2010) or multidimensional (Potts, 2005, 2007) semantics, as both account for the relevant facts for the level of analysis that I am attempting, namely that expressives typically predicate something that is distinctly attitudinal from the perspective of the speaker (although cf. J. A. Harris & Potts, 2009) and which is imposed in conversation rather than proposed. Al-

though many of the observations about different swearing constructions do originate from the literature that uses a multi-dimensional semantics that considers expressives as conventional implicatures (Potts, 2005, 2007; McCready, 2010; Gutzmann, 2011; Gutzmann & Turgay, 2014), I do not believe that these observations are at odds with the presuppositional analysis proposed by Schlenker (2007, 2010), who explains how his analysis accounts for all of the properties proposed by Potts (2007) for expressives.

2.3.3.1 A homogenous set of meanings

Returning to the question of what swearwords actually mean, consider the examples in (1), where (1-a) contains no swearwords and (1-b) contains the swearword *fucking* as a pre-nominal modifier. It might be easy to think that *fucking* is functioning like some kind of degree modifier in (1-b), intensifying the degree to which the speaker hates Swindon Town FC (Football Club). Other intensifiers do not work like this, however. For the meaning denoted by the verb to be intensified, the intensifier must appear before that verb; the sentence *I hate that really Swindon Town FC* would be ungrammatical.

(1) a. I hate that Swindon Town FC
b. I hate that **fucking** Swindon Town FC

The sentences in (1-a) and (1-b) appear to be truth-conditionally equivalent. That is, both sentences are true in a world in which the speaker hates Swindon Town FC. Obviously, the *fucking* must be adding something to the expression, even if it it not affecting its truth-conditions.

This additional meaning is referred to by many semanticists as *expressive meaning* (Potts, 2005, 2007; McCready, 2010; Gutzmann, 2015). In his multi-dimensional analysis of expressives, Potts (2007) identifies several important properties of expressive meaning that differentiate it from *descriptive meaning*. The principal property of interest is *independence* from truth conditional meaning; this property reflects Kaplan's observation that "truth is immune from epithetical colour" (Kaplan, 1999 in Potts, 2007, p. 167). This property is shown in examples like (2). In conversation, if one speaker utters the sentence in (2-a), their interlocutor can felicitously respond as in (2-b), denying the descriptive content of the utterance, namely the proposition that the Royals lost another game. In contrast, the same interlocutor could not felicitously use the sentence in (2-c) in order to deny the content of the expressive. That is, they cannot deny the speaker's heightened emotion. As such, Potts (2007) suggests that examples like (2) are evidence that expressive content is independent from descriptive content.

(2) a. The **fucking** Royals lost another game
b. No they didn't, they won last night

c. #No, you don't feel like that

Several examples suggest this independence not to be absolute, however. Potts (2007), for example, cites an example of Japanese antihonorifics from Potts and Kawahara (2004), in which the antihonorific takes a proposition from the descriptive dimension as its argument. In the sentence *Nesugoshi-chimat-ta* ('I overslept'), the antihonorific *chimat* takes as it's argument that the speaker overslept, conveying the speaker's annoyance at that fact. Potts (2007) resolves to say that the interaction between the expressive and descriptive domains is one-way, such that descriptive operators do not take expressive meanings as their arguments.

Another important property of expressives is *nondisplaceability*. Expressives always predicate something immediate about the utterance situation. Expressive content is "valid only for the utterer, at the time and place of utterance" (Cruse, 1986, p. 272). It cannot, for example, be used to report a past, future or conditional emotion. The evidence for this property is that, unlike descriptive content, expressive content survives standard presupposition holes, i.e., "operators that cannot cancel or modify the presuppositions triggered by items in their scope" (Karttunen, 1973 in Potts, 2007). We can see this in the examples in (3), where the expressive *fucking* survives under negation (3-a), in the antecedent of a conditional (3-b) and a past tense operator (3-c); expressives are therefore said to project (Simons et al., 2010). In each case, the meaning of *fucking* remains tied to the utterance situation.

(3) a. I didn't **fucking** fail the exam (# I feel neutral about the situation).
b. #If I **fucking** failed the exam, I will feel neutral about it.
c. Yesterday I **fucking** failed the exam (# but I feel neutral in the present)

Similarly, expressives are *perspective dependent*. That is, expressives are always evaluated from the perspective of a particular person, typically the speaker. For example, the *fucking* in (4-a) is tied to the perspective of the speaker. This contrasts with descriptive modifiers in the same syntactic position (e.g., *red* for *fucking* in (4-a)), which are not tied to any one perspective. In other cases, as in (4-b), a non-speaker-orientated meaning of an expressive preferable; it is unlikely that the speaker in (4-b) thinks that *Webster* is a *bastard*, given that they intend to marry him (see Hess, 2018). In this case however, the expressive is still evaluated with regard to some concrete perspective, namely that of the speaker's father.

(4) a. John realized that I'd seen his **fucking** shirt
b. My father screamed that he would never allow me to marry that **bastard** Webster.

(Amaral, Roberts, & Smith, 2007)

Expressives also have the property of *immediacy*, whereby simply the act of uttering an expressive is sufficient for conveying its content. Potts (2007) draws comparisons here between expressives and *performatives*. For example, just by uttering (5), a speaker obligates themselves to wash the dishes (Searle, 1969), meaning that the bracketed continuation is infelicitous. Similarly, by uttering a phrase like *that bastard Kresge*, the speaker commits themselves to the emotive performance expressed by *bastard*; as such, the continuation in brackets in (6) is also infelicitous.

(5) I promise to wash the dishes later (# But I refuse to wash the dishes later)

(6) That bastard Kresge was late for work yesterday. (# But he's no bastard today, because today he was on time.)

(Potts, 2007)

The final two properties proposed by Potts (2007) are *repeatability* and *descriptive ineffability*. Repeatability concerns the fact that increasing the number of expressives in a sentence increases the strength of the emotion expressed by the sentence. Potts cites the examples in (7), in which the increasing number of *damn*s from (7-a) through to (7-c) gradually increases the amount of emotion conveyed by the speaker. (7-c) contrasts with (8), in which the descriptive assertion 'I'm angry!' is repeated. In (8), the added uses of 'I'm angry' are redundant; they do not increase the degree of anger that is conveyed.

(7) a. Damn, I left my keys in the car.
b. Damn, I left my damn keys in the car.
c. Damn, I left my damn keys in the damn car.

(8) I'm angry! I forgot my keys. I'm angry! They are in the car. I'm angry!

Descriptive ineffability concerns the inability of speakers to easily articulate the meaning of expressives, with speakers regularly resorting to examples of the appropriate context for a word's usage (Blakemore, 2002). Such explanations fail to account for the variety of contexts in which the same word can be used. While descriptive content can often be precisely defined therefore, the same is not true for expressive content.

In terms of what expressives actually mean, under Potts's (2005, 2007) system, expressives take as their argument descriptive meanings and give back a pair consisting of the same descriptive meaning and an expressive meaning. In (1-b), repeated in (9), for example, under one reading, the expressive *fucking* takes the descriptive meaning denoted by the proper name *Swindon Town FC* and gives back the same meaning, plus an expressive meaning that relates to the speaker's attitude towards the descriptive meaning.

(9) I hate that **fucking** Swindon Town FC

Unlike other modifiers in this position however, expressive modifiers are not restricted to operating over their immediate syntactic sister (see Potts, 2007). That is, in (9), *fucking* does not necessarily have to combine with the meaning denoted by Swindon Town FC. It can instead operate over the proposition denoted by the whole sentence 'I hate Swindon Town FC'. A discussion of the implications of this contrast between descriptive and expressive modification is given in Frazier, Dillon, and Clifton (2015). We can see this contrast best by comparing the minimally different sentences in (10).

(10) a. The black cat is in the garden
b. The damn cat is in the garden

In (10-a), the denotation of the predicate *black* first intersects with the denotation of the predicate *cat* to give the set of *black cats*. This then combines with determiner *the* to yield *the black cat*, and so on until all parts of the sentence have combined to give the meaning of the whole, per the principle of compositionality. The same is not true for *damn* in (10-b). *Damn* can optionally combine with *cat*, *the cat* or *the cat is in the garden*.

Building on Karttunen and Peters (1979), Potts' formal account of expressives uses a mutli-dimensional type-driven system with two types: at-issue and expressive (see Potts (2005) for full details). The three standard types *e*, *t* and *s* can all be either at-issue (written with superscript *a*) or expressive (written with superscript *c*). In simple terms, when applied to an argument like *cat* in (10-b), an expressive like *damn* would be of type $\langle\langle e^a, t^a\rangle, t^c\rangle$. *Damn* would take *dog*, a 1-place at-issue predicate, and return something of type t^c.

The important aspect of Potts system that accounts for the independence of at-issue and expressive (or not-at-issue) meaning is illustrated in the subtree in Figure (11); this is an example of CI (conventional implicature) application. Potts introduces the bullet •, a metalogical device which separates two independent lambda expressions. The key part of this is the independence of the two expressions. While *damn* takes something of type $< e^a, t^a >$ and returns something of type $< t^c >$, the predicate *cat* is not used up. Instead, the meaning is passed up to the mother node, after which the rest of the sentence can be composed as usual.

(11)

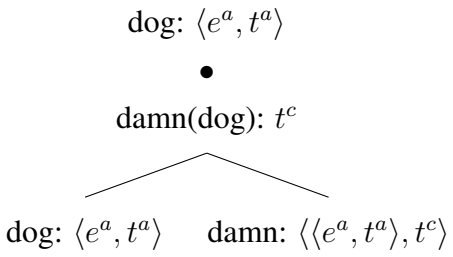

To model the flexibility of expressive adjectives, Potts (2005) allows for syntactic and semantic parsetrees to have different shapes. For example, for a sentence like (12), in which the speaker is likely expressing their heightened emotion about the state of affairs described by the whole sentence, Potts (2005) suggests the syntactic parsetree in (13) and the semantic parsetree in (13). Potts (2005, p. 167) gives a denotation for all *expressive adjectives* in which they can take something of type $\langle \tau^a, t^a \rangle$ and return something of type t^c.

(12) The damn machine didn't come with an electric plug

(13)

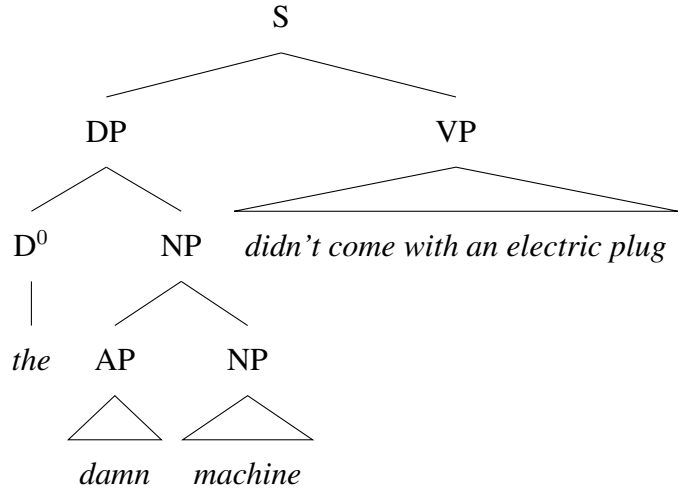

(14)

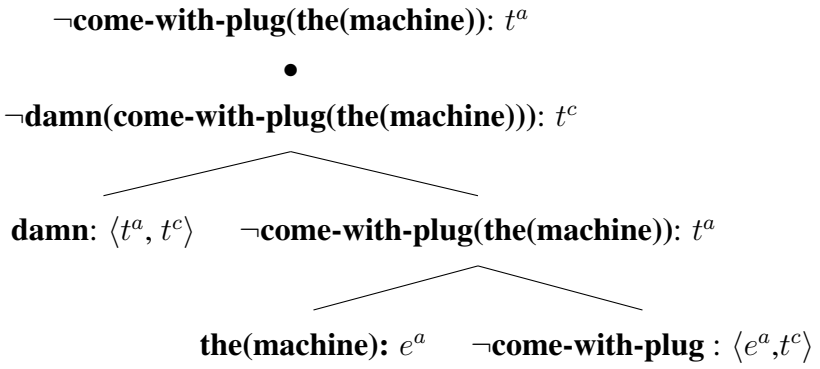

If the type-driven system that Potts (2005) proposes is deterministic, with exactly one legitimate semantic value for a structure, which he suggests it should be (p. 61-62), then the parsetree in (14) is unsuitable, because it doesn't account for purely local readings of expressive adjectives. An alternative posed by Potts, which better accounts for the flexibility of interpretation of expressives, is to treat expressive adjectives like *damn* and *fucking* as *isolated CIs*. The rule for isolated CIs given by Potts (2005) is provided in (15). Potts offers an example subtree for (16), given in (17), in which the expressive *fucking* is treated as an isolated CI.

(15)

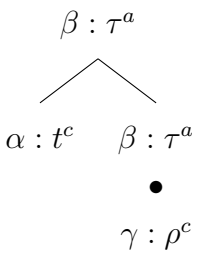

(16) Sue's dog is really fucking mean

(17)

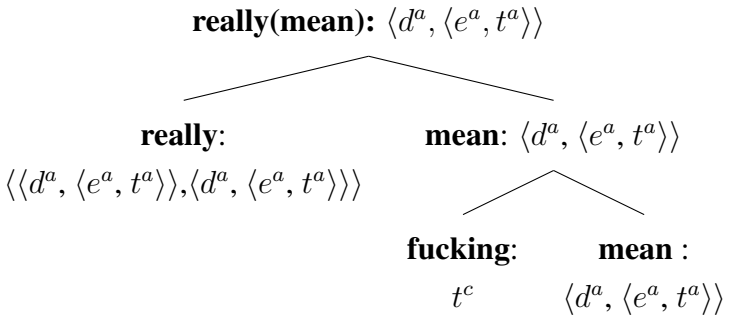

This rule posits no interaction at all between expressive adjectives and at-issue content. Rather, the expressive is interpreted independently as a general expression of the speaker's heightened emotion. The interpretation of an expressive adjective with respect to a particular part of the sentence is therefore pragmatically-derived, rather than a function of semantic composition.

I am proposing to extend this to sentences like (10-b), repeated in (17). The expressive *damn* is treated like an *isolated CI*, behaving independent of the rest of the sentence, contributing a general heightened emotion about the state of affairs. The rest of the sentence composes as expected. The interpretation of *damn* with respect to the rest of the sentence - i.e., whether it is interpreted to be directed towards *cat* or towards the general state of affairs - arises from a listener's pragmatic reasoning over this and alternative sentences (this will be developed in more detail in Chapter 6).

The damn cat is in the garden

The flexibility of interpretation of expressives like *damn* was tested experimentally by Frazier et al. (2015). Frazier et al. presented participants with a series of sentences with *damn* in different positions, for example one of the three in (18), before explicitly asking participants "Which is the speaker most likely to have a negative attitude toward?', with example options being 'the holiday', 'the weekend' and 'the holiday being on the weekend'; these options were known as the *subject*, *object* and *sentence* interpretations respectively. An additional manipulation was used to test whether responses further depended on causality. Such sentences included an animate subject (e.g. *dog*).

(18) a. The holiday is on the damn weekend.

b. The damn holiday is on the weekend.
c. Damn. The holiday is on the weekend.

Overall, Frazier et al.'s participants provided very varied responses. Sentence-level interpretations were chosen for sentences with *damn* in subject and object position (e.g. (18-b) and (18-a)) 39% and 42% of the time, suggesting significant variability in how *damn* was interpreted with respect to the rest of the utterance. This would not occur with an attributive adjective like *black* in a sentence containing the DP *the black cat*, where only a local interpretation is available. In causal sentences, the frequency of sentence-level interpretations dropped significantly, following what Frazier et al. (2015) call the *culprit hypothesis*; the presence of a causal agent provided a more reasonable entity to blame for the situation. Flexibility of interpretation is therefore arguably another important property of expressives, with other factors such as causality also potentially playing a role in how each usage of a swearword is interpreted.

There are limits to the flexibility of expressives however, as has been pointed out by Gutzmann (2019). Firstly, in cases of embedding, expressives are limited in how far up the structure they can go to find their argument. In (19), while there are available readings under which *damn* applies to either *holiday* or *the holiday is on the weekend*, it is not possible for *damn* to apply to either *Mary* or *Mary said that the holiday is on the weekend*. Using a similar design to Frazier et al. (2015), but for German, Gutzmann (2019) showed that sentence-level readings are significantly less available when *damn* is in an embedded clause, as in (19).

(19) Mary said that the damn holiday is on the weekend

Furthermore, the direction of the flexibility appears to be unidirectional. A sentence like (18-a) is ambiguous between object- and sentence-level readings. For sentences like (18-c), however, there is much less ambiguity, with the sentence-level reading dominant. this is emphasised by Gutzmann (2019) using the examples in (20), in which a sentence adverb with a positive evaluation is inserted into the sentence.

(20) a. #Damn! Luckily, the dog has eaten the cake
b. Luckily, the damn dog ate the cake

Inserting *luckily* blocks the sentence-level reading of *damn*. Gutzmann poses a special context in which "somebody dislikes the dog and has a vicious plan which involves the dog eating the cake so that the dog gets into trouble" (p. 104-105). In this context, (20-a) is contradictory; there is no available reading under which *damn* targets either *dog* or *cake* and there is no available sentence-level reading. (20-b) is not contradictory, however; the local reading is available even when the sentence-level reading is blocked. Gutzmann

(2019) suggests that expressive adjectives must look up to get their interpretation (see Chapter 4 of his book for a full syntactic account of expressive adjectives).

2.3.3.2 A heterogeneous set of uses

Swearwords share all of the properties identified by Potts (2007) for expressives, as well as the flexibility of interpretation. They predicate the heightened emotion of the speaker in the context of an utterance in a way that cannot be given an exact definition; this heightened emotion can target a variety of entities in a sentence. In this way at least, swearwords could be said to be a homogenous group. What the exact meaning contribution of each usage of a swearword is in context, however, is more variable. In extensions to Potts's (2005, 2007) account of expressives, the picture becomes more complicated, with the identification of examples of mixed expressives. expressive modifiers and main content.

As already discussed, there are exceptions to the independence of expressive content from descriptive content. One such exception is *mixed expressives* or *mixed content* (McCready, 2010; Gutzmann, 2011). While the expressive *fucking* in e.g., *that fucking Swindon Town FC*, only conveys expressive content, some expressives are said to consist of both expressive and descriptive meanings. The oft-cited example is a slur that appears in predicate position, as in (21). For (21) to be a well-formed sentence, the predicate *kraut* must convey something descriptive when combined with the subject, namely that Hitler was German. In addition, it conveys an expressive meaning, namely the derogation of Germans more generally (McCready, 2010; Gutzmann, 2011).

(21) Hitler was a kraut

(Saka, 2007)

Any use of an expressive that also conveys descriptive meaning could be called a mixed expressive. The case could be made, therefore, that some examples of swearwords are also mixed expressives. One example is literal usages, as in (22). The use of *fucked* in (22-a) must convey descriptive content, or else the sentence would not be well-formed, which it clearly is. The use of *fucked*, rather than the functionally equivalent descriptive expressions *had sex with*, also conveys some expressive meaning, relating to the speaker's emotion about the fact that Mary had sex with John. Similarly, the use of *cunt* or *cock* conveys the descriptive meaning of their respective genitalia, as well as some expressive meaning. The speaker in each case opted for a taboo expression when neutral terms were available. This may express heightened emotion or something about their perspective on the use of these words.

(22) a. Mary fucked John
b. I saw Mary's cunt/John's cock

Another example cited by McCready (2010) as a mixed expressive is *expressive intensification*. A non-swearing example of expressive intensification is the Viennese German intensifier *ur*, as in the sentence in (23-a), with an English paraphrase in brackets. *Ur* has been argued to have both descriptive and expressive meaning components (McCready & Schwager, 2009). Its descriptive contribution is the at-issue intensification of the meaning denoted by *interessant*, meaning that the referent of *that* is *extremely/really interesting*, rather than just *interesting*. Its expressive contribution is heightened commitment from the speaker to the at-issue proposition. The potentially taboo example *sau* (female pig) is discussed by Gutzmann and Turgay (2014), as in (23-b), in which *sau* performs at-issue intensification of *schnell*, as well as signalling not-at-issue speaker commitment to the proposition. Interestingly, the subclass of German expressive intensifiers that Gutzmann and Turgay discuss also have interesting syntactic properties, appearing both internally and externally to a DP (e.g., *der sau coole Party* vs *sau der coole Party*); in both cases, the authors argue, *sau* intensifies the meaning denoted by the adjective *cool*.

(23) a. Das ist **ur** interessant (That is *ur* interesting)
b. Das Ding is **sau** schnell (The thing is *sau* fast)

A similar analysis could be applied to examples involving swearwords in English. Consider, for example, the sentence in (24), where *fucking* appears immediately before the gradable predicate *tall*.

(24) John is **fucking** tall

It could be said that *fucking* here mirrors the functions of *ur* and *sau*, doing at-issue intensification of the meaning denoted by *tall* and expressive intensification of the whole proposition. Exactly what the mechanisms are that lead to this reading of (24) - whether they be semantically or pragmatically driven - will be further discussed in Chapter 6, in which semantic variation is the main focus. The observations of McCready (2010) and Gutzmann and Turgay (2014) regarding expressive intensification suggest that swearwords might, in some uses, affect the descriptive, and therefore truth-conditional, meaning of a sentence.

Gutzmann (2011) also cites examples of *expressive modifiers*, that is, expressives that take as their arguments other expressives and give back expressive meanings. An example from Gutzmann (2011) is given in (25). While Potts (2005) claimed that expressives could only take descriptive meanings as arguments, examples such as these suggest that expressives can also operate over other expressives.

(25) That fucking bastard Burns got promoted.

In other cases, expressives don't appear to modify anything in particular, instead serving as independent expressions that contribute exclusively expressive content, termed *main content* (McCready, 2010). McCready (2010) discusses examples such as (26) from Kratzer (1999). Such expressions do not assert anything and thus have no truth conditions; rather, they have what are described as *use-conditions*, that is, conditions in which their usage is appropriate (Gutzmann, 2015). Expressions like (26) express something about the speaker's mental state about the state of affairs. The same should also be true of a swearword like *fuck* in (27).

(26) Ouch!

(Kratzer, 1999)

(27) Fuck!

For all the properties shared by swearwords, as a group, by virtue of being expressives, the deployment of a swearword in a sentence does not always appear to make the same contribution to that sentence's meaning. Irrespective of whether you take a unidimensional (Schlenker, 2007, 2010) or multdimensional (Potts, 2005, 2007; McCready, 2010; Gutzmann, 2015) view on expressives, there is undeniably variation in the types of meanings that swearwords contribute to an expression. All swearwords predicate speaker-oriented meanings. Some swearwords provide an expressive meaning that accompanies the at-issue meaning in the rest of the sentence, e.g., (1-b). Some swearwords contribute, or appear to contribute, both at-issue and expressive meanings at the same time, e.g., (22) and (24). Some swearwords have no interaction with at-issue meanings at all, contributing only expressive meanings, e.g., (27). Furthermore, as shown in the previous section, the syntactic position of a swearword in a sentence does not automatically dictate how it will be interpreted with respect to the rest of the sentence. While there may be some semantic homogeneity among swearwords therefore, the examples discussed in this section show that there is also significant intra-swearword variation when it comes to usage.

2.3.4 Pragmatics

Work on the pragmatic elements of swearing is typically situated around the concept of *impoliteness*. Building on Brown and Levinson's (1987) theory of politeness, that is, the manner in which interlocutors maintain or promote social harmony, Culpeper (1996) presents a parallel but opposite model for impoliteness, that is, the manner in which interlocutors engage in strategies for social disruption. In it, Culpeper (1996) suggests that, in a minority of cases, including the use of taboo words, language use can be inherently impolite. The use of 'taboo words' was listed by Culpeper under *positive impoliteness out-*

put strategies, that is, strategies designed to damage the addressee's positive face wants. Culpeper (2011, p. 23) later defined impoliteness thusly:

> *Impoliteness is a negative attitude towards specific behaviours occurring in specific contexts. It is sustained by expectations, desires and/or beliefs about specific organization, including, in particular, how one person's or a group's identities are mediated by others in interaction. Situated behaviours are viewed negatively - considered 'impolite' - when they conflict with how one expects them to be, how one wants them to be and/or how one thinks they ought to be. Such behaviours always have or are presumed to have emotional consequences for at least one participant, that is, they cause or are presumed to cause offence.*

In his more recent work, however, Culpeper (2018) explains that describing swearing as just 'impolite' is a significant oversimplification, particularly given the variety of uses and interpretations it can receive. Culpeper (2011) suggests that a better characterisation of swearing is in terms of sociality rights (Spencer-Oatey, 2008), that is, the expectations that people have about whether other people should or should not do certain things in certain contexts. When people come into conflict over such rights, there can be a feeling of injustice, immorality or, at the very least, a feeling of a lack of fair consideration. The restriction to 'certain contexts' is what allows for instances of swearing that are not intended to cause offence or perceived to be impolite. For example, Daly et al. (2004) suggest that instances of swearing in New Zealand are frequently used to express solidarity. Although anecdotal, I myself experienced this first hand during my time playing rugby, where one coach from New Zealand would frequently refer to players he approved of as 'good cunts'. Drawing on work by Terkourafi (2008), Culpeper (2011, p. 29) refers to the *schematic knowledge* that speakers of a language have of conventionalised meanings, that is, "particular expressions [which] are associated in one's mind with particular contexts"; this is similar to the use conditions referred to by Gutzmann (2015).

Furthermore, there is significant variation with respect to impoliteness across different taboo terms. Slurs, for example, appear to be inherently impolite, as they serve to explicitly derogate an individual and their membership to a particular group (e.g., a racial minority group). Common swearwords such as *fuck* and *shit*, however, have no direct connection with a particular identity. On this basis, swearwords would not be inherently impolite. As Culpeper (2018) points out however, words like *twat* denote a female body-part and, as such, could be said to have a metonymic connection with women; it is also likely that they have misogynistic flavour, too (see Sobieraj, 2018 on the use of gendered abuse terms like *cunt* and *bitch* in online hate directed towards women).

As Dynel (2012) argues, therefore, impoliteness may be better served as a concept

that reflects the views of a particular speech community. This is also reflective of the Limbrick (1991) quote I discussed at the start of this chapter in relation to the "social codes" that dictate what is and is not considered swearing. Within particular communities of practice (e.g., a sports team or a building site, or a particular geographical region), as in Daly et al. (2004), swearing can become commonplace and is only perceived to be offensive by a person not belonging to that community. It is therefore important that models of impoliteness capture both the speaker's intention and contextual factors as 'co-determinants' (Dynel, 2012).

Both intention and context have been addressed in research on perceptions of the speaker as a function of swearing (T. Jay & Janschewitz, 2008; Stapleton, 2020). The findings of T. Jay and Janschewitz (2008) and Stapleton (2020) are also relevant to the later discussion of sociolinguistic perception and swearing in Section 2.4.2 of this chapter, but they bear discussion at this point due to their explicit reference to the pragmatic aspects of swearing that give rise to particular social evaluations of the speaker.

Using experimental methods, T. Jay and Janschewitz (2008) asked participants about the likelihood that different swearwords would be uttered by speakers of different statuses (dean vs student vs janitor) and in locations of different levels of formality (dean's office vs dorm room vs parking lot). T. Jay and Janschewitz' (2008, p. 283) main finding is that people are "sensitive to pragmatic variables underlying swearing". Interaction effects between speaker and location showed that swearing was most appropriate for different speakers in different locations (e.g. with the dean in their office, with the student in their dorm). Using a between subjects variable of English experience (native vs non-native), results also showed that native speakers are more sensitive to these variables, with offensiveness and likelihood ratings from non-native speakers showing significantly less variability. While natives and non-natives gave the selected swearwords similar ratings overall, suggesting that non-natives are familiar with their taboo nature, their lower sensitivity to changes in context suggest they do not appreciate 'the diversity of the social functions the word can take on' (T. Jay & Janschewitz, 2008, p. 285). This finding speaks to the contextual complexity of swearword usage.

Across a number of survey studies, Dewaele (2004, 2010, 2015, 2016, 2017) has also shown that non-native speakers typically have different feelings about the (im)politeness/ offensiveness of swearwords. In one study, the effect of nativeness on perceptions of swearwords is shown to be impressively linear; speakers typically find swearwords in their native language to be the most offensive, followed by swearwords in their L2, then their L3, and so on, with speakers of five languages in the study finding swearwords in their L5 the least offensive (Dewaele, 2004); in another study, situational variables were shown to have a smaller effect on the self-reported frequency of swearing of L2 speakers of a language (Dewaele, 2017) . Even among multlinguals who self-report similar profi-

ciency and frequency of use in their L1 and L2, swearwords in the L1 are perceived to be stronger (Dewaele, 2010). In other cases, foreign language users have been reported to overestimate the offensiveness of milder swearwords in their L2 (Dewaele, 2016). Finally, even among speakers of the same language, the exact variety of the language spoken can affect which swearwords are deemed the most offensive. For example, *bollocks* is rated as more offensive in British English than American English, while the reverse is true for *jerk* (Dewaele, 2015).

Returning now to 'intention', listeners also form judgments about speakers who swear based on what they perceive their motivations to be. Pragmatic inferences drawn about speaker intentions, drawn from their swearing, can lead to context-specific evaluations of the speaker. Stapleton (2020) uses a Discursive Psychology approach to analyse online responses to a news report about a celebrity swearing on television. The report concerned Stephen Fry's use of *pissing* and *fucking* while presenting the British Academy of Film and Television Arts (BAFTA) annual awards ceremony. As well as comments linking his swearing to his perceived lack of intelligence/common sense, as well as his social class status, several comments take as his motivation for swearing the desire to be seen as funny and/or cool, and 'down with the kids'. The comments position 'swearing' as something that could be perceived to be funny or cool in particular communities of practice - i.e., among young people - but which, coming from an older man like Fry, is perceived to be inauthentic or contrived (Stapleton, 2020). In assigning social meanings to speakers' swearing therefore, the intention behind the swearing may be inferred based on the hearer's pre-existing knowledge of the speaker.

This relatively brief summary of research from pragmatics on swearing serves to show, as one would expect, that the nature of swearing depends on context; perhaps it is this context-dependence that prevents swearing from receiving a concrete definition in the literature, as discussed in Section 2.2 of this chapter. Although it has classically been grouped under the umbrella of impoliteness, swearing has the potential to be used to mark solidarity between interlocutors. Models of swearing must incorporate the relevance of particular swearing expressions to particular communities of practice. Furthermore, swearers have intuitions about the contexts in which swearing is more or less appropriate, as well as the potential intentions of others' swearers' language choices. In sum, swearing is highly variable across different contextualised usages.

2.3.5 Psycholinguistics

As mentioned in the introduction, research on swearing in psycholinguistics - and psychology more broadly - can improve our understanding of how humans process language and the linguistic factors that play a role. The various studies reviewed in this section have

repeatedly shown that the unique emotional characteristics of swearwords that set them apart from neutral words - namely their properties of arousal, valency and tabooness, among others - have consequences for how they are processed by a listener or reader, particularly with regard to attentional resources.

I will begin by reviewing the empirical data that shows that swearwords have these properties. Work on emotion typically shows that emotional experiences consist of at least two qualities, namely *valency* and *arousal* (Barrett, 1998), with some also arguing for the inclusion of *dominance* (Wyczesany & Ligeza, 2015). Valency concerns the positivity or negativity of an emotion. An example of negative valence is *sadness*, while an example of positive valence is *happiness*. Arousal concerns psychological arousal or activation. An example of a high-arousal emotion is *anger*, while an example of a low-arousal emotion is *calm*. Finally, dominance concerns the degree of control a person has over the emotion they feel. For example, an example of a low-dominance emotion would be *confidence*, while an example of a high-dominance emotion would be *submissiveness*.

Although there is significant within-group variation among swearwords along these three dimensions, studies show that swearwords do pattern in particular directions as a group. In the only swearing-specific study of its kind, Janschewitz (2008) completed a norming study of 460 English words, including 92 *taboo* words in addition to a mixture of emotionally valent and neutral words. The study was completed by 78 American native English-speaking undergraduate students. Among other characteristics, participants were asked to rate words on scales from 1 to 9 depicting *tabooness*, *valency* (happy-sad) and *arousal*. Consistently, the words rated as the most taboo were also rated as significantly more arousing and negatively valent (see also Reilly et al., 2020). In a more general word-rating study by Warriner et al. (2013), swearwords were also found to be significantly more arousing, negatively valent and dominant than average. The norms provided by Janschewitz (2008) and Warriner et al. (2013) are both frequently used in studies on word processing.

The unique properties that swearwords possess, namely their levels of valency, arousal and tabooness, affect speaker performance in both linguistic and non-linguistic tasks. The literature on this topic is extensive. The effects for Stroop tasks, recall tasks and lexical decision tasks have been repeatedly replicated and explored at increasingly finer levels of detail. In each case, I will review the most up-to-date results, as well as referencing previous findings that either provide further support or which differ in some respect.

A common non-linguistic task in which the presence of a emotional stimuli, including swearwords, affects performance is the Stroop task. Participants are presented with individual words on a screen. The words are presented in a variety of different coloured fonts. The participants are required to identify the colour, with their reaction time measured for each trial. The typical experimental manipulation in a Stroop task is the type of word

presented. The challenge to participants is to ignore the content of the word itself and concentrate on the colour. In the original version of the task (Stroop, 1935), participants were slower to identify the colour of the font if the word presented referred to a contradictory colour (e.g., the word *red* in blue font) than if the word matched the font (e.g., the word *blue* in blue font). The finding was taken to a suggest that the meanings of words are processed relatively automatically, in a way that affects other co-occurring cognitive processes such as colour-naming. That is, the content of a word is processed even when it is not relevant to the task at hand, taking up attentional resources that could have been used to complete that task.

The Stroop task has also been frequently employed to test the effects of emotional words on attention. In these cases, researchers have tested the properties of arousal, valency and tabooness on participants' performance in Stroop tasks (Siegrist, 1995; Algom, Chajut, & Lev, 2004; MacKay et al., 2004; Eilola & Havelka, 2011). The words included in the task are grouped categorically as high or low arousal, high or low valency and high or low tabooness, with the addition of neutral words as controls. Consistently, these studies have shown that performance in an emotional Stroop task is negatively affected by words of high arousal, low valency and high tabooness. The most recent iteration of this research is Eilola and Havelka (2011), who combine the Stroop task with a measure of skin conductance response (SCR), a physiological indicator of arousal; the authors also compare native and non-native speakers of English using English stimuli. Both negatively valent and highly taboo words caused slower reaction times during the Stroop task compared to neutral words; this effect was consistent for native and non-native speakers of English. They also found higher SCRs for negatively valent and highly taboo words compared to both positively valent and neutral words, indicating that swearwords also elicit a physiologically detectable response that neutral words do not; this effect was only present for native speakers of English, however, replicating previous work that people are less sensitive to swearwords in their L2s (C. L. Harris, Ayçíçeğí, & Gleason, 2003).

Two notable variations on the Stroop task have also shown that swearwords, compared to neutral words, can interfere in other tasks. The first is an auditory hybrid of the Stroop task. Bertels, Kolinsky, Pietrons, and Morais (2011) recorded four different native speakers of French producing individual words, using the same manipulations of arousal, valency and tabooness as e.g., Eilola and Havelka (2011). Upon hearing each word, participants were required to identify the speaker when presented with one of four names (having been familiarised to each named voice pre-task). The task thus tests the effect of emotional words on voice identification. In blocked trials - e.g., all taboo words presented together, all low valency words presented together etc. - words that were negatively valent and highly taboo caused slower voice identification compared to neutral words. In mixed trials, while there was no within trial effect, words that were presented immediately af-

ter a negatively valent or taboo word received slower responses for voice identification. The results suggest that the effect of swearwords on attention also occurs in the auditory domain.

A second related task is the picture-word interference task used by Dhooge and Hartsuiker (2011). Dhooge and Hartsuiker presented participants with both a picture and a word simultaneously on a screen. Participants were required to name the picture while ignoring the accompanying distractor word. Using both taboo and neutral words, Dhooge and Hartsuiker found that picture naming was significantly slower on trials with taboo distractors. The accuracy of the responses on those same trials improved compared to those with neutral distractors, however. The authors suggest their findings are best interpreted in terms of a verbal self-monitor, claiming that, due to the embarrassing/offensive nature of the taboo words, participants have a greater desire to avoid errors and are thus more careful, resulting in slower but more accurate responses. Similar results were found in an attentional blink task by Mathewson et al. (2008).

Related findings come from tasks with a direct linguistic focus, that is, tasks in which participant focus is on the swearword itself, rather than the swearword acting as a potential distraction from another task. The first of these concerns word recognition, namely the lexical decision task. In this task, participants are presented with words on a screen one at a time. On each trial, participants indicate via button press whether that word is an acceptable word in a given language. Longer response times are taken as an indicator of higher processing costs. Using lexical decision data for over 12,000 words from the English Lexicon Project (Balota et al., 2007), Kuperman, Estes, Brysbaert, and Warriner (2014) showed that valency and arousal each, independently, predict slower response times; for both properties, the effects are stronger among low frequency words than high frequency words. Further ERP evidence is provided by Citron, Weekes, and Ferstl (2013), who find a larger amplitude in early posterior negativity for high arousal and negatively valent low arousal words.

The most comprehensive account of the effect of tabooness on lexical decision is Madan et al. (2017). Using a combination of highly arousing taboo words, moderately arousing negatively valent words, moderately arousing positive words and emotionally neutral words, Madan et al. (2017), replicated previous findings of slower lexical decision times for taboo words compared to neutral and positive words (MacKay et al., 2004; L. Thomas & LaBar, 2005), but not compared to other negative words. After exploring the effect in greater detail, slower lexical decision responses were mostly explained by word frequency, familiarity and personal usage of the word. Once these factors were accounted before, only tabooness and not arousal or valency led to a significant slowdown in responses; this was contrary to previous findings that valency also affects response times independent of non-emotional factors (Kousta, Vinson, & Vigliocco, 2009; Vinson,

Ponari, & Vigliocco, 2014).

The language in which the swearwords are produced appears to modulate the effect of tabooness, however. Combining a lexical decision measure with an fMRI measure, Sulpizio et al. (2019) compare the processing of swearwords by native and non-native speakers of English. In their behavioural results, the authors found that taboo words only led to a slowdown in lexical decision responses for L1 speakers, not L2 speakers. Furthermore, in their fMRI results, swearwords led to lower activations in a range of areas of the brain than non-swearwords, with the authors concluding that "[l]ess effort is needed to process taboo words than non-taboo words, suggesting that, similarly to emotional information, socio-pragmatic knowledge is automatically accessed and facilitates word recognition" (Sulpizio et al., 2019); this effect is also modulated by the language of the swearing, with swearing in an L2 being more effortful to process.

As well as affecting word recognition, emotional word properties can also affect sentence processing. Following work by Scott, O'Donnell, and Sereno (2012), Knickerbocker, Johnson, and Altarriba (2015) tested eye-movement associated with emotional words embedded in full sentences as a measure of reading speed. While negative emotion words appear to lead to slowdowns in individual word recognition, the results of Knickerbocker et al. (2015) suggest that they are read faster in full sentences compared to neutral words; the same was shown for positive emotion words. That is, words like *distressed* and *happy* were read faster than words like *chair* (with a variety of other linguistic factors controlled for including lexical frequency) (see also M. Bayer, Sommer, & Schacht, 2010 for ERP differences between single-word and sentence processing of emotional words). Notably however, no study has tested the specific effect of tabooness in sentence processing.

Finally, multiple studies have explored the effect of tabooness on recall and memory (Grosser & Walsh, 1966; MacKay et al., 2004; Buchanan, Etzel, Adolphs, & Tranel, 2006; Hadley & MacKay, 2006; Guillet & Arndt, 2009; Madan et al., 2017). The most frequently found effect is a recall advantage for emotional and arousing words; such is the strength of this advantage that, when asked to recall randomised lists of emotional and neutral stimuli, memory for words immediately preceding and following emotional words is impaired (Guillet & Arndt, 2009). The most comprehensive multivariate account of this effect is again provided by Madan et al. (2017). Following the previously discussed lexical decision task, participants performed a free recall task in which they had to name as many of the words from the previous task as possible in five minutes. As well as replicating the memory advantage for emotional words, Madan et al. showed that, once non-emotional factors (e.g., lexical frequency, imageability etc.) were controlled for, the strongest predictor of successful recall was the specific property of tabooness.

Taken together, these findings from psycholinguistics are consequential for any form of experimental work that includes swearwords. Swearwords take up increased atten-

tional resources compared with neutral words; they have been shown to impede other simultaneous cognitive processes including colour-, voice- and picture-identification. As a result, swearwords are argued to be subject to more automatic processing than neutral words. Swearwords are more effortful to process individually, but emotional words more generally are easier to process when embedded in full sentences. Finally, swearwords are stored longer in memory than neutral words.

These findings have direct consequences for the experiments discussed in this thesis. For tasks that test the processing of swearwords, these findings provide a context in which emergent effects can be interpreted. This is the case in Chapter 4, in which I explore the role of phonetic variation in the perception of swearing. In Experiment I of that chapter, a variant categorization task, instead of showing the expected bias towards alveolar [ɪn], participants showed a tendency towards velar [ɪŋ]. One of the possible interpretations of this result will be that, like in the numerous tasks discussed in this section, participants' attention was taken up by the swearwords and, as a result, their ability to process the critical (ING) tokens was impaired; as a result, they were more likely to select '*-ing*', as this is the form that is most likely overall (as shown in Vaughn & Kendall, 2018).

2.4 Swearing in sociolinguistics

As discussed in Section 2.3.4 on the pragmatic aspects of swearing, while swearwords are perceived by some to be impolite and/or offensive, the reality of how users produce and interpret meaning in context is more complicated than this traditional classification. Relatedly, work in the field of sociolinguistics, which incorporates pragmatics, goes a step beyond studying swearing in context, shifting the focus onto the individuals who use swearwords to position themselves in the social world. While the fields of pragmatics and sociolinguistics are clearly not entirely distinct (see Holmes, 2018 for a discussion), in the context of this section I will be discussing research that links swearing to particular social identities, personae and stereotypes, both in terms of production and perception. That is, work that shows that the use of and reactions to swearing depend on language-external factors.

Work of this type is highly varied and disparate, coming from fields of study including corpus linguistics, discourse analysis, communications, psychology and sociology, as well as some work from sociolinguistics. I will be couching my discussion of this literature in the context of prominent concepts from sociolinguistics including *indexicality* (Silverstein, 2003; Eckert, 2008). In particular, while research on swearing and social evaluation has not historically been done by sociolinguists, I will be discussing this literature in terms of the social meanings that swearwords have been shown to index and the various language-external factors that have been shown to affect social evaluation as

a function of swearing.

2.4.1 Production data

2.4.1.1 Corpus linguistics

Unsurprisingly, just as with all socially meaningful speech forms, not all types of people swear equally as frequently, nor do they use the same specific swearwords as frequently. McEnery (2004) provided the first comprehensive snapshot of British swearing behaviour as a socially stratified phenomenon in his analysis of the Lancaster Corpus of Abuse, a subset of spoken data from the 1994 British National Corpus (BNC) (see McEnery & Xiao, 2004 for a specific analysis of *fuck*). This work was replicated by Love (2017) (see also Love, 2021) using the 2014 edition of the BNC. In this more recent study, a set of the most frequent swearwords from both corpora - namely *arse*, *bastard*, *bitch*, *bloody*, *crap*, *damn*, *dick*, *fuck*, *God*, *hell*, *Jesus* and *shit*, henceforth the 12 BLWs (Bad Language Words), per McEnery (2004) - was analysed for its sociolinguistic distribution by gender, age and socio-economic status; the requirements for this subset was a minimum of 100 occurrences in both corpora. Love (2017) compares the sociolinguistic distribution of these words in the 1994 and 2014 corpora.

Regarding gender, a significant shift appeared to occur in swearword usage. In the 1994 corpus, the 12 BLWs were used more frequently by men than by women (McEnery, 2004). This trend was reversed in the 2014 corpus, with women swearing significantly more frequently overall than men, although the difference is driven by three specific words (Love, 2017). While most of the words are used equally frequently by men and women, both *bitch* and *God* were used more frequently by women, while only *bloody* was used more frequently by men. Additionally, in the 1994 corpus, both males and females were more likely to utter BLWs to people of the same sex as them (McEnery, 2004).

Regarding age, the group shown to swear the most frequently was people between the ages of 15 and 24. This is in line with Cheshire's (1982) claim that swearing has particular value for teenagers. Similarly, the lowest frequency of swearing is found, in both corpora, in the group aged 60 and above, with McEnery (2004) hypothesizing that this could be attributed to euphemistic replacement terms (e.g., *oh dear* rather than *oh fuck*) being more frequent in this age group. One point of difference was the swearing of the 35-44 and 45-59 groups. In the 1994 corpus, while swearing generally lowered in a linear fashion as speakers got older, the was a significant dip at the 25-44 age group, followed by a rise at the 45-59 age group. In their separate study on *fuck*, which also displays this pattern, McEnery and Xiao (2004) suggest that this is indicative of the 35-44 age group being likely to have young children, meaning they would be more likely to avoid swearing. This trend is not apparent in the 2014 corpus, in which swearing among

the 35-44 age group increased, while swearing among the 45-59 age-group decreased.

Regarding socio-economic status, the 1994 corpus shows frequency of the 12 BLWs to be inversely correlated with the height of social class. That is, the highest socio-economic group swears the least, with frequency gradually increasing, such that the lowest socio-economic group swears the most[1]. The pattern for socio-economic group differs in the 2014 corpus, with the highest frequency of the 12 BLWs occurring in the second lowest (C2) group and the lowest frequency occurring in the second highest (C1) group, with the highest (AB) and lowest (DE) groups having broadly similar frequencies. After breaking down the groups using a more fine-grained Social Grade classification based on occupation (which includes 8 ordered categories - see ONS (Office for National Statistics), 2010), however, Love (2017) finds that the highest frequencies occur in the 4-7 range (with 8 being the lowest socio-economic group by occupation), with the absolute lowest frequency at the highest socio-economic level (1.1).

Using the Irish English component of the International Corpus of English (ICE), Schweinberger (2018) finds broadly similar trends to those found in the BNC. Using mixed effect regression modelling, Schweinberger found that speakers aged 19-33 were significantly more likely to swear than those aged 34 and over. Men were found to swear significantly more frequently than women, although only slightly, and, as was found in the 1994 corpus (McEnery, 2004), swearing was more frequent in same-gender conversations. Finally, the authors found that speakers from the Republic of Ireland were more likely to swear than those from Northern Ireland.

Swearing in online contexts has also received significant attention. Thelwall (2008) analyses the use of swearwords on MySpace, with swearing again concentrated among younger speakers. While male and female MySpace-users did not differ in their swearing usage in the UK, they did in the US, where male MySpace profiles included significantly more 'strong' swearing (i.e., words like *fuck* and *cunt* rather than *damn* or *hell*); Thelwall (2008) believes that this is indicative of an underlying difference in gendered expectations in the two countries with respect to swearing.

Using more detailed collocation analysis on data from Twitter, in which both the gender and age of speakers was systematically inferred using information provided in their bios, Gauthier and Guille (2017) observed a number of gendered word-specific trends in English swearing. While male Twitter users again lead in overall swearing, by-word analyses reveal that, while they lead in usage of words like *fuck*, *shit* and *cunt*, female users lead in usage of *bitch*, *bloody* and, among older females, *crap*. There also appears to be a trend in the gender of the target of particular words. The frequent co-occurrence of male and female pronouns with *cunt* and *bitch* respectively suggest that it is men that

[1]See G. Aston & Burnard, 1998 for an explanation of how social class categories are assigned in the British National Corpus.

are typically described as *cunts* and women who are described as *bitches*. Using similar data, Wang et al. (2014) find an effect for social rank, such that swearing is least common in Twitter users with either very small and very large numbers of followers. Wang et al. (2014) also find that swearing increases in both positively and negatively valenced tweets, compared to those that are emotionally neutral.

As the studies reviewed in the section show, while there are certainly observable trends in different social groups with respect to swearing, the picture is again more complicated than simplistic categorization would allow. The results that are found have changed over time (McEnery, 2004; Love, 2017) and vary across different lexical items (Gauthier & Guille, 2017). In the following section, I will review studies that have explored the social function of swearing in interaction for different communities of practice.

2.4.1.2 Swearing in interaction

As already discussed in Section 2.3.4 of this chapter, swearing is not limited to a form of impoliteness. As the literature discussed in this section will show, swearwords can be used for numerous other social functions, both positive and negative. These functions include solidarity, intimacy, trust, humour, toughness and differing forms of both masculinity and femininity. Alongside the studies on perceptions of swearing discussed in the following section, the literature discussed in this section serves to further illustrate how socially complex swearing can be, as well as identifying further language-external factors upon which its social function can depend.

Swearing appears to have specific social functions in the workplace (Daly et al., 2004; Baruch & Jenkins, 2007; McLeod, 2011). Daly et al. (2004) focus on the swearing behaviour of workers in a soap factory in New Zealand. The group of speakers studied were of mixed gender and ethnicity, with the majority of communication in English. As the authors show, the word *fuck* is used in multiple different face-threatening acts, including complaints, refusal and whingeing (or whining). Inherent in all of these acts is the use of *fuck* to mark solidarity between interlocutors. In whinges in particular, in which the speaker is aiming to elicit sympathy from their interlocutor, *fuck* appears to mark familiarity and an appeal to emotion. Even in interactions between workers and their supervisors, speakers can draw on community norms of speaking by frequently uttering *fuck* when trying to refuse an order to do a task (Daly et al., 2004).

Using theories from organisational psychology, Baruch and Jenkins (2007) studied the swearing behaviour of workers in small British retail company. Longer term workers would adjust their swearing behaviour to the audience, typically using more strong swearwords when talking to temporary workers - as a test of their suitability to join the social group - and avoiding swearing altogether in female company. Swearing was again observed to be a marker of solidarity between workers and to enhance group cohesion, as

well as to relieve stress. A similar situation was observed by McLeod (2011) in a trade workplace in Australia. As well as helping build and maintain rapport among the 'tradies', swearing also served to differentiate the group from the rest of society, particularly with respect to their frequent use of *fuck* and *cunt*.

The social function of swearing among young people has also been examined in interaction, perhaps unsurprisingly given the corpus evidence suggesting it's prevalence among that age group (McEnery, 2004; Love, 2017; Schweinberger, 2018). Drummond (2020), for example, combines a small corpus of swearing in English with ethnographic observations in his study of a Pupil Referral Unit (PRU) in Manchester, UK; PRUs are schools for children that have been excluded from mainstream education. Swearing behaviour in the PRU broadly matches findings in McEnery (2004) and Love (2017); the PRU is home to young people from low socio-economic backgrounds who swear relatively frequently. The lack of a significant gender difference among speakers in a lower socio-economic group also mirrors previous findings (McEnery, 2004; S. E. Hughes, 1992).

The more interesting observation from Drummond (2020) is the unmarked nature of swearing among the young people at the PRU. Studies on swearing have frequently focused on its socio-pragmatic functions, such as expressing emotion, humour and social bonding (Stapleton, 2010). Swearing is just a normal part of the everyday language at the PRU however. Words like *fuck* occur with such frequency in individual conversations that, while a degree of social bonding is still being achieved, the effect of each occurrence of *fuck* is lessened due to how saturated everyday talk is with swearwords. Drummond (2020) states that conversations devoid of swearing are arguably more interesting in such a setting. In other work, Auckle (2017) shows that, for teenagers in Mauritius, combining swearing with code-switching indexes an identity rooted in the localised reality of their speech community.

Teenagers, Stenström (2017) suggests, are the strong innovators in the development and evolution of swearing. Comparing swearing in both British English and Spanish using the *Bergen Corpus of London Teenage Language* and the *Corpus Oral de Lenguaje Adolescente de Madride*, Stenström shows that, for teenagers in at least these two languages, swearing is used for rapport-building, bonding and socialising. There are similar gender and socio-economic differences among the teenage corpora to those observed by McEnery (2004) (but not Love, 2017), with swearing frequency highest among male teenagers and those in lower socio-economic groups. While analysing individual extracts, Stenström (2017) observes numerous functions of swearwords, including emphasis, intensification and frustration, none of which seemed intended to cause offence.

Gender has been examined qualitatively as a factor in the production of swearwords. S. E. Hughes (1992) explores the intersection of gender and social class in her study on

English expletives of working-class women in Salford, UK. Despite folklinguistic beliefs about women swearing less than men, as well as a societal expectation for women to refrain from 'coarse language', the women in S. E. Hughes (1992) swear very frequently. S. E. Hughes explains that these women are adhering to localised norms of language behaviour. For them, there is no incentive to use standard forms, as they carry no prestige in their community and using them carries no social advantages. Rather, use of swearing and other local vernacular forms indexes a *toughness* that is integral to social bonding in their group. Swearing is linked to a localised identity that carries more relevance than the broader social category of 'female'.

Stapleton (2003) finds a similar importance of local norms in her study of swearing in English in an Irish social group centred around drinking in a local pub. Despite swearing being more traditionally associated with the construction of masculinity (De Klerk, 1997), the women in Stapleton's study swear frequently. In a survey, some women cited swearing as a tool for "creating/conveying intimacy and trust" (Stapleton, 2003, p. 29). They also rejected the notion of gender-based differentiation in the use of swearing in their community of practice (the pub), while being aware of the broader societal expectation that they should not swear. Gender differences in the group persisted, however; for example, words depicting female genitalia were disprefered among the women in the group. Therefore, while the women did participate in swearing practices, gendered trends did not disappear entirely (see also Methven, 2020 on representations of gender and swearing in the Australian legal system).

Finally, a contested US study by Feldman et al. (2017) suggests that swearing in English is positively associated with honesty. Firstly, using a controlled experiment, Feldman et al. elicited information relating to participants' swearing frequency, their most-used swearwords and their reasons for swearing. Additionally, they completed a *Lie subscale*, a series of questions to measure for individual differences in lying for socially desirable responding (e.g. *If you say you will do something, do you always keep your promise no matter how inconvenient it might be?*). Feldman et al. found a significant effect for use of swearing on honesty, such that participants with higher self-reported usage of swearwords were deemed less likely to lie based on the Lie subscale. In a subsequent study, Feldman et al. (2017) made similar findings using naturally occurring data. The authors compiled a large corpus of Facebook status updates. The status updates were coded for the presence of swearing and the use of linguistic forms most frequently used by liars, namely fewer first- and third-person pronouns, fewer exclusive words (e.g., *but*, *exclude*), more motion verbs (e.g., *arrive*, *go*), and more negative words (e.g., worried, fearful), following Newman, Pennebaker, Berry, and Richards (2003). Use of swearwords was negatively correlated with this measure, such that users who swore were more likely

to be honest[2].

As the studies reviewed in this section show, while swearing may indeed be intended to be rude and/or impolite in certain contexts, there are numerous other social functions associated with the production of swearwords. These social functions often reflect the speakers' speech communities and are used for building relationships between interlocutors. As will be further emphasized in the following section on social evaluation, the social meanings of swearwords are numerous and varied.

2.4.2 Perception data

In addition to the production data discussed in the previous section, numerous studies have examined the role that swearing plays in social evaluation of the speaker. While production data suggests that swearwords are used by speakers to serve a number of social functions, both positive and negative, it does not automatically entail that the intended social function of the swearword will be successfully interpreted. As the studies reviewed in this section will show, the picture is again very complex, with swearwords having been shown to affect how speakers are perceived with respect to a range of different attributes and characteristics, including trustworthiness, intelligence, professional capability and effectiveness, believability, honesty and persuasiveness, as well as the more expected traits of rudeness and impoliteness. Furthermore, the social evaluation of a speaker as a function of their swearing depends on language-external factors including speaker gender, ethnicity and social status, and the formality of the setting. Finally, as with the production data, perceptions of swearing frequently depend on the norms of the particular speech communities in which the speaker is situated, or to which the hearer perceives them to belong. The majority of the studies reviewed in this section involve controlled experiments, similar those used in this thesis, although a minority involve analysis of naturally occurring reactions to swearing online.

As Beers Fägersten (2012) discusses, many previous studies aimed at measuring the perception of swearwords have started from the assumption that the principle evaluative dimension of swearwords is always *offensiveness* (see citations in Beers Fägersten, 2012). Furthermore, such studies have typically utilised word lists in which swearwords are presented devoid of context, with participants left to interpret their potential usage as they wish. This approach goes against conventional wisdom in the field in the present day, which suggests that swearwords can convey vastly different meanings dependent on context. Beers Fägersten adopts a different approach in her study. Rather than using word

[2]The validity of Feldman et al.'s (2017) has been called into question by R. E. de Vries et al. (2018), who question the author's interpretation of responses to the *Lie subscale*. R. E. de Vries et al. suggest, contra Feldman et al. (2017), that high scores on this subscale are actually indicative of high, rather than low, trait honesty in low stakes settings. See also R. de Vries et al. (2018) for continued discussion.

lists, participants were presented with short dialogues containing a swearword; these were taken from the author's ethnographic observations of naturally occurring speech. The findings were broadly heterogeneous, with ratings varying greatly across participants and different dialogues.

A number of trends were observed that would not have emerged from a word list task. Firstly, literal uses of swear words were deemed more offensive than non-literal uses. This finding, in my view, provides further justification for including literal swearwords in the category of *swearing*, contrary to the opinions of Ljung (2010), among others. Secondly, some words differed significantly depending on the derivation that was used. While *motherfucker* has frequently been rated as the most offensive swearword in word list tasks, the use of *motherfucking* in a dialogue received ratings comparable to *fucking*. Gender and race differences also emerged. For example, white males gave higher offensiveness ratings across the dialogues, followed by black males then Hispanic males. Female participants gave higher offensiveness ratings overall than their male counterparts. Beers Fägersten's (2012) results suggest that the activation of this evaluative dimension depends on not only on the specific word that is used, but also on the conversational context and the demographics of the listener.

Numerous other evaluative dimensions are relevant in the social perception of swearing, however. Recent work by DeFrank and Kahlbaugh (2019) explores the link between profanity in English and impression formation. Participants were presented with two short dialogues. The between-subjects variables were speaker genders (including mixed and same-gender) and the presence of swearing by one, both or neither of the speakers. They were then asked to rate the speakers on a number of scales including *intelligence*, *trustworthiness*, *politeness*, *aggressiveness* and *likeability*. The results showed that swearing negatively affected impression ratings on these scales.

These results occurred despite a pre-experiment questionnaire indicating that almost 50% of participants did not find swearing to be 'profane'. This suggests that, even though some speakers may not take offence to swearing, it still negatively affects their evaluation of a person who swears when compared to someone who does not. A further interaction effect was found between profanity and gender, such that, for both females and males, speakers were rated higher in sociability when swearing in same-gender pairings than mixed-gender. The evaluative dimensions relevant to swearing therefore depend on a hearer's pre-existing ideologies about gendered interaction.

Uses of swearwords has also been linked to increased persuasiveness in the US. Scherer and Sagarin (2006) presented participants with a short video of a speech about lowering tuition fees at a university. Using the swearword *damn*, participants were assigned to one of three experimental conditions. One with *damn* at the beginning of the speech, one with *damn* at the end and one with no *damn* at all. Participants rated the

speaker on a number of scales. In both swearing conditions, the speaker was rated as more *persuasive* than in the condition with no swearing. Swearing had no effect on the *credible* scale however.

Increased swearing in Italian has been shown to improve to overall impression of a politician. Cavazza and Guidetti (2014) created a blog post by a fictitious Italian politician, with participants asked questions relating to a number of factors, including persuasiveness, intensity, informality and overall impression; the measure of overall impression was a composite of six different Likert scales depicting the characteristics *sincere*, *reliable*, *dishonest*, *skilled*, *qualified*, and *uninformed*. A 2 (candidate gender) x 2 (swearing vs no swearing) manipulation was used, with all other elements of the blog post kept constant. Participants' overall impressions of the politician improved when the blog post included swearing; this effect was independent of participant gender, education or political orientation. This effect was only present for the male politician, however, with overall impression of the female politician unaffected by the presence of swearwords. On the scale of persuasiveness, the gender of the politician also affected responses, with the male politician rated as less persuasive when swearing, but the female politician's rating again unaffected. Notably, it was only when swearing was used that the male politician was rated as less persuasive than the female politician.

Similar effects were found for swearing in Dutch on the perceived credibility of victims and suspects in a criminal trial (Rassin & Heijden, 2005). Rassin and Heijden (2005) first asked a group of undergraduates whether they thought that the use of swearwords was a sign of deceitfulness, truthfulness or neither. While almost half of the subjects responded with 'neither', 38% thought it was a sign of deceitfulness and only 16% thought it was a sign of truthfulness. In a controlled experiment, however, in which participants were presented either suspect or victim testimony either containing or not containing swearwords, participants rated both forms of testimony as more *credible* if they contained swearwords. This suggests a dichotomy between people's meta-linguistic beliefs about swearing and the reality of how they react to situated uses of swearwords.

In a study of public perceptions of the police in the USA, Patton, Asken, Fremouw, and Bemis (2017) found a correlation between the use of English swearwords and perceived excessive force. Using an experimental design, participants were shown videos of a state trooper arresting a suspect. Four versions of the video were recorded, with male and female actors taking turns in each role. This was repeated twice to create versions with and without swearwords being used by the arresting officer. In all cases, officers were perceived to be more intense, of a lower quality and were seen to be using more excessive force in the swearing condition. Although there were additional effects for officer and suspect gender, these did not significantly interact with the use of swearwords.

Finally, two studies have examined the effect of swearwords in the consumer domain.

Firstly, Sabri and Obermiller (2012) showed that the use of swearwords in advertising produced a more negative brand attitude and reduced the consumer's intention to buy the advertised product. In another study, Hair and Ozcan (2018) examined the use of swearwords in a corpus of Yelp reviews. Each review received ratings from other reviewers, with the option of pressing buttons labelled 'useful', 'funny' and 'cool'. In a multivariate analysis, Hair and Ozcan (2018) found that reviews containing English swearwords were significantly more likely to be perceived as 'useful'; an interaction effect showed this effect to be particularly strong for the most positive reviews (i.e., those with 5 stars). The authors found similar results using a controlled experiment in which participants read either a positive or negative review that either did or did not contain a swearword. The presence of swearing improved the perceived usefulness of positive reviews, but had the opposite effect for negative reviews, suggesting that the valence of an utterance can affect how a swearword embedded in that utterance is perceived.

The studies reviewed in this section further illustrate the complexity of swearwords as a social phenomenon, with a wide range of evaluative dimensions activated by the use of swearwords. In a number of cases, this has been shown to reflect the localised norms of swearing associated with particular social roles, in-line with the production data already discussed in the previous section. As already alluded to in some of these studies, however, much like swearing in production, the social evaluation of swearing can depend on language-external factors.

2.4.2.1 Language-external factors

The gender of the person swearing is clearly a factor in how they are perceived. In DeFrank and Kahlbaugh (2019), for example, swearing was perceived to be a sign of higher *sociability* if both the speaker and the listener were of the same gender. In Cavazza and Guidetti (2014) too, gender affected the results, with the effect of swearing on overall impression stronger for male politicians than female ones.

Similar effects have been found in the domain of sports coaching. In Howell and Giuliano (2011), participants were provided with a fictitious half-time speech by a male basketball coach to their team. Using a similar research design to Cavazza and Guidetti (2014), Howell and Giuliano (2011) varied the presence of multiple expletives including *damn*, *fucking* and *shitty*, as well as the gender of the team (male vs female). Among male participants, use of swearwords led to the coach being perceived more negatively on scales of *effective* and *capable*. This effect was particularly strong in the female team condition, suggesting that swearing was considered less appropriate by male participants when it was directed towards women by a man; Howell and Giuliano (2011) cite this as an example of *benevolent sexism* (Glick & Fiske, 1996) on the part of the male participants. Female participants, in contrast, were not sensitive to the presence of swearwords or the

gender of the team.

A gender disparity also occurred in Martin (1997). Participants were given transcripts of same-gender or mixed-gender conversations in English, with all gender-identifying language (e.g., pronouns, names) removed. They then had to guess the gender of the interlocutors and state what features had led them to this decision. After the conversational topic (e.g., sports vs relationships), the presence or absence of swearing was the second most frequently mentioned cue to the gender of the interlocutors. Swearing was considered a male behavior, while female speech was characterised as containing more "disfluencies" e.g., discourse-pragmatic *like* or *you know*.

Work by Jacobi (2014) suggests that ethnicity can also have an effect on evaluations of utterances containing the swearwords *asshole* and *fuck* and the ethnic slur *nigger*. Using a similar design to DeFrank and Kahlbaugh (2019), Jacobi presented participants with a dialogue depicting an argument between a store clerk and a customer, with the customer's gender and ethnicity (white or black) varied between subjects. While Jacobi focuses solely on perceived 'offensiveness', the complete sentences and the manipulation of social factors lends further weight to the claim that perceptions of swearing are sensitive to utterance context. While, unsurprisingly, white speakers uttering *nigger* were perceived the most offensive, black speakers uttering *fuck* were deemed more offensive than white speakers doing the same.

As already discussed with respect to work by T. Jay and Janschewitz (2008) in Section 2.3.4 of this chapter, perceptions of swearing can depend on the social status of the interlocutor and the formality of the setting. This has also been studied in work environments in the US by Johnson and Lewis (2010) and Johnson (2012). In Johnson and Lewis (2010), participants were asked to consider the extent to which swearing was surprising, shocking, or unexpected in particular situations. Results suggest a greater expectancy violation for swearing in formal settings (e.g., the workplace) compared to social settings (e.g., social gatherings). There was no effect for gender or social status of the speaker; it was no more surprising to hear a peer swear compared to a supervisor. A follow-up study from Johnson (2012) explored the effect of valency violation, showing that the perceived appropriateness of a message depends on how positively a person reacts to that message's content. For example, if a speaker utters "Damn, I'm tired of Wendy's bullying", a listener might react more positively if they agree with the sentiment about Wendy's bullying. In utterances with lower valencies, the same swearwords were deemed less appropriate.

2.4.3 Social meaning

As mentioned in the introduction to Section 2.4 of this chapter, research on the sociolinguistics of swearing has not often been done using sociolinguistic terminology or

sociolinguistic frameworks. With a view to situating the existing research in a unified framework, as well as setting up the framework in which my own contributions will be considered, the current section introduces the concept of the *indexical field* (Eckert, 2008) for the study of swearing. In adopting this approach to the social meaning of swearwords, I am following work by Christie (2013), who combines a sociolinguistic approach to indexicality with a relevance theoretic (Sperber & Wilson, 1986) approach to pragmatics. Towards the end of this section, I will relativise some of the findings discussed in Sections 2.4.1 and 2.4.2 of this chapter to the sociolinguistic framework.

2.4.3.1 The indexical field

The concept of the *indexical field* has become well-used among sociolinguists across the discipline in recent years. Although the linking of specific speech forms to distinct social groups has long been the focus of sociolinguists, this concept provides a theoretical account of how those forms come to be associated with particular social meanings. In her conceptualisation of the indexical field, Eckert (2008) draws heavily on the concepts of *enregistement* (Agha, 2003) and *indexical order* (Silverstein, 2003).

Enregisterment refers to the process through which particular aspects of an accent come to be associated with particular characteristics, which in turn makes that accent available as a signalling device. Agha (2003) discusses the term in relation to the Received Pronunciation (RP) accent in British English and the way in which the accent has maintained a particular level of prestige through the continued meta-discourse seen on television, in newspapers and in other forms of media. This meta-discourse continues to associate RP with particular cultural values such as good education and manners. This meta-discourse is ideological, Agha argues, because those responsible for it have their own particular social identities and goals; it is the activity of "socially locatable persons" (p. 242). Agha (2003) suggests that a form has become enregistered once "certain regularities of evaluative behaviour can be observed" (p. 242); importantly, these behaviours will not be entirely uniform, even within the same speech community. Once aware of the enregisterment of a particular set of forms, speakers can align or dis-align themselves with the social stereotypes associated with them. In this way, an accent such as RP can become a resource for speakers to index the qualities and characteristics that have come to be associated with it; or, they can distance themselves from those qualities and characteristics by diverging from RP forms.

Indexical order is a theoretical concept used to explain how to "relate the micro-social to the macro-social frames of analysis of any sociolinguistic phenomenon" (Silverstein, 2003, p. 193). A macro-social frame can refer to a larger population-level membership to a particular social group such as socio-economic or gendered groups. It can also refer to more locally relevant group memberships, such as being a *Vinyarder* in Labov's

(1963) seminal study of sound change in Martha's Vineyard. The micro-social - or micro-contextual - refers to the way in which forms that are indexical of macro-social frames are assigned additional meanings in interaction via ideological interpretation. In Silverstein's (2003) terminology, but articulated by Johnstone, Andrus, and Danielson (2006, p. 78), *1st-order* correlations between demographic identities and linguistic usages come to be available for *2nd-order* sociolinguistic "marking". In applying this approach to Labov's (1966) study of (R) in New York department stores, Silverstein (2003) suggests that, via 1st-order associations between higher social class groups and the standard variant, the same usages take up the 2nd-order indexical meaning of "prestige" via an association with that macro-social frame.

Building on both Agha (2003) and Silverstein (2003), Eckert (2008) proposes the *indexical field* as the appropriate theoretical device for understanding the dynamic use of linguistic forms for the indexing of social meanings. Eckert (2008, p. 454) characterises the indexical field as "a field of potential meanings" and as "a constellation of ideologically related meanings, any one of which can be activated in the situated use of the variable". The field is "fluid", and new "ideological connections" can be built at any time. Linguistic variables, Eckert suggests, indirectly index demographic categories through their association with particular "qualities and stances that enter into the construction of categories" (p. 455). The concept of the indexical field starts from the point of social meaning and examines how these meanings come to be associated with both linguistic forms and demographic categories.

Eckert (2008) usefully explains how this model might explain the patterning of sociolinguistic variables using data from her own study on white teenagers at a school in Detroit, Michigan (Eckert, 1989). The data suggested a significant difference in the usage of certain non-standard variants between the school's *jocks* (popular sporty types) and *burnouts* (school-alienated types), with the burnouts leading in use of newer non-standard variants that were more common in areas closer to the urban centre of the city. On the one hand, this led to an "embedding [of] a linguistic opposition between city and suburb within a community to support a local opposition between urban- and school-oriented kids" (Eckert, 2008, p. 458). Eckert argues, however, that the burnout students are not making directly *urban* claims through their use of features common in urban areas; that is, they are not consciously trying to associate themselves with the macro-social category of 'urban'. Rather, they are associating themselves with the socially salient qualities of the kids that live in those areas; kids that are tough, autonomous and street-smart. These qualities belong to the indexical field of the urban speech style in Detroit; they are ideologically related to the identity of 'urban' and other social practices in which members of that group engage. The adopting of features of urban speech therefore represents an attempt by the burnouts to align themselves with these qualities, rather than directly sug-

gesting that they are representative of Detroit's urban areas.

The indexical field has also been employed from a sociolinguistic perception perspective to explain how linguistic variables can activate social meanings in the mind of a listener, dependent on their situated usage. Eckert (2008) uses Campbell-Kibler (2007)'s study on perceptions of variable (ING) as an example. Using the matched-guise technique (Lambert, Hodgson, Gardner, & Fillenbaum, 1960), Campbell-Kibler (2007) compared perceptions of the velar and alveolar realisations of (ING) (e.g., *playing* vs *playin*), showing that the alveolar form is associated with lower levels of education, intelligence, and articulateness. Importantly however, the activation of these social meanings depends on the speech style in which the variant is embedded. While the use of the alveolar variant used with a Southern American accent made that speaker sound more *accented*, it dampened the perceived strength of accents that were perceived as *gay* and *urban*. While many social meanings may be part of a form's indexical field, therefore, their activation depends on the situated use of that form.

2.4.3.2 Swearing and the indexical field

Christie (2013) applies the concept of the indexical field to swearing using both corpus linguistic and discourse analytical methods. Using data from McEnery and Xiao's (2004) study of *fuck* in the BNC, Christie (2013) first identifies population-level trends in usage as evidence that *fuck* is a first-order index of particular demographic groups; *fuck* is most common among male speakers and lower socio-economic groups (McEnery & Xiao, 2004). Christie (2013) argues that strong swearwords like *fuck* have the potential to acquire an *n + 1st* order index, through which they are indirectly associated with these demographic groups via other related ideological meanings.

Christie then focuses on reactions in the media to swearing by people in the public eye. She uses these reactions as evidence of a meta-discourse around swearing, suggesting that swearwords can be argued to have second-order indexical values. One such example is the media reaction to footballer Wayne Rooney's use of *fuck*. After scoring a goal in 2011, Rooney grabbed a television camera, stared down the lens and shouted "you fucking beauty!" in celebration; in the aftermath, Rooney was banned for two matches. Reporting on the incident, many newspapers made reference to particular swearing stereotypes that may explain Rooney's behaviour.

One reporter in the Daily Mail suggested that the football field was a workplace and that "industrial language" was part of its currency. The word *industrial* appears to draw a link between swearing and working-class employment. Christie (2013) suggests that, as second-order indexes, the social roles of 'football player' and 'industrial worker' are also in the indexical field of swearing. Numerous other reports claim that the use of swearing on a football pitch is unremarkable, suggesting that the swearing has become enregistered

as typical of footballers; some reporters also suggested that Rooney's swearing was inappropriate in a public place in front of women and children. This contrasts with another of Christie's examples, namely the use of *cunt* by Gwyneth Paltrow. In this instance, it was suggested in the media that Paltrow was attempting to appear as a "bad girl" and appear "edgy and cool". These contrasting reactions suggest that the indexical fields of these strong swearwords contain social meanings pertaining to particular socially locatable personae (e.g., industrial workers and bad girls, but not other women or children) and characteristics (e.g., edgy and cool).

What Christie (2013) provides in her account is a snapshot of the potential social meanings that a swearword might index. As the body of literature reviewed in the previous sections of this chapter suggest, however, swearwords are associated with numerous other evaluative dimensions that can be made more or less relevant in particular contexts. This fits neatly into what we would expect if swearwords have both first- and second-order indexical values. Consider, for example, the finding of Cavazza and Guidetti (2014) that links the use of swearing to an increase in perceived impression, a composite measure of the attributes *sincere*, *reliable*, *dishonest*, *skilled*, *qualified*, and *uninformed*. Notably, this was only the case for the male politician, not the female politician.

We can interpret the results of Cavazza and Guidetti (2014) using Eckert's (2008) conceptualization of the indexical field. The Italian equivalents of the phrases *pissed off* and *up shit creek*[3] could be argued to have indexical fields containing the social meanings included in this composite measure; these meanings could be said to be second-order indexes (Silverstein, 2003). Near to these social meanings would be a male politician persona, but not a female politician persona. In the context of that persona - which is activated in the context of the experiment using a male name - the use of these swearing phrases activates the social meanings of *sincere*, *reliable* etc. For the female politician persona, because these social meanings are not ideologically linked to both the persona and the swearing phrases, the meanings are not activated.

The same interpretation can be made for other findings reviewed in the previous sections. Jacobi (2014)'s finding that a black speaker was rated as more *offensive* than a white speaker using the word *fuck* might suggest that, for this set of participants at least, the indexical field of *fuck* has a stronger link between the social meaning 'offensive' and the social stereotype of the macro-social category of a black person. Martin's (1997) finding that swearing was considered a *male* behaviour suggests that swearing has a first-order index for the category of *men*, possibly via other second-order associations with masculine stereotypes and stereotypical qualities. Finally, Howell and Giuliano (2011) found that male participants rated a male basketball coach as less *effective* and *capable* when swearing, particularly if the team he was coaching consisted of female players. For the

[3]Cavazza and Guidetti (2014) do not provide the phrases in Italian, only these glosses.

male participants, the indexical field of swearing contains both of these social meanings, ideologically linked to the gendered first-order indexical stereotype that men swearing in front of women is less appropriate.

The indexical field is an appropriate way of conceptualising the social meaning of swearwords for the following reasons. Firstly, this approach accounts for the highly varied social functions of swearing, as well as the highly varied social evaluations it can receive. Secondly, it accounts for the fact that the situated use of a swearword, for example in different settings or from the mouths of different people, affects which social meanings are activated; the social meanings in the indexical field are potential meanings, as they may or may not be made salient depending on context. Thirdly, the fact that the indexical field is fluid accounts for changes in the social meanings associated with swearing that occur over time. Finally, the indexical field accounts for the links the have been shown to exist between swearing, demographic category membership and the qualities that are stereotypically thought to hold of those categories' members. Later in Chapter 7, I will present indexical fields for previous findings in swearing research, as well as for the findings in this thesis.

2.5 The knowledge gap

The research reviewed in this chapter so far can be summarised thusly. Swearing is associated with a wide range of different social meanings, the activation of which depends on an array of language-external factors, that is, the *who*, *when* and *where*. Further to that, swearing, while definable to a certain extent as a homogenous set of words with particular characteristics, is very heterogeneous with respect to usage; in simpler terms, not every production of the same swearword is the same. Rather, each production can vary with respect to pronunciation, word and sentence formation, and meaning contribution. In an effort to use consistent terminology, we might define this as the *how*, as in, *how a swearword is produced*.

Currently missing from research on swearing is work that links variation in the *how* of swearing with the social meanings it can index. That is, we do not yet know whether language-internal factors condition the way in which swearing affects social evaluation of the speaker. For example, would a speaker necessarily be evaluated in the same way when uttering the same swearword but with different phonetic realisations, as in (28-a) and (28-b)? Or in different word formations, as in (29-a) and (29-b)? or in sentences in which the swearword appears to make different meaning contributions, as in (30-a) and (30-b), or (30-c) and (30-d)?

(28) a. John is **fucking** married

b. John is **fuckin'** married

(29) a. That's **fucking** fantastic
b. That's fan-**fucking**-tastic

(30) a. The table is **fucking** long
b. The table is **fucking** wooden
c. The **fucking** kids stained the carpet
d. The **fucking** paint stained the carpet

Each of these language-internal factors will be examined in greater detail in Chapters 4, 5 and 6. The cases for expecting each of these language-internal factors to affect the social evaluation are subtly different, so will be dealt with in the individual chapters. The overall motivation for examining the role of language-internal variation in swearing in this process should nonetheless be clear at this point, however. Swearing is highly linguistically and socially variable, as I have demonstrated through extensive reference to the existing literature in this chapter. Whether these two forms of variation interrelate remains to be seen.

Examining this interaction will push swearing research forward significantly. The majority of the conclusions drawn from studies reviewed in Section 2.4 of this chapter refer to the social meaning(s) of 'swearing' (or other related terms such as *taboo words* or *profanity*) as a general concept. Even studies that focus on the use of specific swearwords (e.g., Scherer & Sagarin, 2006 and Jacobi, 2014) draw conclusions that assume homogeneity in usage of that word. Without understanding the role that internal-linguistic variation plays in the social meaning of swearing, claims such as these remain abstract.

Addressing this gap in the literature will also push the larger field of sociolinguistics forward. Sociolinguists in the variationist tradition have long been interested in the role of linguistic constraints on sociolinguistic variation (see e.g., Labov, 1994). Swearwords, as a distinct set of lexical items, present an interesting puzzle, however, because they are inherently socially and pragmatically meaningful in a way that other words are not. Even before another form of potentially socially meaningful variation is taken into consideration, a swearword is already jumping out to the listener as saying something about the person saying it and/or the context they are in. As already discussed above, this can then be interpreted in a number of different ways. Given the evidence that suggests that swearwords are also cognitively dominant, taking up increased attentional resources compared to neutral words, the social meaning of the swearword itself is likely to be highly salient to the listener. Any additional socially meaningful linguistic variation, such as that illustrated in examples (28)-(30), will therefore be competing for attention with whatever social meaning is already activated by the presence of the swearword. As well as furthering our understanding of swearing as a socially meaningful phenomenon therefore, exploring

the interaction between linguistic variation and social meaning in swearing furthers our understanding of sociolinguistic cognition.

Chapter 3

New methods for studying swearing

3.1 Introduction

As discussed in the previous chapter, research on the social evaluation of swearing is very disparate. The findings that were discussed come from a range of different disciplines. As a result, the experimental methods used in these studies do not follow any one tradition. Furthermore, as none of these studies have examined the language-internal factors that may be playing a role in the social evaluation of swearing, there is no established tradition for doing experimental research on this process.

The aim of this chapter is to introduce the experimental approach I will take in order to address this knowledge gap. I will also introduce the individual tasks that will be employed in each chapter, as well as motivating their use for understanding the role of linguistic variation in the social evaluation of swearing. As will become clear, this thesis involves a two-pronged approach. For each variable of interest - phonetic, morpho-phonological and semantic - there will be two experiments. The first will measure what people *perceive* with respect to the variable. That is, what implicit knowledge do people have of each variable? The second will measure what people *notice* about that variable. That is, does their implicit knowledge of a variable affect the social meanings with which they associate it?

These methods are new to the study of swearing in the way they are used in this thesis. Each task has some basis in the literature on swearing, however. There have already been studies that have examined the effect of swearing on the processing of sound (Bertels et al., 2011), in acceptability judgments (Reilly et al., 2020) and on reading speeds (Donahoo, 2019). There have also been previous studies that have employed matched guise-style tasks using swearwords (Howell & Giuliano, 2011; Cavazza & Guidetti, 2014; Patton et al., 2017; DeFrank & Kahlbaugh, 2019). The innovative methodological approach that I employ in this thesis is using multiple tasks to investigate different aspects of the same linguistic variable.

It is also innovative to systematically investigate three levels of linguistic representation using this method, using the same perceiving/noticing distinction at each level. Even in sociolinguistic research more broadly, it is rare for researchers to compare levels of linguistic representation systematically (see Levon & Buchstaller, 2015 for an exception). While some researchers have examined different levels of awareness of a linguistic variable (Breadmore, Krott, & Olson, 2014; Squires, 2016; McGowan & Babel, 2020), these have been focused at one level of linguistic representation. Swearing makes it possible to compare different levels of representation because it is so inherently socially meaningful.

It is hoped that, by introducing insights and methods from experimental sociolinguistics, this thesis will set an example for how social perceptions of swearing should be measured. Furthermore, with respect to the study of social meaning in language more generally, this thesis employs an approach that investigates both the unconscious, grammatical processing of meaning and the conscious, social processing of meaning for the same linguistic phenomena. Employing this approach furthers our understanding of the different levels of awareness involved in sociolinguistic cognition (Squires, 2016).

I will briefly note here, for the sake of clarity, that the order in which these experiments are reported in this thesis does not reflect the order in which they were conducted. In reality, the first experiments conducted were those in Chapter 5, followed by those in Chapter 4 then those in Chapter 6. This is relevant due to the fact that the demographic data collected from participants is not consistent across all six experiments. As my research developed and as I conducted more experiments, I learnt more about the process and about the best ways in which this data can be elicited. As such, I collected the best demographic data in the final experiment of Chapter 6; in this experiment, I also included a proper debrief phase for each participant, in which they reported what they thought the experiment had been about. I provide this information purely for transparency.

3.2 Perceiving vs Noticing in Sociolinguistics

Not all sociolinguistic variables operate at the same level of consciousness. While some variables are the target of much media commentary - see e.g., discussions of the use of slang in schools (Neal-Holder, 2020) - suggesting that the public has a heightened awareness of them, other variables operate below the level of consciousness. In his early work on the sociolinguistic variable, Labov (1972) drew a distinction between sociolinguistic *indicators*, *markers* and *stereotypes*, three abstract categorizations of sociolinguistic variables that exist along a continuum of awareness. *Indicators* are variables that may be socially stratified in a particular speech community in some way, but which have not yet risen above the level of consciousness in that community. For example, a particular group of speakers might use a particular variant more frequently than another, but they won't be

aware of this. *Indicators* do not vary stylistically. That is, speakers do not use more or less of a variant depending on who they are talking to. Speakers tend to show no awareness of the variable, not even subconscious awareness; as a result, usage of the variant is very unlikely to activate social meanings in the mind of the listener. In the continuum of awareness therefore, an *indicator* is at the bottom end i.e., a variable of which speakers have no awareness.

In the middle of the awareness continuum, there are sociolinguistic *markers*. *Markers* are variables of which speakers show some sub-conscious awareness, although they do not have a meta-awareness of the variable. That is, although listeners may react negatively to stigmatized variants, they would be unlikely to say that X type of person uses variable Y. *Markers* can be socially stratified, and they also vary stylistically, with speakers increasing or decreasing their use of a variant in different situations. Finally, at the top end of the awareness continuum are sociolinguistic *stereotypes*. These are sociolinguistic variables that speakers have conscious awareness of, to the extent that they are subject to significant meta-commentary. Speakers frequently and explicitly mention the variables in connection with particular groups of speakers and their speech styles.

Squires (2016) draws on the work of Schmidt (1990) to address what sociolinguists mean when they talk about different types of consciousness. Schmidt (1990) distinguishes between the processes of *perceiving*, *noticing* and *understanding*. *Perceiving*, in Schmidt's terms, is the creation of "internal representations of external events" (p. 132). *Noticing* involves these internal representations affecting conscious processes, but not themselves rising to the level of consciousness. Schmidt (1990, p. 132) summarises the difference thusly:

> *When reading, for example, we are normally aware of (notice) the content of what we are reading, rather than the syntactic peculiarities of the writer's style, the style of type in which the text is set, music playing on a radio in the next room, or background noise outside a window. However, we still perceive these competing stimuli and may pay attention to them if we choose.*

In another linguistic example, Schmidt suggests that one might *notice* that someone has a regional accent without being able to describe it phonetically. In sociolinguistic perception terms, we might conceptualise this as occasions where listeners show differences in how they socially evaluate a voice as a function of particular linguistic features it contains, but without those participants knowing which features are causing those evaluative differences. They might think that a particular voice sounds more or less articulate, intelligent, working-class etc, without realising that this is driven by the presence of e.g., specific sociophonetic variables. Finally, *understanding* something involves noticing

something about our environment and comparing this experience to our other experiences, to the point where we can "reflect on the objects of consciousness and attempt to comprehend their significance" (Schmidt, 1990, p. 132).

Importantly, *perceiving* is a prerequisite to *noticing*. With respect to sociolinguistic perception, in order for a person to *notice* a sociolinguistic variable and assign some form of social meaning to it, they must first have *perceived* that variable, that is, they must have formed some internal representation of that variable. That internal representation can also be referred to as *implicit knowledge* of the variable.

In her chapter, Squires (2016) explores the extent to which people *perceive* and *notice* the same sociolinguistic variable. Using a self-paced reading task (see Section 3.3.3 of this chapter for a summary), Squires (2016) compares the processing of non-standard English syntax, comparing the standard constructions [plural+don't] (1-a) and [singular+doesn't] (1-b) with the common but non-standard [singular+don't] (1-c) and the uncommon and non-standard [plural+doesn't] (1-d) constructions.

(1) a. After eating, the turtles don't walk very fast
b. After eating, the turtle doesn't walk very fast
c. After eating, the turtle don't walk very fast
d. After eating, the turtles doesn't walk very fast

(Squires, 2016)

The results suggest a significant slowdown for both non-standard constructions, from the *don't*/*doesn't* up to and including the final word in the sentence, compared to the standard constructions. Importantly, the uncommon non-standard [plural+doesn't] construction led to the biggest slowdown. This suggests that participants have an awareness that the other non-standard construction, [singular+don't], is a more likely structure. That is, the participants *perceive* the construction as distinct from both the standard constructions and the uncommon construction.

After the self-paced reading task, participants were asked whether they "noticed anything interesting about the grammar of the sentences". Based on their responses, participants were grouped as *aware* and *unaware*. If participants made explicit reference to *don't*, *doesn't* or subject-verb agreement, they were put in the *aware* group. If they did not, they were put in the *unaware* group. This separation was used to identify those participants that *noticed* the non-standard constructions (the *aware* group) and those that did not. These groupings were then used to re-analyse the results of the self-paced reading task.

While there were subtle between-group differences - for the *unaware* group, only the uncommon [plural+doesn't] structure led to a reading slowdown - both groups ap-

peared to perceive at least one agreement difference. Despite all participants showing some evidence of perceiving a difference however, only one group were shown to *notice* the difference. Squires (2016) suggests this to be evidence for the distinction between the two processes. Similar works include Breadmore et al., 2014 on deaf children *perceiving* but not *noticing* similar constructions and (McGowan & Babel, 2020) on native speakers of Bolivian Spanish's *perceptions* and *understandings* of Spanish-dominant and Quechua-dominant speakers.

This distinction between sociolinguistic variables that are *noticed* and those that are just *perceived* is relevant for any study on sociolinguistic perception. As well as mapping well onto Labov's (1972) prominent distinction between sociolinguistic *indicators* and *markers*, it is a useful framework for considering a sociolinguistic variable from the perspective of the listener/reader. That is, it is useful in assessing what a person experiences about a sociolinguistic variable.

This approach sits alongside work in language attitudes that uses a similar distinction to distinguish between implicit and explicit, or overt and covert, attitudes. Such research has typically focused on the degree to which listeners determine a speaker's social identity based on their accent (Giles, 1970). Implicit attitudes are said to be immediate and automatic, while explicit attitudes are said to take more time and depend on more thoughtful processes (Greenwald & Banaji, 1995). Relatedly, a number of language attitudes researchers have espoused the use of the Associative Propositional Evaluation (APE) Model (Gawronski & Bodenhausen, 2006), a dual process model through which "attitudes can be formed through two distinct but potentially interactive modes of mental processing, namely, associative (automatic or immediate) and propositional (thoughtful) processing" (Pantos & Perkins, 2013, p. 4). The separation of different levels of awareness of socially meaningful language variation is therefore a well-studied of language perception.

3.2.1 Perceiving vs Noticing: Swearing variation

In the case of swearing, given the findings from sociolinguistic perception discussed in the previous chapter, it would be logical to say that the majority of people have reached the level of *understanding* when it comes to the words themselves. Swearwords are highly cognitively and socially salient. They are the subject of significant meta-commentary. Swearwords activate particular social meanings and are linked to particular socially-locatable personae and socio-demographic groups in ways that people can quite clearly articulate (Christie, 2013; Stapleton, 2020). In Labovian terms, they have reached the level of a sociolinguistic *stereotype* (Labov, 1972); they clearly operate above the level of consciousness.

The same is arguably true for individual swearwords. People consistently rate par-

ticular swearwords as more offensive/impolite than others. Some share stronger socio-indexical associations with particular groups of speakers (Stapleton, 2003; Baruch & Jenkins, 2007; Jacobi, 2014). The only context in which it could be said that swearwords operate at a level lower than *understanding* is in Drummond's (2020) study of swearing by teenagers in a Pupil Referral Unit. In that particular speech community, swearing has become so frequent that, among the speakers themselves, it may not be something they are so conscious of doing. Rather, it is an unmarked feature of their speech. In the majority of contexts, however, people are very aware of swearing and the social function is plays.

When it comes to intra-word variation, however, it is less clear whether people *understand*, *notice*, or even *perceive* a difference. If people are to *understand* the social meaning of swearing variation - that is, engage in meta-commentary about it - they must first *notice* it. If people are to *notice* the social meaning of swearing variation - that is, have some sub-conscious social reaction to different forms of swearing - they must first *perceive* swearing variation. To *perceive* the variation is to make some internal representation that differentiates two forms of the same word.

With regard to variables of interest in this thesis, repeated in (2)-(4), I have hypothesized that people may assign different social meanings to sentences like (2-a) and (2-b), (3-a) and (3-b), (4-a) and (4-b), and (4-c) and (4-d). While I will enter into greater detail in the individual chapters as to why one might expect these specific variables to have this effect, the overarching motivation is that a) swearing is highly linguistically varied, b) we already know that listeners assign different social meanings to swearing due to language-external factors and c) listeners are sensitive to language-internal factors when assigning social meaning to other forms of linguistic variation.

(2) a. John is **fucking** married
    b. John is **fuckin'** married

(3) a. That's **fucking** fantastic
    b. That's fan-**fucking**-tastic

(4) a. The table is **fucking** long
    b. The table is **fucking** wooden
    c. The **fucking** kids stained the carpet
    d. The **fucking** paint stained the carpet

In light of the methodological discussion in this section, my hypothesis could be reformulated as: *Listeners will **notice** swearing variation in social evaluation*. Given that *perceiving* is a prerequisite to *noticing* (Schmidt, 1990; Squires, 2016), however, it is important to first test a different hypothesis, namely: *Listeners will **perceive** swearing*

variation. If, in each chapter, I find evidence in support of this second hypothesis, this will motivate an experiment to test the original hypothesis.

Using this two-pronged approach provides a clear linear logic to the thesis. In each chapter, I will first use empirical linguistic data to draw a categorical distinction between two linguistic forms. The nature of this data will differ between each chapter, but will ultimately serve the same goal of systematically differentiating the two forms. I will then employ a first experiment to test whether people *perceive* a difference between the two forms. This experiment will involve automatic processes that will bring to the surface people's implicit awareness of the distinction. Should this experiment suggest that people do indeed *perceive* a difference between the two forms, I will then employ a second experiment to test whether people *notice* a difference between the two forms. These logical steps should create a clear experimental narrative from which my conclusions will ultimately be drawn.

It is worth bringing in a brief caveat at this point, however. On the one hand, and as I will show in the following sections, the concepts of *perceiving* and *noticing* variation pertain to similar processes in relation to phonetic, morphological and semantic/pragmatic variation. It must be noted, however, that there are some differences too. The concept was originally employed in a sociolinguistic context by Squires (2016) using syntactic variation. It has more recently been employed by McGowan and Babel (2020) using phonetic variation. It has not previously been employed, to my knowledge, using morphological or semantic/pragmatic variation.

In particular, one might wonder whether the use of the *perceiving*/*noticing* distinction to measure awareness socio-semantic/socio-pragmatic variation is appropriate. Unlike when it was used by Squires (2016), who was comparing awareness of grammatical and ungrammatical constructions, the contrasts I will be testing in Chapter 6 are more subtle. These contrasts reflect a preference for one reading over another, guided by listeners calculating the utility of different structures for communicating particular meanings. When we talk about a listener perceiving a contrast between e.g., *fucking kids* and *fucking paint*, we are asking whether this preference is sufficiently strong to affect processing, in turn suggesting that the two types of constructions are stored in memory as natural classes. While the experiments in Chapter 6 deal with slightly different contrasts, therefore, I would nonetheless argue that they are getting at similar types of knowledge that may or may not rise to the surface in particular types of tasks.

3.3 Measuring what people perceive

This section will detail the first tasks employed in each chapter aimed at measuring what people *perceive* about swearing variation. Each of these tasks involves a relatively auto-

matic process. In the first, participants have to rapidly categorize words into two groups depending on which variant of (ING) they heard. In the second, participants have to judge novel infixed swearing constructions as either acceptable or not acceptable in English. In the third, participants have to read sentences containing swearwords. In each task, it is expected that participants' performance in the task will reflect some sub-conscious knowledge they have of the two forms in question.

In each sub-section of this section, a task is described in more general terms than is included in the methodology sections of later chapters. The focus is instead on what kind of processes are involved in the task and what conclusions might be drawn from the results.

3.3.1 The variant categorisation task

The variant categorization task which I will describe here is an adapted version of the phoneme monitoring task first used by Foss (1969) and later adapted by Frauenfelder and Segui (1989). In a phoneme monitoring task, participants are presented with linguistic stimuli and are asked to identify, by button press, whenever they hear a particular phoneme, often with the additional instruction to only press when that phoneme appears at a particular position in a word (e.g., word-initially). For example, participants might be directed to press the button every time they hear a word that begins with an /m/ phoneme.

In the original version (Foss, 1969), participants were asked to press a button every time they heard a word that started with a /b/ phoneme. The sentences were manipulated such that the target words appeared either early or late in the sentence and following either a high or low frequency and easy or difficult word. The reaction time for each button press was recorded. Reaction times were significantly longer for target words that followed low frequency and hard words, as well as those appearing early in the sentence. Foss (1969, p. 460) took these results as evidence that "the psychological processes involved in the identification of these two types of entities utilize overlapping or interacting mechanisms"; that is, the authors suggest that the processes of phoneme identification and word recognition can affect one another. As such, the subsystems for phonology, lexicon, etc are not "functionally independent".

In another version of the task, named the *generalised phoneme monitoring task* (GPM task) (Frauenfelder & Segui, 1989), participants are required to respond every time they hear a particular phoneme, irrespective of where it appears in a word. Frauenfelder and Segui (1989) combined their task with a semantic priming task. Participants were presented with two consecutive words auditorily. These words were either semantically related (e.g., *doctor* and *nurse*) or unrelated (e.g., *doctor* and *cook*), with participants required to identify a particular phoneme in the word that they heard second. In a first

experiment, using the original version of the phoneme monitoring task (Foss, 1969) in which only word-initial phonemes needed to be responded to, no effect was found on response times for phoneme identification for semantic relatedness. In a second experiment using the GPM task however, participants were faster to detect both word-initial and word-medial target phonemes when the two words were semantically related than when they were not. This latter finding suggests that improved lexical access can facilitate phoneme identification.

In these and numerous other experiments, there is evidence that bottom-up information, that is, linguistic knowledge, affects phoneme monitoring. This has also been shown for the monitoring of sociophonetic variables. Vaughn and Kendall (2018) examine the role of part-of-speech dominance in the perception of variable (ING). Variable (ING) concerns the alternation between the velar [ɪŋ] and alveolar [ɪn] realisations of the '*-ing*' suffix. The velar realisation is more common in nouns and the alveolar realisation is more common in verbs (Tagliamonte, 2004; Kendall, 2013). Vaughn and Kendall (2018) provide evidence that listeners are aware of the probabilistic distribution of each variant in perception.

Participants listened to sentences containing (ING) words, pronounced with one of the two variants. Sentences contained (ING) words on a part-of-speech continuum, from the most verb-like forms (i.e., verbs) to the most noun-like forms (i.e., nouns), with intermediate forms in the middle (e.g., gerunds). Participants were required to press a button labelled '*-ing*' if they heard a word ending in '*-ing*' and a button labelled '*-in*' if they heard a word ending in '*-in*'. Participants were significantly faster to respond to words with the alveolar [ɪn] variant if they had a smaller bias towards the velar [ɪŋ] variant in production (e.g. nouns and pronoun-3 e.g., *anything* or *everything*), while also being least accurate in correctly identifying the alveolar [ɪn] variant on items with the smallest velar [ɪŋ] bias (the pronoun-3 items). Vaughn and Kendall's (2018) results suggest that listeners make use of bottom-up information from the rest of the word when identifying a sociolinguistic variable.

Elsewhere in sociolinguistics, researchers have used similar tasks to measure the effect of top-down information, that is, cultural or contextual knowledge, when identifying sounds. Both Hay et al. (2006) and D'Onofrio (2018) examine the use of top-down, social information on the categorization of phonemes. Hay et al. (2006) test this premise on New Zealand speech. Many speakers of English in New Zealand have merged SQUARE and NURSE vowels, such that minimal pairs of words containing these two vowels (e.g., *where* and *were*) sound very similar. Notably, the merger is more complete in younger speakers. In the experiment, participants heard one of the words from a minimal pair, pronounced as the un-merged version, and were asked to identify which of the two they heard. For example, on one trial they might have heard *hair*, before choosing whether

they heard *hair* or *hear*. The manipulation of top-down information involved presenting participants with different photos of the 'speakers', varied for gender, age and clothing style. With respect to age, for example, participants were more accurate in the task when presented with a photo of an older person, suggesting that they attended more to the distinction between the two vowels; for photos of younger people, they treated the words as more ambiguous. These results, among others, suggest that listeners make use of top-down information in the categorical perception of sounds.

A similar approach was taken by D'Onofrio (2018) for the TRAP-LOT vowel shift common in California, USA. Production evidence suggests that tokens of the TRAP vowel are becoming backed and lowered, such that they have begun to approximate tokens of the LOT vowel. This shift is most associated with the California-based personae known as the 'Valley Girl', as well as the less region-specific persona known as the 'Business Professional'. In a phoneme categorization task, participants heard minimal pair words containing vowels from a 9-step phonetic continuum from TRAP to LOT (e.g., bat-bot, sack-sock etc). Participants were given one of four possible social primes containing information about the speaker. They were either told that the speaker was described as a Valley Girl, a Business Professional or a Chicago Bears fan, or they were told nothing (the baseline condition). For participants in the Valley Girl and Business Professional conditions, the perceptual boundary between the two vowels was further towards the LOT end of the continuum compared to the baseline condition, suggesting that the social prime caused them to categorize more LOT-like tokens as TRAP.

None of the studies reviewed here provide an exact blueprint for the variant categorization task which I employ in this thesis; a full description of this is available in the relevant methods section of Chapter 4. Rather, these studies point to a more general tradition in speech perception of using tasks that involve the rapid categorization of sounds to reveal some kind of implicit knowledge. These tasks involve the manipulation of bottom-up linguistic information, such as the stimuli's lexical characteristics, and/or top-down social information, such as information pertaining to the speaker. By using rapid categorization, rather than allowing participants to take their time to consider their response, these tasks elicit automatic responses which might be more sensitive to the experimental manipulations. It is through this automaticity that such tasks can measure what a person *perceives*, rather than what they notice.

3.3.2 The acceptability judgment task

An acceptability judgment task involves presenting participants with linguistic stimuli and asking them to examine the extent to which each stimulus seems "good" or "bad" to them. The stimulus in question could be a sentence or an individual word. The experimental

design involves controlled manipulations of particular structural aspects of the stimulus, with other parts of the stimulus kept constant in order is isolate the influence of these manipulations on the perceived "acceptability" of a stimulus. Acceptability judgment tasks are used in morphology, syntax, semantics, phonology and sociolinguistics. The type of task can vary. Some involve a binary Yes-No response, while other versions involve a gradient response, such as a Likert scale.

When we ask participants to give acceptability judgments, it is important to understand what the nature of a judgment is, as well as what conclusions we can draw about language as a result of these judgments. In considering this, I will make repeated reference to points made in Schütze and Sprouse (2014). Firstly, "*acceptability* is a percept that arises (spontaneously) in response to linguistic stimuli that closely resemble sentences (i.e., strings of words)" (Schütze & Sprouse, 2014, p. 28). Just like other percepts (e.g., brightness or pain), there are only indirect measurement tools for acceptability. An acceptability judgment is therefore better thought of as a *reported perception of acceptability*. Reported perceptions of acceptability, although requiring a degree of metalinguistic awareness on behalf of the participant, are typically systematic in ways that have allowed linguists to build grammatical theories that make falsifiable predictions about language. Acceptability judgments are useful because they allow us to learn about phenomena that rarely or even never occur in spontaneous language. Just because a particular construction has not been attested in large corpora does not mean that it is ungrammatical, nor does the existence of a construction in such corpora automatically mean it is grammatical (Schütze, 2009). What acceptability judgments are not is a direct measure of grammaticality. Grammar is an abstract mental concept. It is not accessible conscious awareness and is therefore not something on which a person is able to directly reflect.

One study reviewed in the previous chapter used acceptability judgment tasks with swearwords to measure the influence of particular linguistic characteristics on the acceptability of novel swearing constructions. Reilly et al. (2020) used an acceptability judgment task to show that swearing compounds in which the non-taboo part of the word contains obstruents are considered more acceptable to native speakers of English than those without. Similarly, although not dealing with "acceptability" directly, Tessier and Becker (2018) asked participants about how "satisfying' a series of novel swearing compounds were, finding that compounds with matching vowels (e.g., *fuck-puffin*) were deemed more satisfying than those without (e.g., *fuck-badger*). In Reilly et al. (2020), and to a lesser extent Tessier and Becker (2018), it is very unlikely that the participants had previously heard all of the test constructions they deemed *acceptable* (or *satisfying*). Rather, their responses were probabilistic judgments based on those constructions they had previously heard and the linguistic properties those constructions had.

In this way, using an acceptability judgment task with novel swearing constructions

tells us what people have *perceived* about swearing. In making a judgment about the acceptability of a construction, we are asking whether people have made an internal representation of the grammar of swearing based on their previous experience of hearing and using swearwords. It is unlikely that they know the linguistic reason why one construction is acceptable to them and another is not. Rather, this is just part of their *implicit knowledge* about swearing. The task brings to the surface this implicit knowledge of swearing and the grammatical constraints placed upon it, as well as the non-grammatical properties that they think swearing constructions typically possess (e.g., low valency, high arousal); extra-grammatical factors often influence acceptability effects (Sprouse, 2018). This task will be used in Chapter 5 to test which factors affect the perceived acceptability of infixed swearing constructions.

3.3.3 The self-paced reading task

A self-paced reading task involves participants reading whole sentences on a computer screen. The sentences are presented incrementally, typically one word at a time. The participants control the speed at which the sentence is presented. Every time they press a button, the next word is presented on the screen. The participants are directed to read the sentence at a regular speed. The time between each button press is recorded and is taken to be indicative of how long the participant spent reading and processing the newest word on screen; this *reading time* is the dependent variable. The independent variables are typically manipulations in the sentence to test the processing speed associated with e.g., particular grammatical constructions. Increases in reading time at particular points in the sentence are taken to be indicative of a processing cost incurred by the participant.

Self-paced reading tasks were first used in the 1970s (Aaronson & Scarborough, 1976; Mitchell & Green, 1978) with the aim of finding a real-time method for measuring language comprehension processes in a way that was "as similar as possible to normal reading" (Mitchell & Green, 1978, p. 610). Since then, they have been employed to study a range of different linguistic phenomena in the fields of syntax, semantics, pragmatics and second language research. The majority of studies now use the *moving window* technique, whereby only the newest word is visible on the screen at any one time; this is due to issues with participants sometimes using reading strategies that involve pressing the button repeatedly, then reading the whole sentence at once, rendering the measure of reading time pointless (see Jegerski, 2014 for a review of the method).

For the purpose of illustrating what a self-paced reading task tells us about what a person *perceives* about linguistic stimuli, as well as to focus this section on the use of self-paced reading tasks to examine semantic/pragmatic phenomena, I will review several studies relevant to my own experiment; this will be described in greater detail in Chapter

6.

Although the majority of research using self-paced reading tasks has focused on syntactic variability, semantic variation has also received attention. For example, Schwarz (2007) showed that reading times are sensitive to the presence of unsatisfied presuppositions. In two self-paced reading experiments, Schwarz examined the processing of the English additive particle *also*, and its German equivalent *auch*, and the presuppositions that they introduce. For example, in the German experiment, Schwarz compared sentences like those in (5). While the presupposition introduced by *auch* - that someone else saw the woman - is satisfied in (5-a) by the relative clause, this not the case in (5-b), in which the relative clause says nothing about another individual having seen the woman. Notably, this was not a by-word self-paced reading task. Rather, each press of the button revealed another chunk of words, as indicated by the '/' separator in the examples below; participants read the subject of the main clause, followed by the relative clause, then the remainder of the main clause.

(5) a. Die Frau,/ die der Junge sah,/ hatte auch der Mann gesehen
The woman/ who the boy saw /had also the man seen
'The woman that the boy saw had also been seen by the man'

b. Die Frau,/ die den Jungen sah,/ hatte auch der Mann gesehen
The woman/ who the boy saw /had also the man seen
'The woman that saw the boy had also been seen by the man.'

The results suggest that reading times were heavily influenced by this distinction. The reading times for the final portion of the sentence were significantly longer for sentences like (5-b) in which the presupposition was not satisfied, with a mean difference between the two conditions of around 1.5 seconds. The results suggest that there was a processing cost associated with trying to integrate information in the main clause, in the form of a presupposition, with information in the relative clause if the two did not match (see also Tiemann et al., 2011 and Schwarz, 2015).

In another study, Dörre, Czypionka, Trotzke, and Bayer (2018) tested the influence of German modal particles (MPs) such as *doch*, *wohl*, *nur* and *bloss* on sentence processing (see also J. Bayer, 1991). Each MP can have either an at-issue or not-at-issue reading depending on context. In (6-a), *bloss* has the at-issue meaning of *only*. In this at-issue context, *bloss* only applies to *the corridor*, restricting the elements of a set (i.e., the rooms that have been wiped) and changing the truth conditions of the sentence. In contrast, in (6-b), *bloss* expresses an emotion about the fact that the corridor is still dirty, contrary to the speaker's expectations. In this not-at-issue context, *bloss* applies to the whole sentence without changing its at-issue meaning component.

(6) a. Wer hat bloss den Flur gewischt?
Who has only/PART the corridor wiped?
'Who wiped only/PART the corridor? (The other rooms are also dirty!).'

b. Wer hat bloss den Flur gewischt?
Who has only/PART the corridor wiped?
'Who wiped only/PART the corridor? (Here is still mud from outside!).'

Using a self-paced reading task, Dörre et al. (2018) presented participants with sentences containing ten different modal particles with dual meanings. These sentences were either preceded or followed by another sentence that provided a disambiguating context to prompt either the at-issue or not-at-issue reading of the MP. In trials with context provided before the test sentence, a slowdown was observed for the two words immediately following the MP when the not-at-issue reading was triggered compared to the at-issue reading. The implication is that, when the not-at-issue reading was triggered, participants had to pay greater attention to the rest of the sentence to understand the meaning of the MP, thus incurring a processing cost. In contrast, for the at-issue readings, the participants only had to consider the phrase immediately after the MP (e.g., *den Flur* in (6-a)).

A similar result was found by Donahoo (2019). Using a Maze Task, Donahoo compared reading speeds for swearing and neutral adjectives modifying nouns. For example, the reading speed of sentences such as those in (7), in which a noun like car was preceded either by a lexical, at-issue modifier like *old* or by a not-at-issue expressive modifier like *damn*. Given that expressives have a flexibility of interpretation (see Section 2.3.3.1 of Chapter 2), the *damn* in (7-b) can operate over either *car* or the whole sentence 'the car broke down yesterday'; the modifier *old*, in contrast, can only operate over *car*.

(7) a. The old car broke down yesterday

b. The damn car broke down yesterday

While reading speeds for swearing adjectives were faster than for neutral words, reading speeds for the modified noun were slower when preceded by a swearing adjective than when preceded by a neutral adjective. That is, the *car* in *damn car* was read slower than the *car* in *old car*[1]. The result is a slowdown caused by a word with a not-at-issue meaning (*damn*) compared with a similarly frequent word with an at-issue meaning (*old*). Donahoo suggests that this is potentially caused by the non-localness of expressives. While an at-issue adjective like *old* only composes with the *car*, *damn* can operate at the sentence-level (cf. Frazier et al., 2015), requiring the comprehender to spend more time reading to find out what *damn* is operating over.

The studies reviewed in this section show only a glimpse of the ways in which self-paced reading tasks have been used to further our understanding of language processing.

[1]Donahoo (2019) does not provide data on reading speeds for the remainder of the sentence.

What these studies do demonstrate, however, is that self-paced reading tasks are sensitive to kinds of semantic variation that I am interested in regarding swearing. The work by Squires (2016) discussed in Section 3.2 of this chapter also shows the utility of using the task on socially meaningful grammatical variation. Furthermore, these studies show the way in which a self-paced reading task examines what a person *perceives* in language. Reading is a fairly automatic process. By introducing linguistic variability into the task and measuring how reading comprehension is disrupted, we bringing to the surface the *implicit knowledge* that people have about language and, in my case, swearing.

3.4 Measuring what people notice

This section will detail the second tasks employed in each chapter aimed at measuring what people *notice* about swearing variation in relation to social meaning. The same task is employed in each chapter, namely a matched guise task. In Chapter 4, this will take place in the auditory domain, while the tasks in Chapters 5 and 6 will take place in the visual domain. Unlike the tasks described in the previous section, these tasks involve a controlled process. Participants are asked to consider a stimulus and make a conscious decision about it. If the first task has ascertained that people have some internal representation of the distinction between the two forms of interest - that is, they have *perceived* a difference - this task will test whether this internal representation affects the process of social evaluating speakers using those forms.

3.4.1 The matched guise task

The matched guise task (MGT) is probably the most widely used method in sociolinguistic perception. The technique aims to uncover the covert attitudes that listeners hold about particular speech styles. The MGT involves recording one speaker reading a set passage out loud. Frequently, one or more linguistic features of interest are then manipulated, either by the speaker or in post-hoc artificial manipulation, to create multiple 'guises'. The guises are then presented to listeners as if they are different speakers. Listeners are asked to evaluate the 'speakers' along relevant semantic differential scales (e.g. intelligent-unintelligent, kind-unkind etc). By using the same speaker for all guises, only the feature of interest is manipulated, with all other aspects of the speech signal kept constant across guises. Differences in how particular guises are evaluated can therefore be solely attributed to the presence or absence of that feature.

In the original version of the MGT, there was no manipulation of a feature, however. Rather, the technique was used to uncover covert attitudes towards different languages. Lambert et al. (1960) recorded Canadian bilingual speakers of English and French read-

ing the same passage in both languages to create two guises per speaker, with each guise presented to participants as a different 'speaker'. Participants rated each 'speaker' on a number of Likert scales depicting characteristics related to desirability in regard to friendship. Using a mix of English- and French-speaking listeners, listeners rated the 'speakers' more favourably on the Likert scales if they were speaking their own native language, despite the speaker being the same person (unbeknownst to the listeners). The results were taken by Lambert et al. (1960) to suggest that, in addition to the information we glean from a person when making judgments about them, some of those judgments are driven exclusively by the pre-existing attitudes we have towards particular languages and the people that speak them.

The MGT has remained popular since its inception, with numerous sociolinguists employing it to explore perceptual dialectology (see Preston, 1999). The technique is also of interest to variationist sociolinguists, given their focus on the linguistic variable as a source of social meaning. Particularly notable is the work of Kathryn Campbell-Kibler (2005, 2007, 2010a, 2011) on the sociophonetic variable (ING). In her original PhD thesis, Campbell-Kibler (2005) artificially manipulated recordings of 8 American English speakers to vary the realisation of variable (ING). This was done using re-recorded versions of spontaneous speech containing tokens of variable (ING) provided by the speakers in sociolinguistic interviews. Recordings were made using both variants, with the tokens cross-spliced into one of the recordings; this ensured that all other aspects of the speech signal were kept constant. This resulted in guises of each speaker in which (ING) tokens were realised either as [ɪŋ] or [ɪn].

A survey containing the guises was then completed by 124 participants in which the guises - or 'speakers' - were rated on Likert scales. The characteristics depicted by these scales were chosen based on focus groups who were exposed to the same stimuli. According to Campbell-Kibler's (2005) results, variable (ING) is linked to a network of social meanings in the USA that includes education, articulateness, formality, region and the rural/urban divide. For example, the guises containing the velar [ɪŋ] variant were consistently rated as being more *educated* and *articulate*. Using the MGT, the experimental manipulation - the alternation between [ɪŋ] or [ɪn] - can be solely attributed with causing these changes in social evaluation. This is beneficial compared to using naturally occurring speech, in which a range of different factors could also be playing a role.

A number of follow-up studies by Campbell-Kibler have shown innovations in the MGT. Campbell-Kibler (2007) tested how perceptions of variable (ING) vary when the variable is used by speakers with different accents; the two accents were a Southern US accent and a stereotypically *gay* accent. For the Southern accented guises, the presence of alveolar [ɪn] increased how 'accented' they sounded to speakers. The reverse was true for the *gay* guise. In another study, Campbell-Kibler (2010a) explored the effect of speaker

information on attitudes towards (ING). Using a similar MGT, guises were presented to participants with additional pieces of information about the speaker. The same set of speakers were either presented as professors, politicians or working professionals, with participants rating the guises on a range of Likert scales. An interesting trend emerged regarding the *knowledgeable* scale (among other findings). Speakers presented as professors were rated as sounding more knowledgeable when using alveolar [ɪŋ]. The opposite was true for speakers presented as professionals, with alveolar [ɪn] guises perceived as more knowledgeable.

Finally, Campbell-Kibler (2011) tested perceptions of variable (ING) in the context of other socially meaningful linguistic variables. Again using an MGT, Campbell-Kibler manipulated guises to include various combinations of the sociophonetic variables /s/-fronting or backing and mean pitch, as well as variable (ING). While all variables carry their own social meanings in isolation, the combination of /s/-fronting, a variable associated with stereotypically *gay* voices, and variable (ING) produced unique combinations of social meanings. For example, while ratings of *competent* and *masculine* typically co-varied (i.e. more masculine = more competent), this correlation disappears for guises with velar [ɪŋ] and fronted /s/. These guises are rated as very *gay* and *competent*, but rated very low for *masculine*. In each of these studies, Campbell-Kibler (2007, 2010a, 2011) combines the original experimental manipulation from her 2005 study with additional forms of bottom-up or top-down information, showing that the variable of interest can be studied in a variety of ways using the same experimental technique, namely the MGT.

The list of studies that have employed the MGT is long; many of them have employed a similar approach to that employed by Campbell-Kibler (2005). Given the range of variables investigated in this thesis - one phonetic, one morpho-phonological and one semantic - several other innovations of the MGT first need reviewing before this section is concluded. While the majority of MGTs have been used in the auditory domain, a number of studies have shown that the visual domain is well-suited for studying linguistic variation in other areas of language such as discourse-pragmatic, syntactic and semantic variation.

Using an MGT with two different written transcripts of naturally occurring speech, Buchstaller (2006) explored the social meanings attached to the quotatives *be like* and *go*, as in *Mary was like/went "why are you late?"*. Three versions of each transcript were created, one each with *be like*, *go* or the control quotative *say*. Participants were randomly assigned one version of each transcript. Participants were asked to estimate, for each transcript, the age, gender and social class of the speaker in question, as well as assessing them for a number of personality traits. Responses in the MGT, which assessed covert attitudes, were then compared with the result of a post-task questionnaire that addressed participants' overt attitudes towards the two quotatives of interest. Some questions ex-

plicitly asked participants if they associated either *be like* or *go* with particular categories e.g. genders, ages, social classes etc. Despite questionnaire results suggesting that the quotative *go* is associated with younger and more working-class speakers, results in the MGT showed no significant difference in attitudes between the *go* condition and the *say* condition. Transcripts featuring *be like* were significantly associated with younger people's speech compared to the control; this result mirrored the overt attitudes expressed with regards to age in the questionnaire. A consequence for approaches to social meaning is that overt and covert attitudes can diverge.

In another study, Levon and Buchstaller (2015) tested the social salience and combined social meanings of two variables from different domains of language. The variables were the phonetic variable TH-fronting (8) and the syntactic variable known as the Northern Subject Rule (NSR) (9), a morpho-syntactic variable common in parts of the North of England.

(8) I **f**ink [think] we should go home now.

(9) They really like**s** ice-cream.

Employing an MGT, Levon and Buchstaller tested how the two variables would affect the perceived *professionalism* of a newscaster. Test conditions included versions of the same passage with 0 or 100% TH-fronting, 0 or 100% NSR and combinations each. The results suggest that the two variables have differently stratified social meanings. On the one hand, TH-fronting led to a decrease in perceived professionalism across participants, irrespective of other factors. While NSR did not come out as a significant predictor of professionalism overall, Levon and Buchstaller did find significant interaction effects to suggest that NSR is socially salient to particular sub-populations, namely listeners from the North of England.

In her perceptual study of copula absence in AAVE, Bender (2005) used an MGT to compare perceptions of the same sociolinguistic variable in different syntactic environments. Among AAVE speakers, sentences are often produced with no (overt) copula. For example, a sentence like (10-a) can be produced as (10-b). Similarly, a sentence like (10-c) can be produced as (10-d). Work on production data has previously identified an important non-categorical constraint on copula absence however, showing that it is significantly more frequently used in a pre-verbal position (10-b) than in a pre-nominal position (10-d); its use pre-nominally is therefore the more marked usage.

(10) a. She's teaching me piano
b. She teaching me piano
c. She is my piano teacher
d. She my piano teacher

Using an MGT, Bender (2005) presented participants with sentences with and without an overt copula, either pre-verbally or pre-nominally, spoken by AAVE speakers. Participants rated speakers on Likert scales typically used in studies of stigmatised varieties. Among all groups, the use of a zero copula in both positions resulted in speakers being rated as less *confident*, *educated*, *polite*, *likeable* etc. Additionally, among only the African American participants, the constraint of syntactic position also affected ratings, with this group rating speakers more favourably on these scales when they used pre-verbal copula absence compared to when they used pre-nominal absence. This study serves of an example of the same form being *noticed* in different ways in different linguistic environments.

Finally, Beltrama and Staum Casasanto (2017) employed an MGT to test whether speaker evaluations track semantic variation in intensification. They focused on the intensifier *totally*, which can have both lexical and pragmatic functions. When paired with a bounded adjective, as in (11-a), *totally* brings the lexical meaning of that adjective to the endpoint of an ordered scale i.e. the most bald it is possible to be. When paired with an unbounded adjective, as in (11-b), *totally* cannot do this, as there is no endpoint on the scale of tallness. Rather, *totally* signals the speaker's commitment toward the content of the utterance. As well as these semantic/pragmatic effects, *totally* seems to have socio-indexical associations with particular personae, namely American Valley Girls and cheerleaders.

(11) a. John is **totally** bald
b. John is **totally** tall

Participants each read 12 sentences on a screen, varied according to the type of intensifier and the gradeability properties of the following adjective, and rated them on a number of Likert scales relating to solidarity and status, as well as an estimation of the speaker's age and gender. Results indicated that pragmatic speaker-orientated *totally* is associated positively with solidarity (e.g. friendly and cool) and negatively with status (e.g. intelligence and maturity); the construction also caused participants to evaluate the speaker as lower in age. Beltrama and Staum Casasanto (2017) demonstrate how *stance* can play a role in the formation of social meanings. Speaker-oriented *totally* signals the speaker's epistemic stance towards a proposition, signalling their increased desire to add that proposition to the Common Ground. This particular stance has, over time, taken on socio-indexical meanings linked to particular stereotypes in the manner suggested by Kiesling (2009). Beltrama and Staum Casasanto's results (see also Hunt & Acton, 2022) justify the measurement of social meaning associated with semantic variation using an MGT.

An important element of the matched guise task is choosing which Likert scales to use. In some cases, a precedent has been set by other studies that suggest particular social

meanings are frequently activated by a particular linguistic variable. Variable (ING), for example, has been shown to activate social meanings related to perceived intelligence, articulateness and social class in both the USA (Campbell-Kibler, 2005) and the UK (Schleef et al., 2017); future research testing other factors affecting the social meanings of variable (ING) would therefore be well-placed in using these same characteristics on its Likert scales. In other cases, particular evaluative dimensions are made relevant by the experimental hypotheses and the type of passage used in the task. In both Labov et al. (2011) and Levon and Fox (2014) the passage was of a newcaster. As a result, the evaluative dimension of *professionalness* was of most interest.

In cases where there is less clarity about which social meanings are relevant, or in which the variables have not been tested previously, steps are sometimes taken to make an objective decision about which scales to include. After all, these scales need to reflect the adjectives used by the particular speech community of interest (Al-Hindawe, 1996). This was done by both Clark and Schleef (2010) and Beltrama and Staum Casasanto (2017). In their matched-guise study of social evaluations of linguistic variation in English by Polish-born adolescents, Clark and Schleef (2010) first played their guises to a group of undergraduate students from Edinburgh and London. Each student was asked to describe the guises in their own words. The ten most frequently mentioned words were used to create the Likert scales. In Beltrama and Staum Casasanto's (2017) matched-guise study of *totally*, a preliminary study was conducted on Qualtrics to identify relevant characteristics. A group of crowd-sourced participants saw one of the variants of interest - a use of either lexical or pragmatic *totally* - and provided adjectives they would use to describe the person using this sentence. The most frequently occurring adjectives were used in the eventual MGT.

Another important decision to be taken when employing an MGT is what type of stimuli to use. In much of Kathryn Campbell-Kibler's work, for example, naturally occurring conversational speech has been used (see also Schleef et al., 2017), although this is often re-recorded for the purpose of performing artificial manipulations on the stimuli. In other cases, such as Labov et al. (2011) and Levon and Fox (2014), read speech is used. That is, pre-selected scripts are recorded by voice actors. Although this second option may seem less natural, and therefore less realistic for participants, work by Tamminga (2017) on variable (ING) suggests that both approaches have similar results. While one might expect that the informal alveolar [ɪn] variant would sound less natural in read speech than in conversational speech, this turned out not to be the case. Tamminga used both types of stimuli, with the same Likert scale effects emerging for each. That is, she got a null result when including the stimuli type as a predictor in her statistical model.

The studies reviewed in this section demonstrate the versatility of the matched guise task for examining the ways in which linguistic variables from different domains of lan-

guage can activate social meanings in the mind of the listener/reader. It is an effective method, in both the auditory and visual domains, for controlling for other possible relevant factors in social evaluation. Furthermore, it is a suitable follow-up experiment to all three of the tasks discussed in Section 3.3 of this chapter. In each case, if the first experiment suggests that listeners *perceive* a difference based on my experimental manipulations - i.e., a difference between *fucking* and *fuckin*, between *fan-fucking-tastic* and *fanta-fucking-stic* and between *fucking tall* and *fucking wooden* and/or *fucking kids* and *fucking paint* - this will suggest that they have some subconscious knowledge these distinctions. If that turns out to be the case, it may also be the case that a matched guise task can reveal whether or not they *notice* that difference in the process of social evaluation.

3.5 Participant Recruitment

The decision was made to conduct all participant recruitment for my experiments online. Conducting experiments online has many advantages over in person experiments. Online experiments are cheaper per participant and allow for the collection of large samples of participants in a shorter period of time. As a result, I was able to collect significantly larger sample sizes for my experiments than is typical for these particular methods when collected in person. The use of online experiments also became a necessity once the UK (and the majority of the rest of the world) when into lockdown during the COVID-19 pandemic, restricting access to human participants and lab equipment.

It is worth noting, however, that the nature of participants taking part in online experiments on platforms like Prolific Academic is not representative of the population at large. The participant pool is majority female and typically aged between 20 and 40. Furthermore, the participant pool is restricted to those people with access to either a smartphone, a tablet or a computer, ruling out a portion of the population. This is worth noting, as it means there are limitations on what I can claim about how 'people' perceive or notice swearwords. It is unlikely that many of my participants would be from older age groups, for example. Similarly, I am unable to collect data from children. In order to make broader claims about the perception of swearing as a general phenomenon, therefore, research with these populations would need to be done.

For each experiment, the same piloting process was conducted. Each experiment was first informally conducted using other linguistics PhD students and non-linguists from my own social network. Each experiment was then launched on Prolific with a much smaller participant pool, typically under 20 participants, to check for possible bugs and to test the average amount of time required to complete each experiment. Participants that had issues with the experiment were able to contact me anonymously via Prolific Academic's private messaging service. Following the requisite adjustments, the experiments were

then launched on Prolific Academic with the aim of collecting the final intended number of participants.

3.6 Statistical analysis

I will briefly describe here the modelling strategy employed throughout the thesis. Taking a frequentist approach, I employ linear and logistic mixed effect regression modelling. For each experiment, I start with a full model that includes all relevant fixed effects and random effects, including random slopes. For model comparison, I start with the full model and manually drop variables one-by-one. After a variable has been dropped, a new model is run. This model is then compared for model fit to the full model in a chi-sqaure test using the *anova()* function in R. If the new model does not significantly differ from the full model, then the dropped variable is not included in the final model, thus excluding predictors and interaction terms that do not significantly improve model fit based on chi-square comparisons of the sums of the squares of the residuals. This process continues until all of the variables that did not improve model fit have been dropped. In each chapter, the full model includes all stated fixed and random effects. The final model is then reported in full, with model outputs created using the *tab_model()* function from the *sjplot* package (Lüdecke, 2021). Although each full model began with random slopes included, these often led to convergence errors when running the model. As a result, only random intercepts are included for the final models reported in each chapter.

3.7 Summary

This chapter has introduced the analytical tools that I will use to explore the social meanings association with linguistic variation in swearing. It has introduced appropriate methods for measuring the effect of swearing on sound identification, grammaticality judgments and reading speeds; these methods will be used to test the degree to which people have implicit knowledge about phonetic, morpho-phonological and semantic variation in swearing. It has also introduced an appropriate method for measuring the effect of each of these forms of variation on social evaluation of a speaker. Finally, it has introduced an appropriate analytical framework for combining findings from these different methods; this framework will help me draw conclusions about how people's unconscious linguistic knowledge of swearing relates to their more conscious social knowledge of swearing.

Chapter 4

Swearing and the socio-phonetic variable

4.1 Introduction

The first domain of language from which variation might have consequences for how swearing is perceived is phonetics. As discussed in Section 2.3.1 of Chapter 2, despite receiving relatively minimal attention in the study of phonetics, several studies have shown that people have an awareness of the types of sounds typically contained within a swearword (Yardy, 2010; Gold & McIntyre, 2016; Tessier & Becker, 2018; Reilly et al., 2020). Furthermore, just like any other subset of the lexicon, swearwords are likely to vary in pronunciation. In some cases, this variability might be the source of variation in socio-indexical meaning. While people may think of words like *fucking* and *shitting* as discrete lexical items, in reality they will be produced by speakers in a variety of ways which are likely to pattern according to who and where they are. Whether or not this variation has consequences for how speakers are perceived by their audience remains to be seen, however. To that end, this chapter focuses on the relationship between swearing and a sociophonetic variable that shares a similar socio-indexical profile, namely variable (ING).

The rest of this chapter continues as follows. In Section 4.2, I provide further detail on the study of sociophonetic variation, an established tradition in sociolinguistics. In Section 4.3, I introduce variable (ING), including a summary of previous research on both production and perception of the variable. In line with the methodological logic of this thesis, I then present two consecutive, related experiments that examine different aspects of variable (ING) in relation to swearing. In Section 4.4, I detail Experiment I of this chapter, a variant categorization task aimed at testing whether listeners *perceive* the relationship between swearing and variable (ING). In Section 4.5, I detail Experiment II of this chapter, a matched guise task aimed at testing whether listeners *notice* the relationship

between swearing and variable (ING) when socially evaluating a speaker. Finally, in Section 4.6, I provide a general discussion and the chapter conclusions.

4.2 Socio-phonetic variation

The earliest work on sociolinguistic variation, and the majority of it since, has focused on variation at the level of sound. Labov's early work, for example, focused on a range of different phonetic variables, including the /aw/ and /ay/ diphthongs in Martha's Vineyard, Massachusetts (Labov, 1963) and postvocalic (r) in New York (Labov, 1966). The focus of these and many subsequent works on production in sociophonetics (see E. Thomas, 2010 for a full review) has been on the social stratification of phonetic variation and of sound change in communities across time. Such research has combined the formal study of the linguistic constraints on phonetic variation with the anthropological study of social factors influencing usage. The social component of language has been promoted as equally as relevant as the referential components (Hymes, 1974).

Labov's formal conceptualisation of the *sociolinguistic variable* (Labov, 1972) was tailor-made for phonetic variation: "social and stylistic variation presuppose the option of saying "the same thing" in several different ways: that is, the variants are identical in reference or truth value, but opposed in their social and/or stylistic significance" (p. 271). This is very much the case for socio-phonetic variation. Variable (ING), the variable of interest in this chapter, concerns the alternation between e.g., *playing* and *playin'*, with the only difference being the use of either velar [ɪŋ] or alveolar [ɪn] for the '*-ing*' suffix. The same is true for Labov's production study of postvocalic (r), with the variable alternating between [flɔ] and [flɔr] (*floor*). In both cases, the form changes, but the meaning is kept constant. Both linguistic and social predictors are then used to model these changes in form. As I will discuss in the equivalent sections of Chapters 5 and 6, the situation is more complicated, but ultimately unproblematic, for variation at the levels of morpo-syntax and semantics/pragmatics.

Phonetic variation has also been studied prominently in work on sociolinguistic perception. Much of this work is related to existing research on production. For example, if a particular sociophonetic variable has been shown to pattern as a function of e.g., age, gender or ethnicity, a perception study may be employed to examine the level of awareness of this patterning in a speech community. The studies often employ tasks to target the processes I identified in Chapter 3, namely *perceiving*, *noticing* and *understanding* (Schmidt, 1990; Squires, 2016).

While there have been a few of studies that have examined the phonetics of swearwords (Yardy, 2010; Gold & McIntyre, 2016; Tessier & Becker, 2018; Reilly et al., 2020), none of these could be said to have examined the sociophonetics of swearwords. That is,

no previous study has examined phonetic variation in swearing as a function of where and by whom the swearword is being used, either in production or perception. The closest is Gold and McIntyre (2016) who measure changes in the duration of the /ʌ/ vowel in occurrences of *fuck* on television as a function of its pragmatic function. They find that longer durations are associated with *disbelief*, while shorter durations are associated with *insults*. Their analysis is based in models of Gricean implicature, however, and does not consider the purely social factors that may be playing a role. The experiments discussed in this chapter are therefore the first to draw on the possible relationship between swearwords and socially meaningful phonetic variation.

4.3 The variable: (ING)

This chapter is concerned with how swearing might interact with variable (ING), a phonetic variable that has been subject to a large amount of research in sociolinguistics, both in terms of production and perception, across different varieties of English including including British, American, Australian and New Zealand English (Fischer, 1958; Trudgill, 1974; Houston, 1985; Bell & Holmes, 1992; Kiesling, 1998; Tagliamonte, 2004; Labov, 2001; Watts, 2006; Hazen, 2008; Kendall, 2013). A large body of work has already established the influence of variable (ING) on social evaluation (Campbell-Kibler, 2005, 2007, 2010a; Labov et al., 2011; Schleef et al., 2017). More specifically, it has been shown that social evaluations linked to variable (ING) can depend on the co-presence of other socially meaningful linguistic variables (Campbell-Kibler, 2011).

Variable (ING) concerns the alternation between velar and alveolar realisations of the inflectional '*-ing*' suffix typically found in gerunds and present participles. The variable is specific to this morpheme and does not typically apply to other words that might end in *-ing* (e.g., *sing* or *thing*)[1]. The standard variant is the velar realisation, written in IPA script as [ɪŋ]. The non-standard variant is the alveolar realisation, written as [ɪn]. There also exists a third variant in North Western varieties of British English, referred to as Velar Nasal Plus and written in IPA script as [ɪŋ g] (Wells, 1982). This variant is much less common than [ɪŋ] and [ɪn] as realisations of the '*ing*' morpheme; it is more common in words such as *wrong* (Watts, 2006). As a result, and due to the very region-specific profile of Velar Nasal Plus (see Bailey, 2019), I will leave this aside. As such, reference to variable (ING) will henceforth be restricted to the velar [ɪŋ] and alveolar [ɪn] variants. I will refer to an occurrence of variable (ING), with either realisation, as a *token* of variable (ING).

Production studies on variable (ING) suggest consistent grammatical and social con-

[1] A small subset of exceptions to this include certain pronouns such as *something* and *nothing*, which are frequently pronounced with the alveolar variant (Labov, 2001).

ditioning. In terms of grammatical conditioning, use of variable (ING) patterns along a noun-verb continuum. The velar [ɪŋ] variant is more common among noun forms (e.g. *ceiling*) and the alveolar [ɪn] is more common among verb forms (e.g. *she's playing*), with other forms somewhere in the middle, such as gerunds (e.g. *she enjoys playing with sand*) (Labov, 2001; Tagliamonte, 2004; Kendall, 2013). Work by Vaughn and Kendall (2018) suggests that listeners are sensitive to this grammatical patterning in perception (see Section 3.3.1 of Chapter 3 for details). Studies on variable (ING) frequently categorize *-thing* compounds (e.g. *anything*) and proper nouns separately, as they tend to pattern differently to other nouns.

There are also phonological constraints to variable (ING). In the UK, when the variable (ING) token is immediately followed by a velar consonant (e.g. *playing **c**ards*), the velar [ɪŋ] variant is more likely, whereas the alveolar variant is more common when the (ING) token is followed by an alveolar consonant (e.g. *hearing **n**ews*) (Houston, 1985); variable (ING) is also subject to Progressive Dissimilation. Variable (ING) production is also influenced by recency (Abramowicz, 2007). If a speaker utters two words with the '*-ing*' morpheme in quick succession, the probability of the second word having the velar [ɪŋ] variant is greatly increased if the first word also had this variant.

In terms of social conditioning, the picture differs across varieties. In American English, the alveolar variant is more commonly used by speakers at the lower end of the socio-economic scale and by men (Labov, 2001, 2006); it is also more common in casual speech. In British English, this conditioning is region-dependent, as there is a significant North-South divide with respect to variable (ING) production. In Northern varieties, there is very little social stratification (Tagliamonte, 2004; Watts, 2006). In Southern varieties, however, there is significant social stratification, trending in a similar direction to Labov's (2001) findings in the USA with respect to socio-economic class and gender (Trudgill, 1974; Houston, 1985; Schleef, Meyerhoff, & Clark, 2011), such that alveolar [ɪn] is more common among men and speakers in lower socio-economic classes.

In perception, variable (ING) appears to be robustly socially meaningful, typically indexing a set of related social meanings connected to its social patterning in production. In multiple studies, Campbell-Kibler (2005, 2007, 2009, 2010a) explores the social meaning of variable (ING) among American listeners via matched-guise tests. Her results suggest that the alveolar variant is associated with the attributes *+relaxed*, *+casual*. *+working class* and *-intelligent*, among other related meanings. The emergence of these meanings in context depends significantly on the broader speech style of the particular speaker (Campbell-Kibler, 2007), including the co-presence of other linguistic variables with potentially contrasting social meanings (Campbell-Kibler, 2011); the activation of these social meanings can also depend on what other information the listener has about the speaker (e.g. *professors* vs *working professionals*) and whether this information is

congruent with those social meanings (Campbell-Kibler, 2010a). Use of alveolar [ɪn] in the USA has also been linked to a decrease in perceived professionalism (Labov et al., 2006, 2011).

In the UK, perception studies on variable (ING) have found different but related results. Firstly, Levon and Fox (2014) attempted to replicate the findings of Labov et al. (2011) with British listeners; their replication was unsuccessful, suggesting that, in the UK, increased use of alveolar [ɪn] is not linked to a decrease in perceived professionalism. Elsewhere, Schleef et al. (2017) explored the social meanings attached to variable (ING) in three different regional accents, namely London, Manchester and Edinburgh, using listeners from those respective regions. For Londoners listening to London accents and, to a lesser degree, Mancunians listening to Manchester accents, use of the alveolar variant was associated with *-articulate -educated*, *-hard-working*, *-rich*, *-posh*, *+working-class* and *+casual*. This was contrasted with Edinburgh-based participants listening to Edinburgh accents, for which use of alveolar [ɪn] was associated with *+friendly*, *+down-to-earth* and *+trendy*. Schleef et al. (2017) suggest a degree of cohesion between results found in the USA (Campbell-Kibler, 2005) and their own results for London, Manchester and middle-class listeners in Edinburgh.

The ultimate aim of this chapter is to test whether listeners *perceive* and/or *notice* sociophonetic variation in swearing. To do this, I used both swearwords and neutral words, comparing the two to see whether swearwords pattern differently. Variable (ING) was the most suitable phonetic variable for this task. It can be attached to the majority of swearwords to create attested morphologically complex words; the same is true for neutral words. It therefore meant that stimuli could be easily constructed and matched for other factors that may affect the perception of the variable.

As I will explain in greater detail in the following section, swearwords share some of the same characteristics as the alveolar [ɪn] variant. Like the alveolar [ɪn] variant, in production, swearwords are used more frequently in informal speech styles and by speakers in lower socio-economic classes ((McEnery, 2004; McEnery & Xiao, 2004; Love, 2017); cf. Love, 2021). Furthermore, like alveolar [ɪn], swearwords are associated with particular socio-indexical meanings, such as lower intelligence (DeFrank & Kahlbaugh, 2019), informality (Stapleton, 2010; Beers Fägersten, 2012; Cavazza & Guidetti, 2014) and working class speech (De Klerk, 1997; Gordon, 1997; Lawson & Milani, 2015). As I will now discuss, these similarities between swearing and one particular variant of variable (ING) might lead one to make particular predictions about how they interact in perception.

4.4 Experiment I: Perceiving phonetic variation in swearing

The previous section described variable (ING) and the way in which its production and perception are linguistically and socially conditioned. I will now turn my attention to examining the extent to which people have an implicit knowledge of variable (ING) in relation to swearwords, as distinct from non-swearwords. That is, do people *perceive* the relationship between swearing and a particular variant of (ING) as a function of their shared socio-indexical meanings?

4.4.1 Methods

I tested people's implicit knowledge of variable (ING) in swearwords using a variant categorization task. I reviewed this and similar tasks in Section 3.3.1 of Chapter 3, but I will briefly repeat the aim of the task here. The task involves rapidly presenting participants with stimuli in the auditory domain. Each time they hear a word, they must state via button press whether the word contained a particular phonetic variant. In my case, this was a forced choice between '*-ing*' and '*-in*', with the stimuli in the experiment consisting exclusively of words ending in a variable (ING) token.

As I will discuss in greater detail below, the stimuli in the experiment came from a 7-step nasal continuum, from maximally velar to maximally alveolar. As such, items from the middle of the continuum were acoustically ambiguous between the velar and alveolar variants. The stimuli were a mixture of phonetically matching swearwords, neutral words and phonotactically licit non-words (e.g., *fucking*, *ducking* and *nucking*). With other aspects of the words matched, including lexical frequency and part-of-speech dominance, I measured for changes in participants' responses for the different word types.

4.4.1.1 Experimental Hypothesis

This experiment had the following hypothesis:

H1 In the middle of the nasal continuum, swearwords will be more likely to be heard as '*-in*' compared to neutral words.

As discussed previously, swearwords and the alveolar [ɪn] variant share some of the same socio-indexical associations and are similarly socially stratified in production. It is possible that people have an internal representation of this relationship. That is, it may be that, alongside their linguistic knowledge of both swearwords and alveolar [ɪn], people also store social information about the types of speakers that use those forms and the types of situations in which they are used. If this is the case, it may be that the presence of one

might cue a listener to expect the other. That is, this stored information may mean that when a person hears a swearword, it might cue them to expect to hear the alveolar variant.

4.4.1.1.1 Expectation in word processing It is widely accepted that social information and linguistic knowledge are retained in memory alongside one another (Foulkes, 2010). A prominent model used to account for this is *exemplar theory* (see Pierrehumbert, 2001; Bybee, 2002; Jackendoff, 2007). Exemplar theories posit that individuals compare novel stimuli to similar instances of those stimuli that they have previously encountered. The inclusion of exemplar theories in sociolinguistics has been argued for by Foulkes and Docherty (2006), among others. Such approaches concern the encoding of social information in a speaker's memory traces of linguistic tokens and the use of that information in linguistic processing.

The influence of social information on linguistic perception concerns expectation and prediction. If a listener thinks that a speaker belongs to a particular social category (e.g. *female*), their predictions about their speech will converge towards the speech styles associated with that category. Following an exemplar model, the social information encoded in the memories of linguistic tokens helps listeners identify novel tokens. Expectations relating to a speaker's identity can influence semantic and pragmatic processing (Van Berkum, Van den Brink, Tesink, Kos, & Hagoort, 2008; Fairchild, Mathis, & Papafragou, 2020), as well as word processing (see Hay, 2018).

A particularly relevant example from sociophonetics is Strand (1999) on the effect of perceived speaker gender on phoneme categorization. Strand's study used minimal pair consonant-vowel-consonant tokens on a synthesized fricative continuum from /ʃ/ to /s/ (e.g., *shod* to *sod*), with participants required to categorize each token accordingly. Participants were presented with a visual prime of either a male or female face. The prime influenced participants' perceptual boundary between /ʃ/ and /s/. The female face pushed the boundary further towards /s/, while the male face pushed the boundary further towards /ʃ/, in-line with production trends for men and women. Essentially, participants decisions were informed by their social expectations for how they expect men and women to sound, based on previous experience.

Expectation also plays a key role in word recognition. In spoken word recognition, information is delivered sequentially to the listener. Each piece of information helps narrow the set of possible candidates for the target word. In Bayesian models of speech perception, listeners combine the available perceptual evidence with knowledge of the prior probabilities of words (Norris & McQueen, 2008). This evidence can include social information. For example, word processing can be influenced by nationality via accent (Cai et al., 2017), emotional tone (Nygaard & Lunders, 2002) and voice age (Walker & Hay, 2011), supporting claims that experiences of words are stored as exemplars with

detailed social and contextual information (Pierrehumbert, 2002; Foulkes & Docherty, 2006).

For example, Walker and Hay (2011) showed that 'older' words, such as *shilling* or *inclined*, were recognised faster in a lexical decision task when spoken by an older speaker than by a younger speaker; the opposite was true for 'younger' words such as *internet* or *physio*. Walker and Hay suggest that words are more easily processed if they contain "characteristics that most resemble the listener's accumulated past experience with that word" (p. 219).

This is the dynamic that could cause implicit knowledge about the shared socio-indexical meanings of swearing and variable (ING) to rise to the surface in a variant categorization task. The task is rapid and involves an automatic process, namely word recognition. On items from the middle of the nasal continuum - for which the acoustic cue to the (ING) token is unclear - listeners would rely on what cues they do have to help them identify the (ING) token. Having successfully processed the first part of the word, the listener could use their knowledge and past experience of that word to make a decision.

We see this dynamic at play with variable (ING) in Vaughn and Kendall (2018) with regard to part-of-speech, albeit with categorical velar and alveolar variants in the stimuli. Participants were significantly faster giving responses to items with the alveolar variant if the item came from a part-of-speech category that was more frequently pronounced with that variant, such as a verb. Similarly, participant accuracy when identifying an alveolar variant was significantly weaker for items from a part-of-speech category that rarely has that variant, such as a noun. Participants' knowledge of the word and its probabilities for appearing with each variant, coupled with the available acoustic evidence, influenced their decision-making process.

4.4.1.1.2 Swearing and variable (ING) Swearing is broadly associated with a particular subset of speakers and speech styles. Swearwords typically index informality (Stapleton, 2010; Beers Fägersten, 2012; Cavazza & Guidetti, 2014). This is shown in work by Bayard and Krishnayya (2001), in which swearing occurred more frequently in unstructured conversations, compared to during task-orientated discussions. Swearing is more tolerated in private or in-group settings compared to more formal and public settings (Mercury, 1995). Swearing is frequently used in stand-up comedy, conversations about sex or when telling stories (T. Jay, 2009a; Seizer, 2011). People also experience a greater *expectation violation* for workplace swearing compared to social settings (Johnson & Lewis, 2010). Finally, in the British National Corpus, swearing is more frequent in the spoken section than the written section (McEnery, 2004; Love, 2017).

In the English-speaking world, swearing is also associated, by some speakers, with

working-class speech (Cheshire, 1982; S. E. Hughes, 1992; Romaine, 1999; Stapleton, 2010). Swearing is often considered to index particularly working-class masculinities (Eckert & McConnell-Ginet, 2013 in Pietilä, Tainio, Lappalainen, & Lahelma, 2020). For working-class men, swearing can be seen as a way of avoiding effeminacy (De Klerk, 1997; Lawson & Milani, 2015). Swearing has particular relevance for working-class women too, however, argues S. E. Hughes (1992). While women in the upper and middle classes have historically avoided swearing (Gordon, 1997), swearing has long been part of the language of working-class women, where swearing indexes localised norms of communicating (S. E. Hughes, 1992).

With regards to social class, the corpus data on swearing is more mixed. In the 1994 edition of the BNC, swearing was inversely correlated with socio-economic group, with swearing most frequent among the lowest socio-economic group (McEnery, 2004). More recent work by Love (2017, 2021) presents a more complicated picture, however. Using a more fine-grained scale of social class used by the Office for National Statistics (2010), with 1 being the highest and 8 being the lowest, Love (2017) finds the highest frequency of swearing in the 4-7 range. In a follow-up article, Love (2021) finds swearing among 'working class' and 'middle class' speakers to be roughly equal; his collapsing of the five different categories of social class used by McEnery (2004) under these binary labels may be hiding some of the variation, however.

As discussed in Section 4.3, the alveolar [ɪn] variant is also associated with a particular subset of speakers and speech styles. In fact, there is considerable overlap with swearing. Both swearing and alveolar [ɪn] are associated with informal speech styles and working-class speakers. Furthermore, Bailey (2015) provides evidence that swearing and alveolar [ɪn] are stylistically linked in orthographic representations on Twitter. While velar [ɪŋ] is preferred overall (e.g., *They are going*), the exception was swearwords and swearing euphemisms (e.g., *fuckin*, *flippin*, *friggin* etc), which showed a preference for the alveolar variant. Admittedly, Bailey's data is taken from Twitter, meaning there is a degree of sample bias with respect to the type of people who use the site. Furthermore, that Twitter users show this bias in writing does not automatically mean they would do so verbally. The study does suggest, however, that people have some stored knowledge of a link between swearing and the alveolar variant.

Additionally, work by Kiesling (1998) on use of variable (ING) in an American fraternity showed higher uses of the vernacular [ɪn] variant in 'confrontational stances', which is also a feature of swearing. Finally, a recent Buzzfeed quiz (C. Aston, 2021) asked participants to answer a series of questions depicting social stereotypes to find out whether they were more likely to say *fucking*, *fuckin* or *fucken*. This suggests a degree of metalinguistic awareness regarding particular personae associated with non-standard pronunciations of a swearword.

If social information is encoded in linguistic tokens, and if this social information can influence listener expectation, we might expect the swearwords to cue listeners to expect to hear the alveolar [ɪn] variant, given that they share some of the same social meanings. That is, if a listener hears the word *fuck*, followed by an '*-ing*' morpheme that is acoustically ambiguous, stored social information relating to swearing may cause the listener to expect the alveolar variant rather than the standard velar variant.

I would not expect this to be the case for a similar-sounding neutral word, such as *duck*, which shouldn't activate the same shared socio-indexical meanings. In general, for neutral words, one would expect a bias for selecting '*-ing*'; Vaughn and Kendall (2018) observed an overall bias towards '*-ing*' using a similar procedure, with the authors suggesting that, unlike '*-in*', '*-ing*' is always available and is unmarked. In the case of swearwords however, an '*-in*' bias was predicted in Experiment I of this chapter.

4.4.1.1.3 A caveat There is one significant piece of evidence missing from this prediction, namely production evidence of an alveolar [ɪn] bias in swearwords in spoken English. The predictions made by Vaughn and Kendall (2018) for part-of-speech category and variable (ING) were motivated by previous production studies showing the grammatical conditioning of variable (ING) production (Labov, 2001; Kendall, 2013). The prediction that participants would use their knowledge of that conditioning in perception was motivated by participants' own experience with the variable.

I do not have equivalent data for swearwords, beyond Bailey's (2015) work on orthographic representations of swearwords and swearing euphemisms on Twitter. Freely available phonetically transcribed corpora of British English, or any variety of English for that matter, are rare. The majority of those which are accessible feature read speech, which is much less likely to contain swearwords than spontaneous speech. While the British National Corpus does include phonetic transcription of the spoken portion, these are regularized to the standard pronunciation, such that every transcription of an (ING) word is transcribed with the velar variant.

In the absence of such data - the collection of which could be a research project in its own right - I am relying on indirect evidence of an association between swearwords and the alveolar [ɪn] variant via shared socio-indexical meanings and similar frequencies of usage among particular groups of speakers.

4.4.1.2 Stimuli

This experiment employed a 3x7 experimental design. The 3-level factor was the categorical variable Item Type, with the levels Swearword, Neutral Word and Non-word. The 7-level factor was the ordinal variable Continuum Step, starting with Step 1 (maximally velar) and ending with Step 7 (maximally alveolar). The task used a repeated measures

design, such that every participant heard at least one item in each combination of the experimental manipulations.

4.4.1.2.1 Selecting the items The items for this experiment were phonetically-matched monosyllabic swearwords, neutral words and phonotactically licit non-words combined with the suffix '*-ing*'. Nine swearwords were selected from previous studies on British swearing (McEnery, 2004; Love, 2017). All nine were included in Ofcom's guide to offensive language in broadcast media (2016). In a normed study of tabooness (Janschewitz, 2008), all nine were well above the upper quartile (1.57). All items also formed licit words when combined with '*-ing*'. Neutral items were taken from the UK SUBTLEX database (Van Heuven, Mandera, Keuleers, & Brysbaert, 2014) and non-word items from the ARC non-word database (Rastle, Harrington, & Coltheart, 2002). Non-word items were included to act as a second baseline. A number of other word-specific factors might influence participant expectations regarding variable (ING), such as lexical frequency and part-of-speech dominance (Vaughn & Kendall, 2018). While the chosen swearwords and neutral words have been matched, where possible, for such factors, the non-word items were included to be used as an alternative comparison class in the event that, for neutral words, other factors drown out any possible effects that result from the swear vs non-swear distinction. Non-words should not contain the same expectancy biases that real words do.

Swearwords were phonetically matched with neutral and non-words to create 9 minimal triplets, totalling 27 test items. Both the swearword and neutral items in each triplet were coded for log lexical frequency and part-of-speech dominance using the UK SUBTLEX database and valency and arousal using measures from Warriner et al. (2013). All nine swearing items, including the less common *cunting* and *twatting*, are in regular usage on Twitter and have their own entries in the Oxford English Dictionary (2020) and the Urban Dictionary.

The list of 27 test items is provided in Table 4.4.1.2.1. As metadata was not available for all of the full items (e.g. *fuck+ing*), information on valency and arousal relates to the word stems without the '*-ing*' morpheme; no frequency or part-of-speech dominance information was available for *cunting*, although this is nonetheless attested (*Oxford English Dictionary*, 2020).

Table 4.1: Experiment I: List of items

Item	Triplet	Item Type	SUBTLEX Log Freq	Dom POS	Stem Valency	Stem Arousal
hitching	bitching	word	2.767617734	verb	5.29	4.22
bitching	bitching	swear	2.508649083	verb	2.55	5.95
vitching	bitching	nonword	NA	NA	NA	NA
flapping	crapping	word	3.389462675	verb	5.74	3.45
crapping	crapping	swear	2.26393745	verb	3.44	4.62
blapping	crapping	nonword	NA	NA	NA	NA
blunting	cunting	word	1.841863762	verb	4.65	4.61
cunting	cunting	swear	NA	NA	3.7	6.1
yunting	cunting	nonword	NA	NA	NA	NA
slamming	damning	word	3.001087096	verb	4.79	4.61
damning	damning	swear	3.397303711	noun	4.32	5.1
framming	damning	nonword	NA	NA	NA	NA
picking	dicking	word	4.550163232	verb	5.91	3.62
dicking	dicking	swear	2.038158407	verb	3.06	5.6
zicking	dicking	nonword	NA	NA	NA	NA
ducking	fucking	word	3.165557742	verb	6.11	4
fucking	fucking	swear	4.689171957	verb	5.23	7.14
nucking	fucking	nonword	NA	NA	NA	NA
kissing	pissing	word	3.855904019	verb	7.78	6.05
pissing	pissing	swear	3.102275907	verb	2.73	4.74
tissing	pissing	nonword	NA	NA	NA	NA
gritting	shitting	word	2.967577333	verb	4.55	4.22
shitting	shitting	swear	3.03419222	verb	3.91	5.67
plitting	shitting	nonword	NA	NA	NA	NA
chatting	twatting	word	4.030591416	verb	5.75	4.27
twatting	twatting	swear	1.598825713	verb	4.32	5.39
yatting	twatting	nonword	NA	NA	NA	NA

4.4.1.2.2 Recording the items Stimuli were recorded by a speaker of Standard Southern British English. Recordings were made in a sound attenuated booth using a Neumann TLM103 microphone via the RME fireface UX audio interface. Recordings were made at a 16 bit, 44.1kHz sample rate. For each item, three versions were recorded using the alveolar realisation and three with the velar realisation. To achieve the best continuum, the alveolar and velar tokens best matched for amplitude and duration were chosen. Pro-

nunciations were kept constant and, where possible, devoid of sociophonetic variables that could bias listeners. For example, t-glottaling is socially stratified in a similar way to (ING) in the UK (Williams, Kerswill, Foulkes, & Docherty, 1999). The choice was therefore made not to glottal in any of the items where this was possible (e.g. in *shitting*).

One important element of the items to check before they were manipulated was the qualities of the vowel sound in the (ING) tokens. Velar realisations of (ING) typically have longer, higher and fronter vowels, more closely approximating the FLEECE vowel, than alveolar realisations, which more closely resemble the KIT vowel (Vaughn & Kendall, 2018). If this was true of these stimuli, this wouldn't have been problematic. As detailed below, the program used to create a continuum between the 2 sounds morphs the recordings in their entirety; this includes the pre-nasal vowel. As such, the potentially ambiguous tokens of (ING) contained an equally ambiguous pre-nasal vowel.

More important in the acoustic analysis were possible differences in the pre-nasal vowels between the different item types. If listeners were getting different auditory cues to either swearwords, neutral words or non-words, it would interfere with any word-level effects. If, for example, all of the swearwords had a consistently different value for F2 in their pre-nasal vowels, it is possible that listeners could have been sensitive to this as a cue. In sum, I wanted to know whether the presence of swearword would bias a listener to pick [ɪn] for ambiguous tokens; if (ING) forms were not uniform across all items, participants may have reacted to auditory cues rather than word-level cues.

To test this, all items were entered into Praat (Boersma, 2019) for acoustic analysis. Using Textgrids, the /ɪ/ vowel from the (ING) of each item was segmented manually to include between 40 and 50 milliseconds in the middle of the vowel. All items were then run through a Praat script to automatically extract the first and second formants at the midpoint of the vowel; F1 is mostly determined by the height of the vowel and F2 by the frontness/backness (E. Thomas, 2010, p. 145). F1 and F2 were extracted for both variants of all items (N = 27), resulting in 54 measurements of F1 and F2 respectively. To test for acoustic differences, both the F1 and F2 measurements were analysed in R using group-wise t.tests, with each group (Swearword, Neutral word and Non-word) compared manually. There were no significant differences in F1 or F2 between any of the three Item Type groups (all $p > 0.05$).

4.4.1.2.3 Manipulating the items To create authentic-sounding nasal continua, items were manipulated using the MATLAB-based program TANDEM-STRAIGHT (Kawahara et al., 2008). TANDEM-STRAIGHT is regularly used in studies of paralinguistic speech perception (Andics et al., 2010; Bestelmeyer, Rouger, DeBruine, & Belin, 2010). Recently it has also been employed in sociolinguistic perception studies on race (Zheng & Samuel, 2017), sexuality (Kachel et al., 2018) and dialectology (Bukmaier, Harrington, &

Kleber, 2014). For more information on the technical aspects of TANDEM-STRAIGHT, see Kawahara, Morise, Banno, and Skuk (2013).

TANDEM-STRAIGHT analyses the input sound, breaking it down into an interference-free spectrogram, an aperiodicity map and a fundamental frequency (F0) trajectory. This process was performed for both the velar and alveolar recordings of each test item (e.g., *fucking* and *fuckin*, *ducking* and *duckin*). TANDEM-STRAIGHT then uses these factors to generate natural-sounding synthesized voices by morphing two sounds together using manual mapping of time-frequency anchors on their respective spectrograms (Skuk & Schweinberger, 2014). The two sounds in a pair were loaded into a distance matrix (see Figure 4.1) to compare their STRAIGHT spectrograms. Frequency anchors were then manually placed along the centre line of the matrix to map together the meaningful segments of two sounds e.g., the beginning and end of each phoneme and significant formant transitions. TANDEM-STRAIGHT then morphed the paired sounds together, in their entirety, generating a continuum of a pre-selected number of steps (see Kawahara et al., 2009 for more detail). Through this process, all aspects of the word were morphed into a continuum, including the pre-nasal /ɪ/ vowel. For this experiment, 7-step continua were constructed. This process was completed for all 27 test items, generating 189 test stimuli.

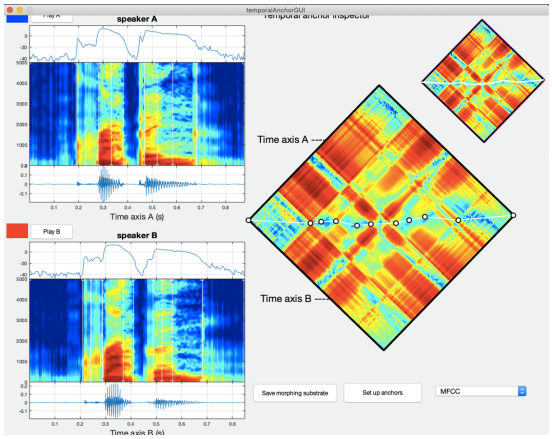

Figure 4.1: TANDEM-STRAIGHT: Distance Matrix for two sounds

4.4.1.2.4 Norming the items To test for by-group acoustic differences, the nasal continua were normed on a separate set of participants. All 189 test stimuli were manually segmented in Praat to remove the first syllable. For example, the minimal triplet *fucking*, *ducking* and *nucking* became *king*, *king* and *king*, at each of the 7 continuum steps. 45 participants, paid £0.76 and recruited via Prolific Academic (2022), completed an online variant categorization task following the same procedure as the main task (see Section 4.4.1.3), including headphone check and practice trial elimination criteria. On each trial, participants were required to select, via keyboard button press, whether the word they

heard ended with '*-ing*' or '*-in*'. Participants completed 54 test trials, including 2 trials per item. Item and Continuum step were counterbalanced across participants.

Results were analysed using a logistic mixed-effect regression model in R with the *lme4* package (Bates, Mächler, Bolker, & Walker, 2015). Response was included as the dependent variable (1 = '*-ing*', '*-in*' = 0). Item Type (Swear vs Neutral vs Non-word) was treatment coded and included as a categorical predictor. Continuum Step (centred) was included as a continuous predictor. By-Participant and By-Word random intercepts were also included; random slopes were initially included, but were removed due to convergence issues. The model summary is provided in Table 4.2.

Table 4.2: Experiment I: Norming model summary - Response ~ Continuum Step + Item Type + (1|Participant) + (1|word)

	Response			
Predictors	*Odds Ratios*	*std. Beta*	*Statistic*	*p*
(Intercept)	1.94	1.94	1.94	0.052
Continuum Step	0.34	0.12	-25.87	**<0.001**
Type: Non-word	0.74	0.74	-0.62	0.534
Type: Neutral word	0.57	0.57	-1.17	0.244
$N_{Participant.Private.ID}$	45			
N_{word}	27			
Observations	2404			

No significant differences were found between the baseline type (Swear) and neutral or non-word types. The graph in Figure 4.2 suggests that the nasal continuum construction in TANDEM-STRAIGHT was successful. This graph plots the mean response for all items (1 = '*-ing*', 0 = '*-in*'), as well as two times the standard error, at each step in the nasal continuum. Step 1 is the maximally velar end of the continuum and Step 7 is the maximally alveolar.

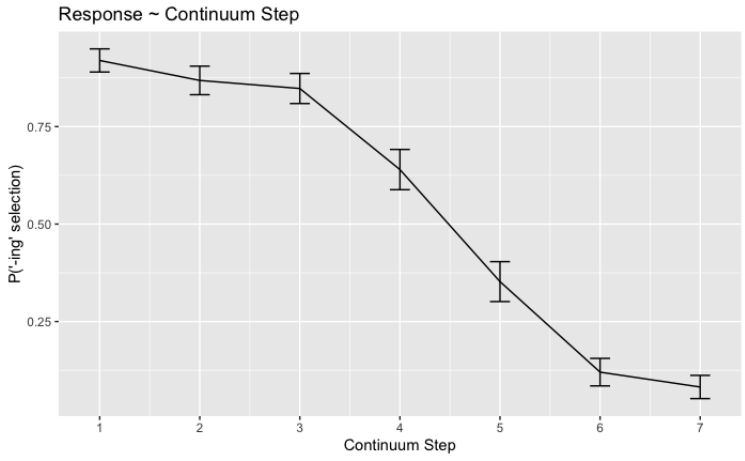

Figure 4.2: Experiment I: Norming experiment line-graph: Response ~ Continuum Step

4.4.1.2.5 The final items The 189 test items were divided into seven lists of 27, with exactly one occurrence of each word per list. The lists were counterbalanced for continuum step, such that each list contained the same number of items from each step. This was done using a Williams Latin square. An example list of test stimuli is included in Table 4.3.

Item	Step	Item	Step	Item	Step
fucking	1	ducking	4	nucking	2
bitching	7	hitching	6	vitching	3
crapping	5	flapping	6	blapping	7
cunting	5	blunting	4	yunting	1
damning	2	slamming	3	framming	3
dicking	6	picking	2	zicking	7
pissing	1	kissing	5	tissing	3
shitting	4	gritting	7	plitting	4
twatting	3	chatting	1	eating	2

Table 4.3: Experiment I: Example trial sheet

The same 50 fillers were added to each list. These consisted of a mixture of other neutral words and non-words also ending in '*-ing*'; fillers were unambiguously pronounced with either velar [ɪŋ] or alveolar [ɪn]. All stimuli were then normalized to 70db in Praat.

4.4.1.3 Procedure

The task was created and hosted on Gorilla (Anwyl-Irvine, Massonnié, Flitton, Kirkham, & Evershed, 2018). Participants were first required to pass a headphone check to ensure concentration (Woods, Siegel, Traer, & McDermott, 2017). The headphone check was a 3-alternative forced choice "Which tone is quietest?" task with 200Hz pure tones. Participants were asked to listen to three sounds and choose which one was quietest; they did this three times. On each trial, one of the three sounds was played in antiphase across the stereo channels. As a result, the sound would be heavily attenuated if played through speakers, but not if played through headphones. It would therefore have been very difficult to complete three correct trials without using headphones (Woods et al., 2017). Participants were allowed two attempts at the headphone check; this was due to high rates of failure in pilot tests.

They were then given instructions for the experiment. As this task required participants to make explicit judgments about a sociophonetic contrast, rather than a phonemic one, the instructions included an introduction to the variable. Instructions were as in (1).

(1) *In this task you will hear a series of words. Every word in the task ends with the letters '-ing' e.g.* ***smashing***, ***gazing***, ***taking*** *etc. As you might know, there are two ways of pronouncing words that end with '-ing'. One way is to pronounce the full sound e.g.* ***taking***. *Another way is to only pronounce the '-in' sound e.g.* ***takin'***. *The words you will hear will either have the '-ing' or '-in' pronunciations. Your job is to tell us which one you hear. Using the* ***left*** *and* ***right*** *buttons on your keyboard, you need to select which ending you heard for each word. You need to do this quickly, as you will be timed out if you take too long to decide. To complete the task, you will use the arrow keys on your keyboard. At the bottom of the screen there are two labels. One is labelled '-ing' and one is labelled '-in'. To select the label on the left side of the screen, press the left arrow key. To select the label on the right side of the screen, press the right arrow key.*

Participants then completed eight practice trials to familiarise them with the task. The stimuli for these trials were four high frequency neutral words, not included in the main task, pronounced categorically with either the velar or alveolar realisation of variable (ING), totalling 8 practice trials. No spectral manipulation was performed on the practice stimuli. Participants were required to pass 6 out of 8 practice trials to progress. If they failed, they were given a second chance with the same pass threshold.

On each practice trial, participants were automatically played a sound clip of one of the eight practice items with the order randomised. On the computer screen, participants were presented with 2 buttons: one button labelled '-ing' and the other '-in', symbolising the orthographic representations of the velar and alveolar realisations of (ING). The order of the buttons was randomised across participants, but not within the experiment; changing the order of the buttons between trials would have interfered with the automaticity of processing. Participants controlled these buttons using the left and right arrow keys on their keyboard. On each practice trial, participants received instant feedback on whether they were right or wrong via either a green tick or a red cross at the bottom of the screen.

Once participants had completed the practice trials, they were given more instructions. These instructions reminded them to complete the task quickly and thoroughly. They then began the test phase of the experiment. The procedure for test trials was the same as for the practice trials unless otherwise stated. Each participant was randomly assigned to one of the seven trial sheets consisting of 27 test stimuli plus 50 fillers. Stimuli were pseudo-randomised to ensure that, while the order was random, no two 2 items from the same minimal triplet appeared immediately after one another. Participants were given 2500 milliseconds in which to provide a response before the next sound was played automatically. On test trials, participants received no feedback on performance. Figure 4.3 illustrates the experimental procedure.

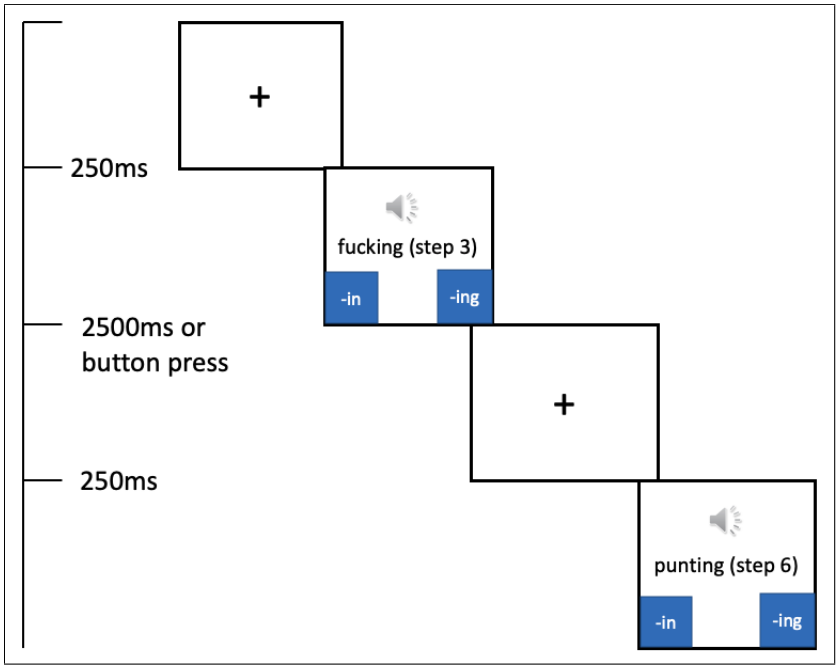

Figure 4.3: Experiment I: Experimental procedure

After completing all 77 test trials, participants were asked to provide basic demographic information, including age, sex, sexuality, level of education and yearly household income (in blocks of 10,000). Finally, they were asked a question relating to their own social class identification. They were asked to rate, on two separate Likert scales, the degree to which they felt the labels 'working class' and 'middle class' applied to them. This final question aimed to measure the participants' own self-orientation towards the categories of 'middle class' or 'working class'.

4.4.1.4 Participants

456 British English speakers were recruited via Prolific Academic (2022). 57 participants were rejected for failing the headphone check and 19 participants were rejected for failing practice trials, leaving 385 participants (F = 211, M = 170, Other = 4[2]) for statistical analysis. Participants were paid £0.55.

4.4.2 Results

Incomplete trials (n = 400) and responses below 500ms (n = 39) were discarded. Within Item Type, swearwords received the most '*-ing*' responses (58%), followed by neutral

[2]These four participants consisted of two non-binary, one trans and one genderqueer participant. For statistical purposes, they have been grouped together under *Other*. I must acknowledge that this is not ideal, but it is preferable for the purpose of testing for differences between participants with normative and non-normative genders.

words (50.7%) and non-words (46.1%). There was a marginal bias overall for '*-ing*' responses (51.3%) across participants.

A logistic mixed effect regression model was run in R (R Core Team, 2018) using the *lme4* package (Bates et al., 2015), with '*-ing*' vs '*-in*' as the categorical dependent variable ('*-ing*' = 1, '*-in*' = 0). The full model included Item Type (Swear, Neutral Word, Non-word), Dominant Part-of-Speech (Dom POS) (Van Heuven et al., 2014), Participant Gender, Participant Sexuality and Participant Level of Education as categorical predictors (all treatment coded) and Continuum Step (centred), Duration (in milliseconds), Preceding Continuum Step (1-7), log lexical frequency (Van Heuven et al., 2014), stem arousal and valency (Warriner et al., 2013), stem tabooness (Janschewitz, 2008), and Participant Income as continuous predictors. Factors which improved model fit after model comparison were further tested as interacting predictors with Item Type. Random intercepts were included for Participant and Word. A random slope for Participant over Word was originally included, but was dropped due to issues with model convergence.

The model reported in Table 4.4 excludes predictors and interaction terms that did not significantly improve model fit based on chi-square comparisons of the sums of the squares of the residuals. This was done using the *anova* function in R. There were significant main effects for Continuum Step ($\beta = -1.2$, $p < 0.001$), Item Type: Non-word ($\beta = -0.74$, $p = 0.001$), Item Type: Neutral Word ($\beta = -0.54$, $p = 0.014$) and Gender: Other ($\beta = 0.9$, $p = 0.005$).

Table 4.4: Experiment I: Model summary - Response ~ Continuum Step + Item Type + Pre-step + Participant Gender + (1|Participant) + (1|word)

	Response		
Predictors	*std. Beta*	*Statistic*	*p*
(Intercept)	0.46	2.88	**0.017**
Continuum Step	-1.20	-43.71	**<0.001**
Item Type: Non-word	-0.74	-3.34	**0.001**
Item Type: Neutral Word	-0.54	-2.45	**0.014**
Pre-Step	0.04	1.50	0.134
Gender: Male	0.01	0.15	0.883
Gender: Other	0.90	2.83	**0.005**
$N_{ParticipantID}$	380		
N_{word}	27		
Observations	10012		

Figure 4.4 plots the mean probability of a listener selecting '*-ing*' for each Item Type at each step in the nasal continuum (see Figure 4.5 for a colour-blind friendly version) . Error bars display two times the standard error. This graph logically follows the results

of the model output, suggesting that, at steps 3 to 5 in the nasal continuum, there was a difference between swearwords, on the one hand, and neutral and non-words, on the other.

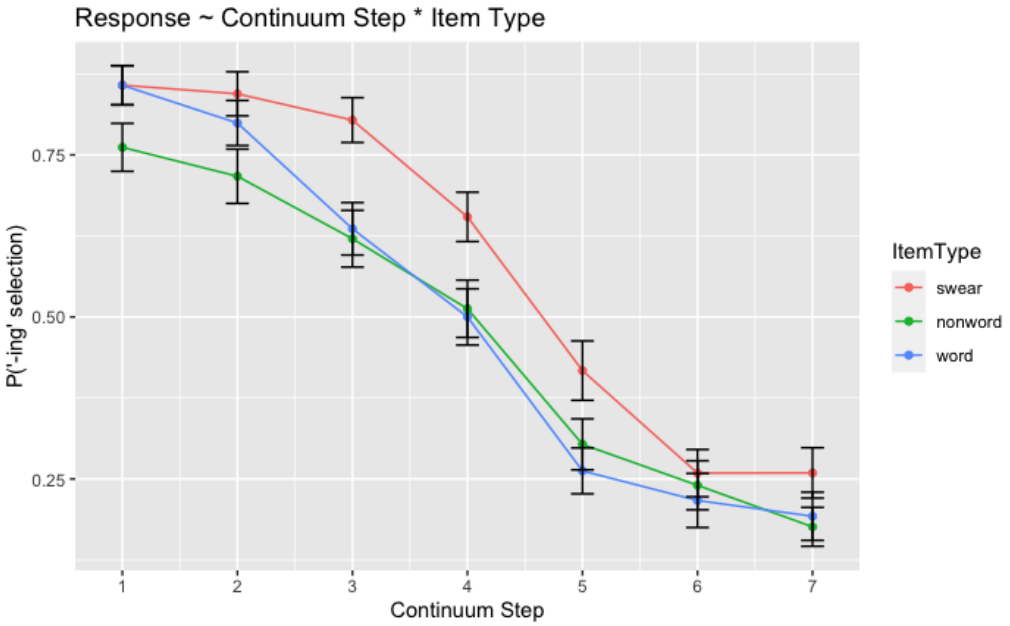

Figure 4.4: Experiment I: Mean probability of '-ing' selection ~ Continuum Step * Item Type

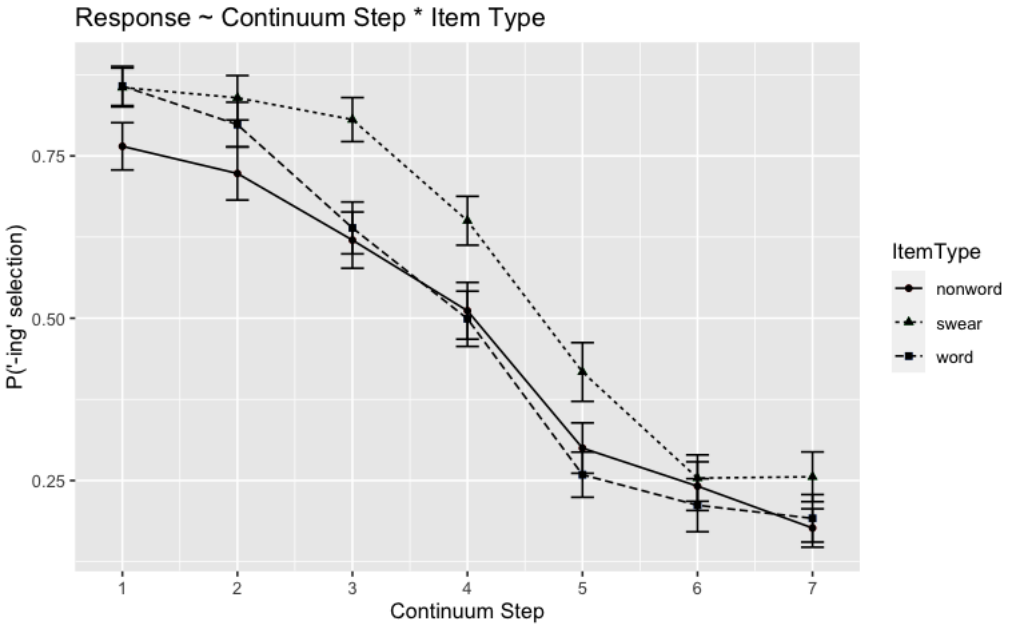

Figure 4.5: Experiment I: Mean probability of '-ing' selection ~ Continuum Step * Item Type - colour blind-friendly version

4.4.2.1 Reaction time data

In light of these results, further exploratory analysis was done using the reaction time data collected on each trial. A linear mixed effect regression models was constructed in R using the *lme4* package (Bates et al., 2015), with Reaction Time (RT) in milliseconds (centred) as the dependent variable. Item Type (Swearword as baseline) and Continuum Step (centred) were included as predictors, with Participant and Word as random intercepts; random slopes were initially included. but were dropped due to convergence issues. The model summary is provided in Table 4.5.

Table 4.5: Experiment I: Model summary - RT $\sim$ Item Type + Continuum Step + (1|Participant) + (1|word)

		RT	
Predictors	*std. Beta*	*Statistic*	*p*
(Intercept)	0.10	1.36	0.175
Item Type: Non-word	-0.06	-0.67	0.502
Item Type: Neutral Word	-0.21	-2.28	**0.023**
Continuum Step	-0.03	-3.30	**0.001**
$N_{ParticipantID}$	385		
N_{word}	27		
Observations	10267		

There was a significant main effect for Item Type: Neutral Word ($\beta = -0.21$, $p = 0.023$), such that, compared to swearwords, responses to neutral words were significantly faster; there was no significant difference in reaction time between swearwords and non-words. There was also a significant main effect for Continuum Step ($\beta = -0.03$, $p = 0.001$), such that, as the continuum step increased (i.e., became more alveolar), reaction times also increased.

4.4.3 Interim discussion

The hypothesis for Experiment I is repeated below:

H1 In the middle of the nasal continuum, swearwords will be more likely to be heard as '*-in*' compared to neutral words.

The logic behind this hypothesis was that, upon hearing a swearword with an ambiguous (ING) token, listeners would rely on their knowledge about swearwords to make an accurate prediction about the (ING) token. As swearwords and the alveolar [ɪn] variant have been shown to index similar social meanings and have similar usage patterns across

macro-demographic categories (social class), it was predicted that, for swearwords with ambiguous (ING) tokens, participants would be more likely to select the '*-in*' button.

In reality, the opposite effect occurred. Swearwords with acoustically ambiguous (ING) tokens were more likely to be identified as containing the velar [ɪŋ] variant compared to their neutral counterparts. This was also true for swearwords compared with non-words to an even greater extent, with non-words receiving an '*-in*' response on 53.9% of trials; non-words were the only items to receive a greater proportion of '*-in*' responses than '*-ing*' responses. The social characteristics of swearwords did not facilitate access to the alveolar variant. A possible explanation is that, despite evidence from the literature linking alveolar [ɪn] and swearwords to similar speakers and speech styles, as outlined in Section 4.4.1.1.2, this association was not strong enough to prompt listeners to expect the alveolar variant on swearwords.

A tendency towards the velar variant for swearwords is unexpected. This effect would suggest a socio-indexical association between swearwords and velar [ɪŋ]. One possible explanation for this finding is that swearwords and velar [ɪŋ] share an association with careful and effortful speech. While swearwords do index informality, this may be a different type of informality from the casual/relaxed informality indexed by alveolar [ɪn]; it may be that swearwords reflect a more performative and controlled informality. Furthermore, swearwords can function to create *verbal emphasis* (Stapleton, 2010). In Stapleton (2003), a study on swearing in Irish English, participants reported the creation of verbal emphasis as one of the most common reasons for swearing.

In the attention-paid-to-speech model used by Labov (1966, 1972), speech styles are categorized based on the degree to which speakers are thought to be consciously thinking about the way in which they are talking. The categories principally used are *casual*, *careful* and *read*. Importantly, the use of vernacular forms is typically lowest in read speech, the style in which speakers pay the least attention to their speech, with casual speech typically having the highest frequency of vernacular forms (Labov, 1972). For example, velar [ɪŋ] is most common in read speech, followed by *careful* then *casual* speech. If swearwords too are associated with careful, deliberate speech, in the form of verbal emphasis (Stapleton, 2010), rather than casual speech, then this may have primed listeners to expect the velar [ɪŋ] variant on swearwords.

There is an alternative cognitive explanation of the findings, however. Swearwords take up increased attentional resources compared to neutral words, as has been demonstrated in tasks where participant concentration is on a non-linguistic stimulus, including modified Stroop and attentional blink tasks ((MacKay et al., 2004; Guillet & Arndt, 2009; Eilola & Havelka, 2011; Mathewson et al., 2008; Dhooge & Hartsuiker, 2011); a review of such studies is available in Section 2.3.5 of Chapter 2). In Stroop tasks for example, participants are presented with individual words on a screen. They are required to identify

the colour of the writing, ignoring the content of the word. Performance in such tasks is significantly slower for swearwords than neutral words (MacKay et al., 2004; Eilola & Havelka, 2011). A similar effect was shown in the auditory domain by Bertels et al. (2011) in French. Having been familiarised with four different voices labelled with names, participants heard individual words and were required to identify the speaker on each trial. For blocks of either taboo words or neutral words, participants were significantly slower at voice identification for low valency and high taboo words. In each of these examples, the additional attentional resources required to process swearwords impeded the performance of another cognitive task.

This is also a possible explanation for the observed tendency towards velar [ɪŋ] that was found for swearwords in the current study. On swearing items, it is possible that participant attention was drawn away from the variable (ING) token, towards the taboo stem. With their ability to focus on the critical sound impaired, participants may have defaulted to the most likely option overall: velar [ɪŋ]. The reader will recall that an overall bias for the velar variant was also found by Vaughn and Kendall (2018). While it was predicted that the social significance of swearwords would make the alveolar variant more likely for those items, this may not have been an sufficient cue. This explanation receives some support from the reaction time data showing that participants were significantly slower to respond to swearing items than neutral items.

Finally, it is possible that the nature of the participant pool is responsible for the observed effect. As noted in Section 3.5 of Chapter 3, the participant pool available on Prolific Academic is not necessarily representative of the population at large; the effects predicted in Section 4.4.1.1.2 of this chapter may therefore not have held for these participants. Furthermore, there was no check in the experiment for whether the participants themselves were predominantly users of velar [ɪŋ] or alveolar [ɪn]. If the majority of the participants were, in fact, velar [ɪŋ] users, this could explain the effect; it is plausible that such participants might hypercorrect to velar [ɪŋ] for swearwords, particularly if alveolar [ɪn] was less natural to them. Further scrutiny of the results, or the collection of further data with more detailed information about participants, might reveal interesting patterns; this is a concern for future research, however.

While both explanations are plausible, further work is required to isolate the exact cause of the effect. In the absence of positive evidence for either of them I can only conclude the following: rather than having a bias towards alveolar [ɪn] as predicted, swearwords with ambiguous (ING) tokens were in fact more likely to be identified as containing velar [ɪŋ], while the neutral words with ambiguous (ING) tokens were roughly at chance (see Figure 4.4).

4.5 Experiment II: Noticing socio-phonetic variation in swearing

Experiment I established that listeners do *perceive* the relationship between swearing and variable (ING), albeit not in the way I expected. That is, the result of Experiment I suggests that the listeners had created some internal representation of the relationship between swearing and the velar [ɪŋ] variant. In isolation, listeners showed a tendency towards the velar [ɪŋ] variant for swearwords when the (ING) token was acoustically ambiguous. This either suggests that they may have some implicit knowledge of the relationship between the two or that they perceived the swearwords differently from the neutral words to an extent that this affected their perception of the (ING) token.

Experiment II extends this to the use of swearwords in full sentences. Experiment II tested whether listeners *noticed* the relationship when conducting the conscious process of social evaluation. If they did, this would suggest that listeners' perception of swearing and variable (ING) has social significance.

The conditions of Experiment II more closely matched those of real life. Rarely do individuals hear swearwords ending in '*-ing*' spoken in isolation. Rather, exclamations of *fuck!* or *shit!* are perhaps more common as stand-alone swearwords. People are much more likely to encounter *fucking*, *shitting* and *pissing* used in full sentences. Presenting words in isolation also removes potentially relevant information that might guide listener expectation, such as syntactic category (Vaughn & Kendall, 2018). Furthermore, unlike the relatively unnatural process of variant categorization, in which listeners' attention is explicitly directed towards a specific part of a word, social evaluation is an everyday process that has real-world consequences (see e.g. Roberts, Davies, & Jupp, 2014). People are constantly forming judgments about people, consciously or unconsciously, as a function of how they speak. They do so on the basis of a person's entire speech style, including the presence and combination of sociolinguistic variables, as well as their pre-existing expectations of the speaker based on their identity.

If listeners perceive variable (ING) tokens differently on swearwords than on neutral words, and thus assign different traits and characteristics to speakers as a result, this has consequences what we can say about linguistic variables as carriers of social meaning. The potential for a form to activate particular social meanings can be modulated by the co-presence of other socially meaningful linguistic variables (Campbell-Kibler, 2011; Levon, 2014). If this is also true for the socio-indexical meanings indexed by particular sets of lexical items - such as swearwords - this must be accounted for in models of sociolinguistic cognition.

This would also have consequences for what we can say about swearwords themselves as carriers of social meaning. As I established in Section 2.4 of Chapter 2, swearing is

hugely socially meaningful, affecting how people are perceived on a broad range of evaluative dimensions. If swearing also affects how other linguistic sources of social meaning are perceived, however, this would suggest that swearwords are socially salient to the extent that, unlike neutral words, their exact pronunciation is inconsequential. This, in turn, would have consequences for the generalizability of previous findings from sociolinguistic perception.

4.5.1 Methods

Experiment II used an auditory matched-guise task (Lambert et al., 1960). Participants listened to a short narrative recorded by speakers of Standard Southern British English. The passage contained a mixture of swearwords and neutral words ending with '*-ing*. The speakers were recorded using both the velar and alveolar realisations of variable (ING). The recordings were then artificially manipulated in Praat (Boersma, 2019) to create four experimental conditions: fully velar (*All-ing*), fully alveolar (*All-in*), only swearwords as velar (*Swear-ing*) and only neutral words as velar (*Swear-in*); the conditions are fully schematised in Table 4.6. Participants heard multiple speakers in multiple conditions, rating each speaker on a series of Likert scales depicting traits shown to be sensitive to variable (ING) in previous perception studies.

Recording	Swearing words	Neutral words
All-ing	velar	velar
All-in	alveolar	alveolar
Swear-ing	velar	alveolar
Swear-in	alveolar	velar

Table 4.6: Experiment II: Test recording key

4.5.1.1 Experimental Hypothesis

Experiment II had the following hypotheses:

H1 Recordings in the *All-ing* condition will be rated higher on the test scales *articulate*, *hard-working*, *posh*, *rich* and *educated* and lower on the test scales *working class* and *casual* than recordings in the *All-in* condition.

H2 Responses for recordings in the *Swear-in* condition will approximate responses for recordings in the *All-ing* condition, but responses for recordings in the *Swear-ing* condition will not.

Hypothesis 1 follows from the results found in Schleef et al. (2017) for British English for a southern English accent. If the alveolar [ɪn] variant is socio-indexically associated

with traits linked to e.g., lower social status, then recordings that are saturated with that variant (*All-in*) should be rated higher for those traits and lower for traits with the inverse association. The null hypothesis was that there would be no statistically significant difference between *All-ing* and *All-in*; this would mean that, for these listeners and these voices, variable (ING) is not socially meaningful on these evaluative dimensions. Hypothesis I was the minimal hypothesis; if no support was provided for Hypothesis 1, it would not be possible to find support for Hypothesis 2.

Hypothesis 2 follows from the results of Experiment I of this chapter. If the effects identified in Experiment I hold in full sentences, we would expect an increase in perceiving the velar variant for swearwords as compared to neutral words. For neutral words, we would not expect this to be the case - neutral words were roughly at chance in the middle of the nasal continuum in Experiment I (see Figure 4.4). This was expected to have different consequences for the *Swear-ing* and *Swear-in* guises.

For the *Swear-in* guises, in which the swearwords had the alveolar [ɪn] variant, it was expected that participants would tend towards hearing these as the velar [ɪŋ] variant instead due to effects observed for swearwords in Experiment I. The neutral words, which had the velar [ɪŋ] variant, were expected to be unaffected by this effect. All of the critical words in the *Swear-in* guises would therefore tend towards being heard as velar tokens of variable (ING). As such, responses to those guises were expected to be similar to the guises in which all the critical words really did have velar tokens of variable (ING), namely the *All-ing* guises.

The same was not expected to occur for the *Swear-ing* guises. As the swearwords in these guises already had the velar [ɪŋ] variant, the effect found for swearwords in Experiment I should not have made a difference. The neutral words, which had the alveolar [ɪn] variant, were expected to be heard as such. Listeners should therefore have heard the *Swear-ing* guises as 50% [ɪŋ] and 50% [ɪn]. As such, the availability of alveolar [ɪn] tokens should have elicited different responses to the *All-ing* guises. The null hypothesis in this case was that responses to both the *Swear-in* and *Swear-ing* guises would be significantly different to those to the *All-ing* guises, as both had 50% of each variant.

The shift from words in isolation to words embedded in full sentences is significant. In Experiment I, when words were presented in isolation, the critical stimuli had acoustically ambiguous tokens of (ING) taken from the middle of a nasal continuum. In Experiment II, the '*-ing*' words are either categorically velar or alveolar realisations. Importantly, however, while one might therefore expect listeners to be able to easily identify the unambiguous (ING) token on each word, their presentation in full sentences means that the participants' concentration would not be trained on the (ING) token specifically. The participants were instructed to listen to the whole passage and, unlike in Experiment I, the instructions made no mention of variable (ING). While the tokens of (ING) would have

been easier to identify, therefore, participants would have also been processing the rest of each speaker's speech, potentially meaning that the effect observed in Experiment I might also emerge in Experiment II while their attention was spread across the whole passage.

4.5.1.2 Stimuli

The stimulus used in this task was a passage taken from a previous study on sociolinguistic perception in the UK using speakers from a similar area (Levon, 2014). The passage recounts a mildly dramatic episode (someone falling down the stairs in a London underground station). This passage was suitable because, while consisting of read speech, the nature of the narrative made the inclusion of swearwords realistic; this type of passage contrasts with the newscaster passage used in similar studies on (ING) (Labov et al., 2006, 2011), in which the inclusion of swearwords would have been surprising.

The original version from Levon (2014) was edited for the current study to include swearwords and to ensure there was an equal number of swearing and non-swearing ending in variable (ING). Not all swearwords appear equally as frequently in naturally occurring speech, however. For example, while attested, the use of e.g., *cunting* or *twatting* may be distracting for participants when included in a narrative. As such, two instances of *fucking* and *shitting* and one instance of *dicking* were included in the passage. The neutral (ING) words were all verbs, the grammatical category in which variable (ING) is most common (Labov, 2001; Tagliamonte, 2004; Kendall, 2013), except for the first *fucking* which was a modifier. The swearing and neutral (ING) words were ordered alternately throughout the passage. The passage is included in (2), with (ING) words in bold:

(2) So last Thursday I was **walking** down the steps to the tube and there were these two guys on the other side. One of them kept **dicking** around, and then he slipped and all of a sudden he was **falling** backwards. And for like three seconds he was just sort of balanced there, and I thought he'd pull himself up. But then he tipped even further back and just **fucking** tumbled down the stairs and landed on the floor, his head right on the tiles. He seemed alright at first, but then blood started **streaming** from his head. The guys mate just stood there and looked at him like 'are you **shitting** me?'. For a minute it didn't look like he was **breathing**. I shouted at him to call an ambulance, and they came pretty soon but the guy was **shitting** himself. He turned out to be alright - he just had a big cut on the side of his head from **hitting** the tiles. Serves him right for **fucking** around on the stairs.

Four speakers of Standard Southern British English were recorded reading the passage aloud. The four speakers were from the same geographical area - Kent, near South-east London - and attended the same school. Due to limitations on face-to-face research

brought about by the COVID-19 pandemic, the four speakers recorded themselves in their own homes using headphones with a built-in microphone. They recorded directly into their mobile phones using voice recording apps while sat in a quiet room on their own. Each speaker first recorded the passage once through with no directions on how to pronounce particular words. Participants then recorded themselves reading the passage twice more, first with all (ING) words pronounced using the velar [ɪŋ] realisation and then with all (ING) words pronounced using the alveolar [ɪn] realisation.

Read speech was chosen over spontaneous, conversational speech to allow for a sufficient level of control over the inclusion of the number neutral words and swearwords that included the '*-ing*' morpheme. While read speech has the disadvantage of being less authentic, work by Tamminga (2017) using variable (ING) suggests that matched guise tasks using both speech styles can produce the same results, validating the use of read speech in sociolinguistic perception research.

Recordings of each speaker were manually manipulated in Praat (Boersma, 2019), creating four test recordings per speaker. For each, the first 'naturally read' recording was used as the base. For the first test recording, the *All-ing* guise, all of the speakers' carefully articulated velar [ɪŋ] tokens were copy and pasted into the base recording, following the approach of Campbell-Kibler (2005), resulting in a recording exclusively containing velar [ɪŋ] realisations. If the speaker had already pronounced an (ING) word with a clear and unambiguous velar [ɪŋ] token, a new one was not copied in. For all recordings, the copied-in tokens included the preceding consonant, e.g., the [kɪŋ] segment in *fucking*. For the second test recording, the *All-in* guise, I followed the same process, but using the carefully articulated alveolar [ɪn] tokens, resulting in a recording exclusively containing alveolar [ɪn] realisations. For the third test recording, the *Swear-ing* guise, velar [ɪŋ] tokens were copied in for swearing (ING) words and alveolar [ɪn] tokens were copied in for neutral (ING) words. For the fourth test recording, the *Swear-in* guise, alveolar [ɪn] tokens were copied in for swearing (ING) words and velar [ɪŋ] tokens were copied in for neutral (ING) words. A key for the test recordings is included in Table 4.6.

These manipulations resulted in a total of sixteen test recordings, with four recordings for each speaker and four recordings in each recording condition (*All-ing*, *All-in*, *Swear-ing*, *Swear-in*). These recordings were divided into four test blocks. Each test block contained exactly one recording from each speaker and exactly one recording in each test condition. For example, one test block contained *All-ing* for speaker 2, *All-in* for speaker 3, *Swear-ing* for speaker 4 and *Swear-in* for speaker 1. Each recording was scaled for intensity to 70db in Praat (Boersma, 2019). The recordings were informally normed for naturalness on a set of native British English speakers.

The Likert scales included in Experiment 2 were based on the results of previous work on perceptions of variable (ING). In Schleef et al. (2017), (ING) variation was shown to

influence perceptions of scales of *articulate*, *hard-working*, *posh*, *rich*, *working class*, *casual* and *educated*. All of these were main effects for one (ING) variant over the other that were found for speakers with London accents. As such, all of these scales were included in the current task. In addition, 2 distractor scales were chosen, namely *tall* and *attractive*. All scales were from 1 to 8, with the main scale label (e.g. *articulate*, *hard-working* etc) at the end closest to 8. The label for the other end of the scale was taken from Schleef et al. (2017) e.g., *inarticulate* or *lazy*.

4.5.1.3 Procedure

This task was constructed using Qualtrics. After giving their consent for completing the task, participants were given the instructions provided in (3).

(3) In this task you are going to hear a group of people telling a story. They will each tell the same story of something they saw happen at a tube station recently. In total you will hear the story told by 8 different speakers. Each clip will last around 45 seconds. After listening to each clip, we want you to use the scales below to evaluate the speaker. For example, if you think that the speaker sounded more intelligent than average, you might choose a number on the scale from *unintelligent* to *intelligent* that is close to the *intelligent* label. Please listen to each clip at least once through to the end before starting the scales. Once you've rated that speaker on each scale, please click the *next* button to move on to the next speaker. Once you've listened to all 8 speakers, we will ask you a few questions about yourself. If you're happy you understand these instructions, please click on the *continue* button to start the task.

After reading the instructions, participants began the task. Participants were randomly assigned 2 of the 4 test blocks, meaning that each participant heard 8 recordings, including exactly 2 recordings of each speaker and exactly 2 recordings in each condition. The order and assignment of the blocks was randomised between participants; this was counterbalanced such that every combination of test blocks was assigned to a roughly equal number of participants. The order of trials within each block was randomised between participants but kept constant between blocks within each participant. That is, if a participant heard speaker 1 first in their first block, they heard speaker 1 first in their second block. There was no break between blocks, with all 8 recordings presented to participants as 8 distinct speakers.

On each trial, participants were presented with a play button and a set of scales. The play button allowed the participant to begin playing the sound clip at their own discretion. It also allowed participants to pause the clip if necessary. Participants were able to rewind

the clip and listen to it multiple times if they wished. Participants were instructed to listen to the entirety of each clip before beginning providing a response on the Likert scales; participants who failed to do this were excluded without pay, as they had been warned would happen.

Below each play button were 10 scales (all 8 test scales and 2 distractor scales). The order of scales was randomised between trials. Participants were required to provide a response for every scale; they were unable to continue to the next trial if a scale was left blank. Participants were also asked to estimate the age of each speaker. On 2 of the 8 trials, there was an additional scale to control for participant attention when responding to the sound clip. This scale only included a label at the top end of the scale which stated 'Move to number 6'. Participants were warned before beginning the task, via instructions on Prolific, that there would be attention checks and that their payment would be withheld if they failed these checks. Any participant that failed these checks was rejected and their responses were excluded from statistical analysis. Unlike in Experiment I, a headphone check was not used, as the variable (ING) tokens were all categorical velar or alveolar realisations, rather than sitting on a nasal continuum with slight gradations between each token.

Once participants had completed all scales for all 8 recordings, they then completed two short post-experiment surveys. The first was a swearing survey, measuring their own swearing behaviour, their attitudes towards swearing and their associations of swearing with particular social groups (e.g. *men*, *working-class speakers*, *younger people*). Participants were presented with the following questions and statements:

- How often do you swear? (1 = never, 5 = very frequently)
- How often do you think swearing is appropriate? (1 = never, 5 = very frequently)
- Swearing is more common among men (1 = Completely disagree, 5 = Completely Agree)
- Swearing is more common among working class people (1 = Completely disagree, 5 = Completely Agree)
- Swearing is more common among less articulate people (1 = Completely disagree, 5 = Completely Agree)
- Swearing is more common among younger people (1 = Completely disagree, 5 = Completely Agree)

The aim of this survey was to examine whether participants' responses to the test stimuli were modulated by their pre-existing attitudes towards swearing and belief in

myths about how swearing patterns socially. The second survey was a basic demographic survey that elicited participants' gender, age, yearly household income and self-defined social class.

4.5.1.4 Participants

200 speakers of British English were recruited for this task via Prolific Academic (2022). They were paid £1.20 for their participation.

4.5.2 Results

Results were analysed using linear mixed effect regression models which were constructed in R (R Core Team, 2018) using the *lme4* package (Bates et al., 2015). Responses for 8 participants were removed from the analysis for failing catch trials. Separate models were constructed for each response scale. In each, the response scale (1-8) was included as the continuous dependent variable. The full models in each case had the treatment-coded categorical predictors Condition (*All-ing*, *All-in*, *Swear-ing*, *Swear-in*), Participant Gender (Female, Male) and Participant UK Location (e.g. South of England, North of England, Scotland etc), and the continuous predictors Participant Age and Participant Income. Participant Education, originally an ordinal variable from 'no formal education' to 'PhD', was re-coded as the categorical variable University (yes, no: sum coded) based on responses. Participant Social Class, originally a 5-level ordinal variable from 'lower working class' to 'upper class' was re-coded as a categorical variable to include only Working Class and Middle Class as factor levels (sum coded). In addition, participant responses to questions and statements in the swearing survey, with responses from 1 to 5, were included as continuous predictors. Random intercepts were included in all models for Participant (ProlificID) and Speaker. A random slope for Condition over Speaker was initially included, but this prevented models from converging, so were dropped.

From the full models, variables that did not significantly improve model fit based on chi-square comparisons of the sums of the squares of the residuals were removed. This was done manually in R using the *anova* function. The full model was gradually stepped down to simpler models using an alpha level of 0.05. For all scales except for *articulate* and *working class*, only Condition significantly improved model fit. For *articulate*, the model was also significantly improved by responses to the statement 'Swearing is more common among less articulate people' (Swearing = LAP). For *working class*, the model was also significantly improved by including Participant Social Class. Condition was treatment coded, with *All-ing* as the baseline factor level. Where multiple variables improved model fit, interactions between these variables were also tested; no interactions improved model fit for any scale. Observations for 3 participants are missing for the *work-*

ing class model, as these participants failed to answer provide an answer for the Social Class survey question.

Table 4.7 summarises the results for the scales *articulate*, *rich*, *working class* and *educated*, on which the expected effects were found. For *articulate*, there were significant main effects for Condition: *All-in* ($\beta = -0.24$, $p < 0.01$), Condition: *Swear-ing* ($\beta = -0.22$, $p < 0.01$) and Swearing = LAP ($\beta = -0.12$, $p < 0.05$). For *rich*, there were significant main effects for Condition: *All-in* ($\beta = -0.32$, $p < 0.001$) and Condition: *Swear-ing* ($\beta = -0.17$, $p < 0.05$). For *working class*, there were significant main effects for Condition: *All-in* ($\beta = 0.32$, $p < 0.01$), Condition: *Swear-ing* ($\beta = 0.29$, $p < 0.01$) and Social Class: Middle ($\beta = -0.4$, $p < 0.001$). For *educated*, there were significant main effects for Condition: *All-in* ($\beta = -0.28$, $p < 0.001$) and Condition: *Swear-ing* ($\beta = -0.22$, $p < 0.01$). On these four scales, there were no significant effects for Condition: *Swear-in*.

Table 4.7: Experiment II: Model summaries for *articulate*, *rich*, *working class* and *educated* - Response ~ Condition + 'Swearing = LAP' + 'Social Class' + (1|Participant) + (1|Speaker)

	articulate			**rich**			**workingclass**			**educated**		
Predictors	*Estimates*	*Statistic*	*p*	*Estimates*	*Statistic*	*p*	*Estimates*	*Statistic*	*p*	*Estimates*	*Statistic*	*p*
(Intercept)	5.42	17.75	**<0.001**	4.74	14.76	**<0.001**	4.08	12.00	**<0.001**	5.47	17.79	**<0.001**
Condition: Swear-in	-0.10	-1.17	0.243	-0.07	-0.85	0.394	0.18	1.77	0.077	-0.12	-1.57	0.117
Condition: Swear-ing	-0.22	-2.62	**0.009**	-0.17	-2.09	**0.037**	0.28	2.84	**0.005**	-0.22	-2.75	**0.006**
Condition: All-in	-0.24	-2.89	**0.004**	-0.32	-4.00	**<0.001**	0.31	3.13	**0.002**	-0.28	-3.52	**<0.001**
N	192 $_{\text{ProlificID}}$			192 $_{\text{ProlificID}}$			192 $_{\text{ProlificID}}$			192 $_{\text{ProlificID}}$		
	4 $_{\text{speaker}}$			4 $_{\text{speaker}}$			4 $_{\text{speaker}}$			4 $_{\text{speaker}}$		
Observations	1536			1536			1536			1536		

To visualise these effects across the four experimental conditions, Figure 4.6 plots the mean response for each condition on each of the four response scales for which significant effects were found. Error bars represent two times the standard error. The conditions are ordered such that *All-ing* and *All-in*, the conditions in which (ING) realisations were 100% velar and 100% alveolar, are at either ends of the plots.

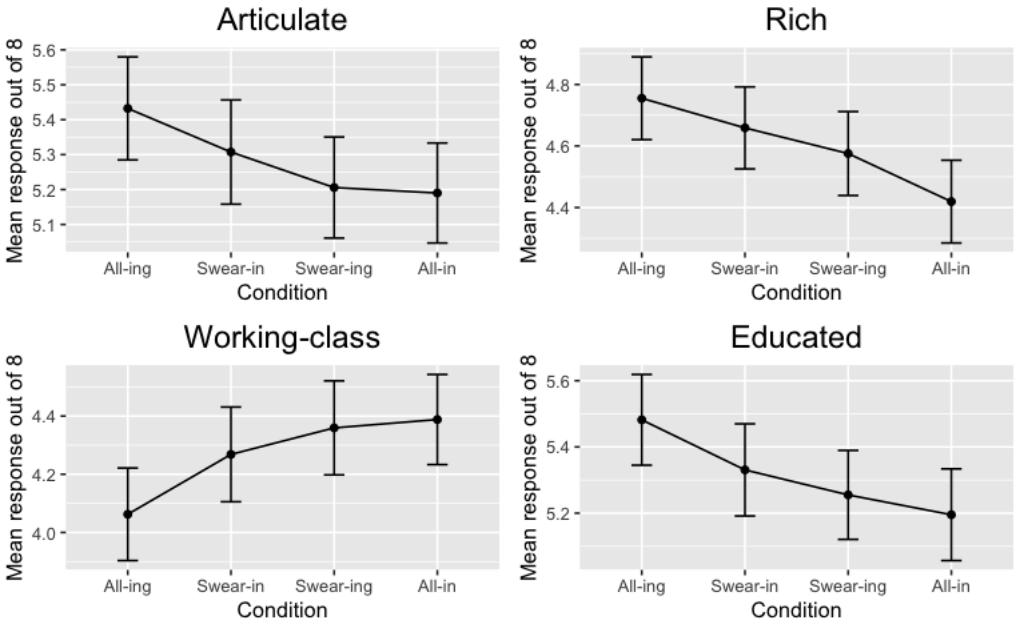

Figure 4.6: Experiment II: Mean responses for the scales *articulate*, *rich*, *working class* and *educated* by condition. *All-ing* = 100% velar, *All-in* = 100% alveolar, *Swear-ing* = swearwords are velar, *Swear-in* = neutral words are velar.

Table 4.8 summarises the results for the scales *hard-working*, *casual* and *posh*. There were no significant main effects for *hard-working* or *casual* for any condition. For *posh*, there were significant main effects for Condition: *All-in* ($\beta = -0.41$, p < 0.001), Condition: *Swear-ing* ($\beta = -0.27$, p = 0.005) and Condition: *Swear-in* ($\beta = -0.22$, p = 0.025). Figure 4.7 plots the mean and two times standard error for each of these scales by condition.

Table 4.8: Experiment II: Model summaries for *posh*, *hard-working* and *casual* - Response ~ Condition + (1|Participant) + (1|Speaker)

	hardworking			**casual**			**posh**		
Predictors	*Estimates*	*Statistic*	*p*	*Estimates*	*Statistic*	*p*	*Estimates*	*Statistic*	*p*
(Intercept)	4.92	29.19	**<0.001**	5.07	13.98	**<0.001**	4.88	11.77	**<0.001**
Condition: Swear-in	-0.03	-0.43	0.668	0.16	1.52	0.128	-0.41	-4.24	**<0.001**
Condition: Swear-ing	-0.02	-0.31	0.759	0.20	1.85	0.064	-0.27	-2.84	**0.005**
Condition: All-in	0.03	0.43	0.669	0.08	0.77	0.438	-0.22	-2.24	**0.025**
N	192 $_{\text{ProlificID}}$			192 $_{\text{ProlificID}}$			192 $_{\text{ProlificID}}$		
	4 $_{\text{speaker}}$			4 $_{\text{speaker}}$			4 $_{\text{speaker}}$		
Observations	1536			1536			1536		

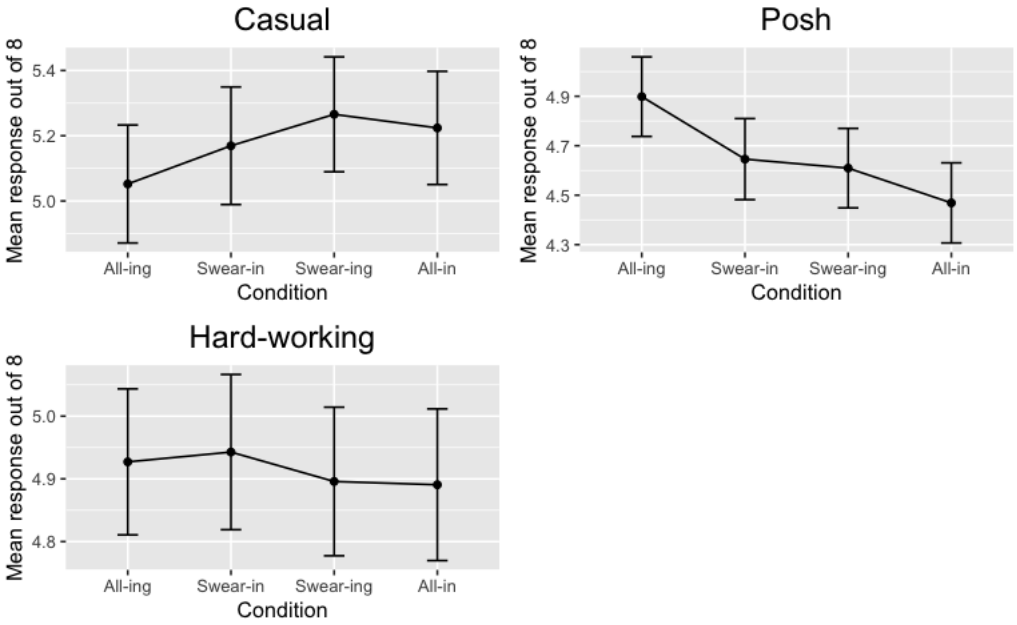

Figure 4.7: Experiment II: Mean responses for the scales *casual*, *posh* and *hard-working* by condition. *All-ing* = 100% velar, *All-in* = 100% alveolar, *Swear-ing* = swearwords are velar, *Swear-in* = neutral words are velar

4.5.3 Interim discussion

These results provide some support for both experimental hypotheses. Support for Hypothesis 1 is provided by the significant main effects found on five different scales, namely *articulate*, *rich*, *working class*, *educated* and *posh*, for Condition: *All-in*. These effects trend in the expected direction. On each scale, the guise saturated with the alveolar variant (*All-in*) was rated as significantly less articulate, rich, posh and educated and significantly more working-class than the guise saturated with the velar variant i.e., the *All-ing* guise. This replicates the finding of Schleef et al. (2017) that, in a London accent, variable (ING) can influence how speakers are socially evaluated on social class- and education-related scales; this is also in line with perception studies in the US showing similar evaluative trends (e.g., Campbell-Kibler, 2005). Unlike Schleef et al. however, variable (ING) had no effect on the scales *hard-working* or *casual*.

Support for Hypothesis 2 is provided on four of the scales, namely *articulate*, *rich*, *working class* and *educated*. On each of these scales, there was a significant main effect for Condition: *Swear-ing* but no significant main effect for Condition: *Swear-in*. Compared to the velar-saturated *All-ing* guises, the guises with velar [ɪŋ] on swearwords and alveolar [ɪn] on neutral words (the *Swear-ing* guises) were considered less articulate, rich and educated and more working-class. The same was not true for the guises with alveolar [ɪn] on swearwords and velar [ɪŋ] on neutral words (the *Swear-in* guises), which were not rated significantly differently from the velar-saturated guise. The similarity between

responses to the *All-ing* and *Swear-in* guises is not absolute, as illustrated in Figure 4.7, suggesting that the effect of an increased probability of hearing the velar variant is not as strong for all swearwords and/or for all participants; rather, this reflects a more general trend in the data.

To illustrate the effect, consider the sentences in (4) taken from the test passage. The sentence in (4-a) is an orthographic representation of how the sentence was pronounced in the *Swear-in* guise. The results suggest that, in this guise, even though the word *dickin* had the alveolar [ɪn] variant, participants may have heard it as *dicking*; the same is true of the other four swearwords in the full test passage. With all five swearwords having some increased probability of being hard as '*-ing*', and the five neutral words actually being pronounced with '*-ing*', the whole passage may have been heard as being similar to the *All-ing* passage, hence the similarity in responses across the four Likert scales.

(4) a. One of them kept dickin' around, and then he slipped and all of a sudden he was falling backwards

b. One of them kept dicking around, and then he slipped and all of a sudden he was fallin' backwards

For the sentence in (4-b), an orthographic representation of how the sentence was pronounced in the *Swear-ing* guise, the neutral word *fallin'* would have been unaffected by the '*-ing*' effect that was expected for swearwords; the same is true for the other four neutral words. The whole passage would therefore be heard as having 50% [ɪŋ] and 50% [ɪn]. The presence of five alveolar [ɪn] tokens was sufficient to cause a difference in responses from the velar saturated *All-ing* guise.

Notably, support was not provided on the scales *hard-working*, *casual* or *posh*. For *hard-working* and *casual*, as already discussed, there were no effects of any kind as a function of manipulations in variable (ING). For *posh*, which had shown the expected effect for the *All-ing* and *All-ing* guises, both the *Swear-in* and *Swear-ing* guises differed significantly from the *All-ing* baseline condition. The beta coefficients for these two guises suggest that they might be following the expected trend; the *Swear-in* guise is closer to the *All-ing* guise and the *Swear-ing* guise is closer to the *All-in* guise. Pairwise comparisons suggest that the two are not significantly different however ($p > 0.05$) (see also Figure 4.7, which suggests that responses to the *Swear-in* and *Swear-ing* guises were very similar). The lack of the expected null effect for the contrast between the *All-ing* and *Swear-in* guises prevents me from rejecting the null hypothesis in this case.

As with Experiment I, the nature of the mechanism causing the increased probability of hearing [ɪŋ] for swearwords in Experiment II is unclear without further research, with the same two competing explanations posited in Section 4.4.3 being plausible. On the one hand, the use of swearwords for deliberate verbal emphasis (Stapleton, 2010) in the

passage may have cued the velar [ɪŋ] variant. This might be particularly plausible here given that the task used read speech, rather than spontaneous conversational speech. On the other hand, the effect may result from the alternative explanation of the increased attentional resources required for processing swearwords, leading to a blocking effect of the (ING) token. In the *Swear-in* guise for example, with the alveolar [ɪn] variant on swearwords impaired but the velar [ɪŋ] variant on neutral words unimpaired, listeners may have assumed that the speaker would be consistent in their realisations of (ING); having heard the 5 velar [ɪŋ] tokens on neutral words, they may have generalised to the impaired tokens and assumed that all critical words were realised with [ɪŋ]. Each of these explanations is rather speculative however, so I will remain neutral on the issue.

4.6 Chapter discussion and conclusions

The results of Experiment I suggest that listeners *perceive* the relationship between swearwords and variable (ING), such that swearwords caused listeners to expect the velar [ɪŋ] variant when they were presented in isolation. The results of Experiment II further suggest that listeners *notice* the relationship when conducting social evaluation of a speaker, such that swearwords also lead listeners to expect the velar [ɪŋ] variant when they are embedded in full sentences, which in turn can influence the social information they extract from the speech signal to inform their judgments about the speaker.

Returning to the overall aims of this thesis - to examine whether people are sensitive to language-internal factors in the social evaluation of swearing - the findings of this chapter provide some insight from the perspective of phonetics, although it must be noted that the direction of the effect is not what was originally predicted.

When planning these experiments, the overarching hypothesis had been that variable (ING) might affect how a swearword was perceived. Previous work had suggested this to be true for neutral words (Campbell-Kibler, 2005; Labov et al., 2011; Schleef et al., 2017), with the alveolar variant linked to low status speakers and informal speech styles. For example, someone saying [plaɪjɪŋ] (*playing*) would be perceived differently from someone saying [plaɪjɪn] (*playin*), with the latter expected to be perceived as more informal, less articulate, more working class etc, provided that it occurred in the right speech style (Campbell-Kibler, 2007) and provided that the variable was socially meaningful in that way for the people listening (Schleef et al., 2017).

A similar effect was predicted for swearwords, for example the comparison between someone uttering [fʌkɪŋ] and someone uttering [fʌkɪn]. It was expected that the effect for swearwords might differ from that for neutral words, however, due to the social significance of the swearword itself. Previous work on variable (ING) has shown that the copresence of other socially meaningful linguistic variables can affect how variable (ING)

is perceived (Campbell-Kibler, 2011) (see Section 3.4.1 of Chapter 3 for a review), with two co-present sociolinguistic variables activating social meanings that were more than the sum of their parts. With swearwords and the alveolar [ɪn] variant sharing some of the same socio-indexical meanings, it was predicted that this combination would produce a unique profile of social meanings.

The results of Experiment I were expected to provide evidence of a perceptual association between swearwords and alveolar [ɪn]. Experiment II was then expected to test whether this association would lead to instances of swearwords realised with the alveolar [ɪn] variant being perceived as even less polite, capable, intelligent etc than swearwords usually are (Howell & Giuliano, 2011; DeFrank & Kahlbaugh, 2019). In Experiment II, an association between swearing and alevolar [ɪn] might have been expected to prompt listeners to think they heard e.g., '*fuckin*' when they actually heard '*fucking*', leading to the inverse of the effects that were actually observed in Experiment II.

The results of Experiment I pointed to a different association between swearing and variable (ING) however. As a result, in the design of Experiment II, rather than expecting variable (ING) to affect how swearwords would be perceived, it was expected that swearwords would affect how variable (ING) was perceived. This turned out to be the case, with swearwords more likely to cause listeners to think they heard the velar [ɪŋ] variant, even when they actually heard the alveolar [ɪn] variant. The direction of the effect, as well as the nature of the effect, was therefore the opposite of what had been expected.

Regarding the question I posed in Section 1.1 of Chapter 1 as to whether all pronunciations of a swearword are perceived the same, the answer seems to be they sometimes they are, at least in the case of pronunciations of (ING). While *playing* and *playin'* activate different social meanings (Campbell-Kibler, 2005; Schleef et al., 2017), the same does not always seem to be true of *fucking* and *fuckin*, *shitting* and *shittin*, *dicking* and *dickin* and so on. In cases where there is evidence that the speaker uses the velar variant elsewhere in their speech, my findings suggest that swearwords with the alveolar variant are sometimes treated as if they contain the velar variant.

In Chapter 7, in which I will bring together the findings from Chapters 4-6, I will expand further on what this result means for swearing research and for the overarching aims of this thesis. Before moving on however, I will briefly discuss the consequences that the findings of the current chapter have for sociolinguistics more broadly. These findings shed further light on the complicated nature of sociolinguistic perception with respect to variable (ING). We know based on previous research that perceptions of variable (ING) depend on the accent of the speaker (Campbell-Kibler, 2007; Schleef et al., 2017), the information the listener has about the speaker (Campbell-Kibler, 2010a) and, as already discussed, the co-presence of other sociolinguistic variables (Campbell-Kibler, 2011). My findings on swearing suggest that they also depend on information contained in the

word to which the variable (ING) token belongs.

This is not the first finding of this kind. Recent work by Vaughn (2021) has shown that the perceptual effects for part-of-speech category tested in Vaughn and Kendall (2018) also affect the social meanings indexed by the variants of (ING). Vaughn (2021)'s matched-guise study on variable (ING) suggests that the typical grammatical category of a word ending with the alveolar [ɪn] variant can affect the perceived *professionalism* of a speaker; when the overall rate of [ɪn] was low, there was a bigger social penalty when alveolar [ɪn] appeared on words that were typically nouns than on those that were typically verbs. Use of [ɪn] on nouns is less frequent (Kendall, 2013). Furthermore, while [ɪn] is more common in words with a higher lexical frequency, this effect is dampened in words that are typically nouns (even if the word is functioning as a verb) and amplified in words that are typically verbs (see Forrest, 2017 on frequency in favorable contexts). Occurrences of [ɪn] on words that were typically nouns were more marked than those on words that were typically verbs. The more marked occurrences of [ɪn] were more available for indexing lower levels of *professionalism* (Vaughn, 2021). In sum, this internal linguistic constraint affects the social meanings that are associated with an individual usage of a sociolinguistic variable.

The findings reported in this chapter suggest that the presence of swearwords may have a similar effect, although this is less directly linked to patterns of usage than grammatical category. On neutral words, the alveolar [ɪn] variant is more likely to activate the social meanings one would expect of that variant than if it attached to swearwords. By some mechanism, the alveolar variant is more noticeable on neutral words; something about swearwords, be it stylistic or attentional, rendered the [ɪn] variants on swearwords less noticeable to participants for social evaluation. This, combined with the findings of Vaughn (2021), suggests that word-level information affects the activation of social meaning by variable (ING). My findings, therefore, motivate further work into the role of such information in the perception of sociolinguistic variation more broadly.

Chapter 5

Swearing and the socio-morphological variable

5.1 Introduction

I will now move on to the second domain of language where I expect variation in swearing to have consequences for social evaluation of the speaker, namely morphology. As briefly discussed in Section 2.3.2 of Chapter 2, there is a considerable amount of morpho-syntactic variation in swearing. Numerous examples provided by Hoeksema (2019) point to the innovative ways in which swearing can be used in idiomatic set phrases; the formulaic nature of a lot of swearing expressions was included by Ljung (2010) as one of the four criteria he proposed for swearing. This chapter will focus on a form of expressive morphology (Zwicky & Pullum, 1987) which can act as a source of social meaning for speakers to modulate how their swearing is perceived, namely infixation.

The rest of this chapter is structured as follows. In Section 5.2, I discuss morpho-syntactic variation in the context of sociolinguistics and explain how infixation can be treated as a sociolinguistic variable in the same way that variable (ING) was in the previous chapter. In Section 5.3, I provide a formal morpho-phonological account of expletive (swearing) infixation. In line with the methodological logic of this thesis, I then present two consecutive, related experiments that examine different aspects of swearing infixation. In Section 5.4, I detail an acceptability judgment task that tests whether people *perceive* the linguistic constraints on swearing infixation, as well as generating the most natural-sounding infixed constructions to be used in a sociolinguistic perception experiment. In Section 5.5, I detail a matched-guise task that tests whether people *notice* swearing infixation, and its linguistic constraints, in the process of social evaluation. Finally, in Section 5.6, I discuss the findings of this chapter and give my conclusions.

5.2 Socio-morphogical variation

The previous chapter dealt with variation at the level of sound, a form of variation that has been widely studied in sociolinguistics, in both production and perception. Historically, structural variation has received significantly less attention in sociolinguistics, particularly in perception. This is particularly true for morphological variation, with a few notable exceptions to be discussed later in this section.

The reason for the relative scarcity of research on morphological variation in sociolinguistics is that changes in morphology typically either carry meaning (as with the majority of derivational morphology and compounding) or affect grammaticality (as with most inflectional morphology). As a result, morphological variation does not conform to the expectation that a sociolinguistic variable involves two ways of saying "the same thing" (Labov, 1972, p. 271). The variable studied in the previous chapter fit this definition; I compared two ways of saying a swearword like *fucking*, with either the velar or alveolar realisations of the '*-ing*' morpheme. Both [fʌkɪŋ] and [fʌkɪn] *express* the same thing in a truth-conditional sense. The same does not tend to be true for morphology.

One type of morphology for which this is not the case is *expressive morphology* (Zwicky & Pullum, 1987). Zwicky and Pullum (1987) differentiate between plain and expressive morphology. Expressive morphology, they say, "is associated with an expressive, playful, poetic, or simply ostentatious effect of some kind" (p. 335). Zwicky and Pullum (1987) suggest that expressive morphology can be differentiated from plain morphology based on several interesting characteristics including pragmatic effects, promiscuity of input, interspeaker variation and special syntax. One example that they cite is *shm-reduplication* (see also Nevins & Vaux, 2003), as in (1).

(1) transformations shmansformations

The expression in (1) gives rise to a dismissive not-at-issue meaning, roughly summarised as 'I don't care about transformations'. Meaning can be described as at-issue if it is relevant to the Question Under Discussion (QUD), a semantic question corresponding to the current discourse topic (Roberts, 1996). Following Simons et al. (2010, p. 316), "an assertion is relevant to a QUD iff it contextually entails a partial or complete answer to the QUD." While *transformations*, on its own, could constitute an assertion relevant to the QUD "what did John enjoy about the fashion show?", it is unclear whether the meaning expressed by (1) asserts anything, as its meaning cannot be challenged or denied by another discourse participant. Rather, it expresses something about the speaker's attitude; this attitude cannot be challenged or denied by others. For example, in (2), B cannot felicitously respond to A by trying to deny the content of their utterance.

(2) A: transformations shmansformations!
B: # That's not true, you do care about transformations

A number of other examples of expressive morphology have been studied elsewhere in the literature. Steriopolo (2009) analyses two types of expressive affix in Russian, including *attitude affixes*, which can express e.g., affection or vulgarity, and *size affixes*, which constitute mixed expressives, expressing that something is a particular size and that the speaker has a particular attitude towards that something. Fortin (2011) provides a similar account of *connotative affixes* in Spanish, in which he analyses the expressive meaning they contribute using Potts's (2005, 2007) multidimensional semantics for expressives. Finally, Meibauer (2013) discusses a number of examples of expressive compounding in German, including *Arschgesicht* ('arse face') and *Reformscheisse* ('reform shit'), in which the same expressive morpheme can have pejorative (*Arschgesicht*) and meliorative (*Arschgut* - 'arse good') connotations.

Another example cited by Zwicky and Pullum (1987) which displays some of these same expressive properties, and which will be the focus of this chapter, is swearing infixation. Swearing infixation, sometimes called expletive infixation, concerns the insertion of a swearword, such as *fucking* or *bloody*, into the middle of another word. The most common example of this is *fan-fucking-tastic* (see, for example, the numerous definitions on Urban Dictionary (2021) for *fan-fucking-tastic*). Like shm-reduplication, swearing infixation has some of the properties of expressive morphology suggested by Zwicky and Pullum (1987). It can apply to words of all grammatical categories, as shown in (3). Swearing infixation also gives rise to a pragmatic meaning. In addition to the social meaning contributed by the presence of a swearword, the manipulation of the word structure by means of swearing infixation is an example of the *playful* or *poetic* language discussed by Zwicky and Pullum (1987).

(3) a. Any-fucking-where
b. Com-fucking-puter
c. Disa-fucking-pear
d. Hi-fucking-larious

Starting from the assumption that the examples in (4) are semantically equivalent, the presence or absence of infixation in swearing could be considered an example of a sociolinguistic variable, by the traditional Labovian definition (Labov, 1972). The only difference between them is the expressive meaning contributed by the infixation that occurs in (4-a), but not in (4-b). The nature of the expressive meaning contributed by swearing infixation is likely to vary due to both the expression being infixed and the context in which it is used. One could imagine, for example, a speaker using (3-d) in one context

to accentuate how hilarious something was, while doing so in another context to add sarcasm to their utterance to suggest that they are, in fact, not finding the present situation hilarious.

(4) a. That's fan-fucking-tastic
b. That's fucking fantastic

The question for this chapter is whether the expressive meaning contributed by the infixation in (4-a) activates social meanings that differentiate it from the minimally different non-infixed construction in (4-b). While I am unable to provide a complete account of the array of potential social meanings that could be indexed by swearing infixing infixation, I hope to discover whether infixed and non-infixed constructions are in some way differentiated in the effect they have on social evaluation, with a view to opening the door to more research that can isolate more context-specific social meanings associated with the former.

A limited amount of previous research has examined the social information that can be contained in a morpheme. Work by Sanchez (2008), following the Labovian principles of quantitative sociolinguistics, identified a social predictor of morphological borrowing in creole Papiamentu in Aruba and Curaçao, namely the speaker's degree of bilingualism. Säily (2011) examined the effect of gender variation on the morphological productivity of the suffixes *-ness* and *-ity* in the British National Corpus. Both suffixes are typically used to derive abstract nouns from adjectives (e.g., *prescriptiveness and prescriptivity*). While the productivity of *-ness* did not differ by gender, *-ity* was shown to be significantly less productive in the writing of women than men. That is, the productivity of the *-ity* morpheme is socially stratified.

Finally, work by Needle and Pierrehumbert (2018) suggests that speakers store socio-indexical information related to speaker gender in their knowledge of different morphemes. Needle and Pierrehumbert first used data from the British National Corpus to find sets of morphemes that were more commonly used by men (e.g., *-cide*, *-ium*), by women (e.g., *-ette*, *-ful*), or by both equally (e.g., *-ance*, *-ist*). Using these groups, they created stimuli consisting of simple real words, complex real words, and complex pseudowords, counterbalanced for the presence of a gendered morpheme and the gendered bias of the whole word; for example, while *-ist* was considered a gender-neutral morpheme, *soloist* was found to have a female bias, while *loyalist* was found to have a male bias.

Needle and Pierrehumbert (2018) presented participants with these stimuli. They were also presented with a male face and a female face above the item; on each trial, they had to indicate whether the male or female face was most likely to use the word in question. For real words, items with a greater whole word bias for female speakers were significantly

more likely to be attributed to the female face. Whole gender bias also interacted with morpheme bias to predict responses, such that "among words containing more woman-biased morphemes, the man-biased whole words were judged to be more man-biased, and the woman-biased words were judged to be more woman-biased" (p. 11). For pseudowords, the gender bias of the morpheme also influenced participants' responses in the same direction. That is, participants used their implicit knowledge to generalize gender associations of pseudowords based on their component morphemes.

Returning to the infixed swearing constructions to be discussed in this chapter, it may be that infixation is a morphological process that also has particular socio-indexical associations. Rather than gendered associations however - which may well exist but are not of interest here - I tested whether swearing infixation was linked to particular socio-pragmatic meanings, pertaining to conversational politeness and sincerity. I will go into greater detail about this dynamic in Section 5.5 and beyond.

5.3 The variable: infixation

The sociolinguistic variable of interest in this chapter is infixation. A number of prominent linguists have attempted to account for expletive infixation in English. The majority of these accounts come from phonology, including Siegel (1974) and Aronoff (1976) (see also McCawley, 1978). Both Siegel and Aronoff provide a rule for expletive infixation, stated below.

Table 5.1: Rule for expletive infixation from Siegel (1974) and Aronoff (1976).

$$\begin{array}{l} \overset{3}{} \quad \overset{1}{} \\ [\ \mathrm{X}\ \mathrm{V}\ \mathrm{Q}\ \mathrm{V}\ \mathrm{Y}\ ] \\ \ 1\ \ 2\ \ 3\ \ 4\ \ 5 \quad \rightarrow 1\ 2\ 3\ \textsc{expletive}\ 4\ 5 \end{array}$$

Condition: Q does not contain $\overset{3}{V}$

The word that hosts the infixation, they suggest, must have the stress pattern *3 1*, that is, the primary stress must be on the second vowel. The rule stipulates that infixation must then appear immediately before the primary stress and be preceded somewhere in the word by a tertiary stress. This accounts for common examples such as *fan-fucking-tastic*. This account was later updated by McCarthy (1982), however, whose own account I will detail here.

McCarthy cites a number of examples that deviate from Siegel and Aronoff's rule. For example, the constructions in (5) are well-formed, but they do not follow the rule; in each case, the infixed swearword does not immediately precede the primary stress of the word; primary stress is indicated here and throughout with capital letters.

(5) a. aMALga-bloody-mated
b. eMANci-motherfuckin-pator
c. HANdi-bloody-cap

To replace the rule provided by Siegel and Aronoff, McCarthy opts for a characterisation that is based in a theory of metrical phonology and which, conveniently, accounts for expletive infixation via the universal conditions of prosodic well-formedness, rather than requiring some additional and elaborate phonological condition. The characterisation assumes the following characteristics of stress: (a) Syllables (denoted by σ) are grouped hierarchically into binary-branching labelled categories called FEET (denoted by Σ) and (b) Feet, which exhaustively partition the syllables of a word, are gathered into a similar word-level metrical structure. In any tree, the leftmost syllable in a foot will bear some degree of stress, with all other syllables in a foot being unstressed. An example structure for *fantastic* is given in Figure 5.1, taken from McCarthy (1982).

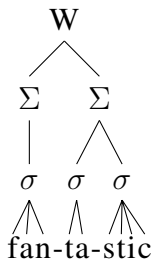

Figure 5.1: Metrical structure tree for *fantastic*

This structure illustrates the fact that *fantastic* consists of three syllables divided between two metrical feet. As the primary stress in *fanTAstic* falls on the second syllable, the second and third syllable belong to the same foot, with the first syllable in a separate foot; this follows from the fact that the leftmost syllable must bear stress. We can contrast this with the structure assumed for *FABulous*, a word with word-initial stress, in Figure 5.2. As the primary stress falls on the first syllable, with tertiary stress on the third syllable, the first and second syllables must belong to the same foot.

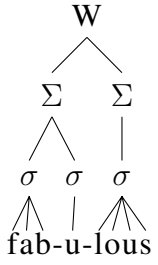

Figure 5.2: Metrical structure tree for *fabulous*

McCarthy makes the generalization that an infix may lodge only at the edge of a

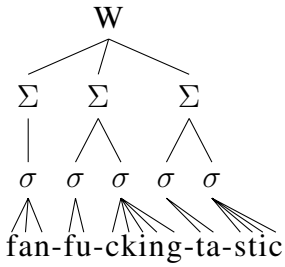

Figure 5.3: Well-formed metrical structure for fan-fucking-tastic

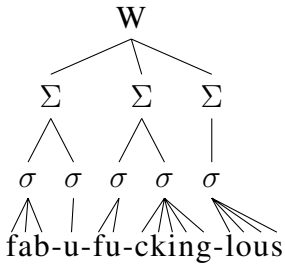

Figure 5.4: Well-formed metrical structure for fabu-fucking-lous

metrical foot. The rule is formalized below. This rule permits the structures in 5.3 and 5.4 for *fan-fucking-TAstic* and *FABu-fucking-lous* respectively, each of which only allow for one site of infixation.

Table 5.2: Rule for expletive infixation from McCarthy (1982).

X [Y]$_\Sigma$	
1 2	$\rightarrow$ 1 EXPLETIVE 2

As the partitioning of syllables into feet is an exhaustive process, all syllables in a word must be dominated by a foot. The well-formedness conditions of English exclude overlapping prosodic domains. McCarthy (1982) therefore suggests that the generalization is better expressed as "an expletive can be inserted in any position not internal to a foot" (p. 579). The generalization correctly rules out the following, ill-formed structures for *fanTA-fucking-stic* (Figure 5.5) and *FAB-fucking-ulous* (Figure 5.6), due to their overlapping domains.

McCarthy (1982) suggests that native speakers have no difficulties in making judgments about the well-formedness of infixed swearing constructions, despite encountering relatively few examples. As yet, the acceptability of different infixed swearing constructions has not been tested experimentally, however, despite the fact that many other factors beyond prosodic well-formedness may contribute to a construction's acceptability (Sprouse, 2018). It is unknown whether the stress pattern of the infixed word affects people's judgments about well-formedness; are stress initial infixes as acceptable as stress-medial infixes? Furthermore, a number of other linguistic factors could also contribute

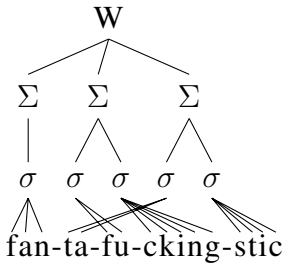

Figure 5.5: Ill-formed metrical structure for fanta-fucking-stic

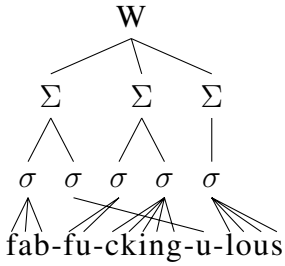

Figure 5.6: Ill-formed metrical structure for fab-fucking-ulous

to the perceived 'acceptability' of an infixed swearing construction, such as lexical frequency, valency and arousal and other morpho-phonological factors (Tessier & Becker, 2018; Reilly et al., 2020). Indeed, despite evidence suggesting little difference between informally and formally collected acceptability judgments (Sprouse & Almeida, 2012; Sprouse, Schütze, & Almeida, 2013), the latter allow for the testing of extra-grammatical factors that might be influencing judgments (Sprouse, 2018).

5.4 Experiment I: Perceiving morphological variation in swearing

The previous section described swearing infixation and its linguistic constraints. I will now turn my attention to examining the extent to which people have an implicit knowledge of these constraints, as well as other potential factors affecting the perceived acceptability of an infixed swearing construction. That is, do people *perceive* a difference between e.g., *fan-fucking-tastic* and *fanta-fucking-stic*?

The reader will note that I am not testing whether people *perceive* the difference between *fan-fucking-tastic* and *fucking fantastic*. That is, unlike in the previous chapter, I am not testing whether people perceive the difference between the two variants of my sociolinguistic variable. This is because the difference between infixed and non-infixed constructions is likely to be more salient to people than the difference between e.g., *fucking* and *fuckin*. As such, an experiment to test whether people *perceive* infixation, compared to a lack of infixation, wouldn't serve much purpose.

A more interesting endeavour is to test whether people are able to generalize the linguistic constraints on swearing infixation proposed by McCarthy (1982) to novel constructions, while also testing whether other linguistic factors may be playing a role in whether or not an infixed construction is considered *acceptable* to a native speaker of English. If people do *perceive* these linguistic constraints, it may also be possible that they *notice* these constraints in the process of social evaluation; that is, it may be that whatever social meanings are associated with infixation are modulated by whether or not the construction is well-formed.

5.4.1 Methods

I tested people's implicit knowledge of the linguistic constraints on swearing infixation using an acceptability judgment task. I reviewed this task in Section 3.3.2 of Chapter 3, but I will briefly repeat the aim of the task here. When we ask someone to judge a sentence's acceptability, we are eliciting a *reported perception of acceptability* (Schütze & Sprouse, 2014). *Acceptability* is a percept that arises spontaneously as a reaction to a string of words that resembles a sentence. The perceived *acceptability* of a sentence is informed by person's previous experience of a sentence and their knowledge of the language in which the sentence is appearing. While we cannot directly access the cognitive system in which this information is stored, we can ask native speakers to report on how acceptable a sentence feels. Using this data, linguists can make inferences in order to build theories of grammar.

As well as providing some experimental verification of the empirical observations underpinning McCarthy's (1982) generalization, Experiment I also helped identify a set of natural-sounding infixed swearing constructions to be used in Experiment II. Although Experiment I was also set-up to test participants' linguistic knowledge of swearing infixation, the principal aim of this chapter, and indeed this thesis, is to test the effect of linguistic variation in swearing on social evaluation of a speaker. To do this, the constructions used in Experiment II needed to sound as natural as possible, despite being constructed examples rather than occurring naturally.

5.4.1.1 Experimental Hypothesis

Experiment I had the following experimental hypothesis. Although it was expected that other factors might also affect the perceived acceptability of the test items, there were no precise predictions about the direction of these effects. As such, no further experimental hypotheses were proposed.

H1 Constructions in which the infix lodges exterior to a metrical foot will be more

likely to be considered acceptable than constructions in which the infix is internal to a metrical foot.

5.4.1.2 Stimuli

Stimuli for this task consisted of 201 adjectives, including 87 stress-initial and 114 stress-medial words taken from an online word finder (Yougowords.com, 2019) (see Appendix A.1). For each adjective, two versions of each construction were created by infixing the adjective with the word *fucking*. In what I will refer to as the *well-formed* version, *fucking* was inserted exterior to a metrical foot. In what I will refer to as the *ill-formed* version, *fucking* was inserted interior to a metrical foot. For example, stress-initial words like 'positive' became 'posi-fucking-tive' (well-formed) and 'pos-fucking-itive' (ill-formed). Stress-medial words like 'indecent' became 'in-fucking-decent' (well-formed) and 'inde-fucking-cent' (ill-formed). The resulting 402 test items were divided equally into two lists, counterbalanced for stress and formedness, such that each list contained exactly one version of each adjective (either well- or ill-formed) and the same number of stress-initial and stress-medial and well- and ill-formed test items. Test items were additionally coded for log lexical frequency from the SUBTLEX UK database (Van Heuven et al., 2014), valency and arousal from Warriner et al. (2013) and number of morphemes (*Nmorph*), concreteness, orthographical neighbourhood density and mean lexical decision accuracy (*LD*) from the English Lexicon Project (Balota et al., 2007); all of these measures were based on the adjective, minus the infix.

Following Tessier and Becker's (2018) finding for vowel harmony in novel swearing constructions - in which constructions such as *fuckpuffin* were rated as more acceptable than constructions such as *fuckbadger* - items were also coded for whether or not their first vowel matched the /ʌ/ vowel in *fucking (First Vowel Match*: No vs Yes). Although Tessier and Becker found no effect for consonant harmony, this does occur in the most common infixed phrase *fan-fucking-tastic*; this was therefore included (*First Consonant Match*: No vs Yes). Phonetic transcriptions were taken from the English Lexicon Project (Balota et al., 2007).

5.4.1.3 Procedure

The task was created and hosted on Gorilla (Anwyl-Irvine et al., 2018). Before starting the task, participants provided brief demographic information (Gender and Age). Participants first completed four practice trials to familiarise them with the procedure. On each trial, participants were presented with one of four versions of the word *delightful*, three of which were infixed with *fucking* in three different positions. Before both the practice and test trials, participants were given the instructions in (6).

(6) *This task is about possible swearwords in English. Sometimes we are creative with the way we swear by adding swearwords to other words to create new words. An example of this is the word 'motherfucker'. However, this isn't necessarily possible with all words. For example, for one reason or another, you might be less likely to think that 'fuckermother', 'mothershitter' or 'moth-fucker-er' are possible swearwords. In this task, you will be presented with a number of words. You will be asked whether this is a possible word in English. You are then required to answer 'yes' or 'no'. When you press continue, you will go into a brief practice version of the task. This will involve performing the task on 4 different words. After this, you will see another short set of instructions before starting the experiment.*

Participants were assigned to one of the two test stimuli lists. Participants then completed 201 test trials, with order randomised between participants. On each test trial, participants were presented with a test item in the centre of the screen. At the top of the screen appeared the question 'Is this a possible word?' and at the bottom of the screen were two buttons labelled 'Yes' and 'No'. Participants were required to use their mouse or touch screen to answer 'Yes' or 'No' on each trial. A break screen appeared every 20 trials to give participants the chance to look away from the screen. Participants also completed a task on the acceptability of swearing compound phrases, the results for which will not be discussed here.

5.4.1.4 Participants

27 British English speakers were recruited and paid £2.20 via Prolific Academic (2022).

5.4.2 Results

The mean values for the four test conditions suggest that the most acceptable test items were the well-formed stress-medial items (50.9% Yes), followed by well-formed stress-initial items (26.8% Yes), ill-formed stress-initial items (10.5%) and ill-formed stress-medial items (9.4% Yes). Figure 5.7 plots the mean values for adjective by Stress and Form.

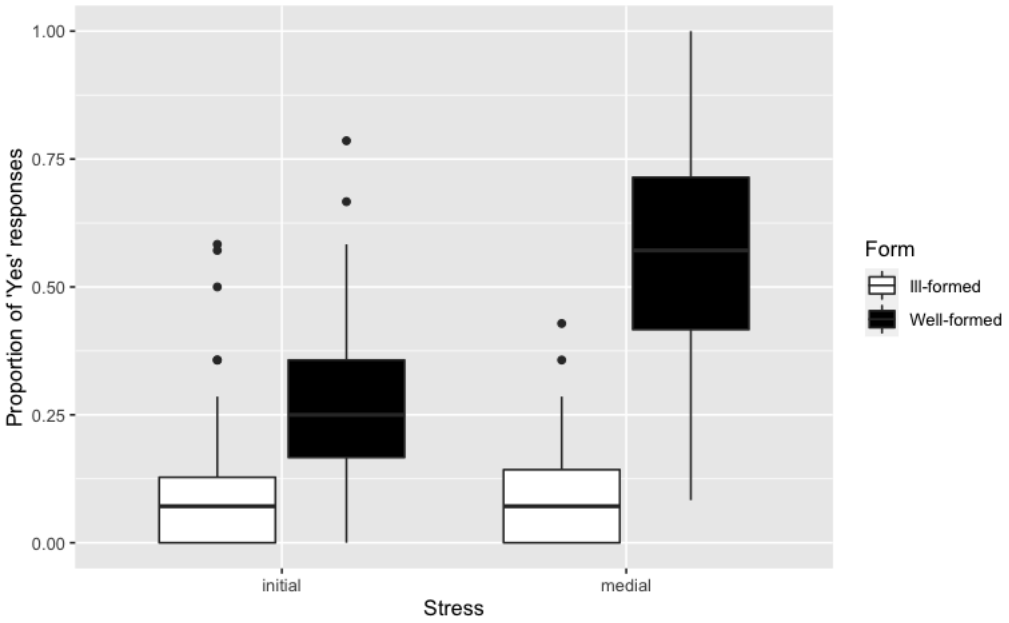

Figure 5.7: Experiment I: Boxplot of mean responses by Stress and Form

Results were analysed in R using a logistic mixed effect regression model using the *lme4* package (Bates et al., 2015), with the binary response variable Yes or No (Yes = 1, No = 0). The full model included the sum-coded categorical predictors Stress (initial, medial) and Form (Well-formed, Ill-formed), as well as their interaction, First Vowel Match, First Consonant Match and the continuous predictors LOG frequency, Valency, Arousal, Number of Morphemes, Mean Lexical Decision Accuracy, Orthographic Neighbourhood Density and Conreteness. Age and Gender of participants were also included as predictors. Participant and Word were included as random intercepts. All categorical predictors were sum-coded.

Variables that did not significantly improve model fit based on chi-square comparisons of the sums of the squares of the residuals were removed. The final model is reported in Table 5.3. As there were not normed measures of valency for all 201 adjectives, the final model includes observations for 142 adjectives.

Table 5.3: Experiment I: Model summary - Response ~ Stress * Form + Valency + Orthographic Neighbourhood Density + First Vowel Match + (1|Participant) + (1|Word)

	Response			
Predictors	*Odds Ratios*	*std. Beta*	*Statistic*	*p*
(Intercept)	0.39 *	-1.45	-4.19	**0.025**
Stress: Medial	4.09 ***	1.41	8.18	**<0.001**
Form: Ill-formed	0.17 ***	-1.80	-10.57	**<0.001**
Word Valency	0.92 *	-0.15	-2.01	**0.044**
Orthographic Neighbourhood Density	0.57 **	-0.23	-2.88	**0.004**
First Vowel Match: Yes	3.61 ***	1.28	4.93	**<0.001**
Medial Stress * Ill-formed	0.14 ***	-1.94	-8.38	**<0.001**
Random Effects				
σ^2	3.29			
$T_{00\ Word}$	0.41			
$T_{00\ ParticipantID}$	2.76			
Observations	3750			
Marginal R^2 / Conditional R^2	0.282 / 0.634			

** p<0.05 ** p<0.01 *** p<0.001*

There were significant main effects for Stress: Medial (β = 1.41, p $<$ 0.001), Form: Ill-formed (β = -1.8, p $<$ 0.001), Word Valency (β = -0.15, p $<$ 0.05), Orthographic Neighbourhood Density ($/beta$ = -0.23, p $<$ 0.01) and First Vowel Match: Yes (β = 1.28, p $<$ 0.001). There was also a significant interaction effect between Stress: Medial and Form: Ill-formed (β = -1.88, p $<$ 0.001).

5.4.3 Interim discussion

The hypothesis for Experiment I is repeated below:

H1 Constructions in which the infix lodges exterior to a metrical foot will be more likely to be considered acceptable than constructions in which the infix is internal to a metrical foot.

Support is provided for this hypothesis. Items in which the infix appeared immediately before a metrical foot were more likely to be considered acceptable than those that did not. This preference was not absolute, however, with ill-formed test items receiving a Yes response on around 10% of trials. The preference for well-formed constructions was significantly stronger for items with word-medial stress, as shown by the significant

interaction effect between Stress and Form. The result suggests that participants did have an implicit knowledge of swearing infixation that reflected McCarthy's (1982) generalization, as outlined in Section 5.3. That is, the results of Experiment I suggest that native speakers are able to generalise the prosodic constraints of expletive infixation to novel constructions.

The results suggest a preference among participants for test items with word-medial stress. As is evident from Figure 5.7, participants were much more consistent in rating well-formed constructions as 'acceptable' when the infix was in a stress-medial word than when it was in a stress-initial word. There was also a significant preference for test items with low scores on normed measures of adjective valency. In essence, the 'sadder' the adjective was, the more likely participants were to consider the whole infixed construction a possible word. While this is not the case for the most commonly infixed adjective *fantastic* - *fan-fucking-tastic* was the only word to get 100% *acceptable* responses - it is consistent with the fact that swearwords, themselves, typically score low for valency (Janschewitz, 2008). Other high-scoring items included *un-fucking-pleasant*, *un-fucking-stable* and *cold-fucking-hearted*. The highest scoring ill-formed test items (*im-fucking-polite*, *trust-fucking-worthy* and *a-fucking-moral*) involved the infix lodging directly between two morphemes; while this was not included in the statistical models, it is possible that this also had an effect on responses.

Finally, there were additional main effects for Orthographic Neighourhood Density and vowel harmony. Regarding Orthographic Neighbourhood Density, this finding suggests the opposite trend to that found in Reilly et al. (2020), as an increase in this variable led to a decrease in the likelihood that the word was found acceptable by participants. For vowel harmony, the result provides further support for the previous finding by Tessier and Becker (2018) that novel swearing constructions are funnier and more satisfying when they contain matching vowels, showing that this effect also improves acceptability judgements. While participants did have an implicit knowledge of the linguistic constraints on swearing infixation, therefore, this was not the only type of knowledge they appeared to have regarding these constructions. Rather, participants' knowledge of the words that were being infixed also played a role.

5.5 Experiment II: Noticing morphological variation in swearing

Experiment I established that native speakers do *perceive* the linguistic constraints on swearing infixation. This suggests that people have some internal representation of these constraints. This representation is below the level of consciousness - it is unlikely that my

participants would have been able to articulate the rule that governs swearing infixation that I detailed in Section 5.3.

In Experiment II I will test whether or not people *notice* swearing infixation, and its constraints, in the process of social evaluation. That is, is swearing infixation linked to particular social meanings and is this modulated by the well-formedness of the infixed construction?

While these types of constructions have not yet been studied with respect to social meaning, a similar construction has received some attention. Constructions such as those in (7), referred to as *expletive insertion*, have been studied by Hoey et al. (2020). Using a conversation analytical approach, Hoey et al. suggest that expletive insertion is frequently used in a series of subsequently positioned actions, in which the speaker of e.g., (7-b) is escalating their attempt resolve some interactional trouble.

(7) a. Who the fuck is that?
b. No really, what the fuck is it?

This is illustrated in a short portion of one of Hoey et al.'s extracts, provided in (8), minus the conversation analysis annotations. The series of events is characterised by the authors in terms of three positions. In position 1, a sequence-initiating action orients to difficulty between participants and provides for its resolution (e.g., *who are you then?*). Position 2 is an uncooperative or inappropriate response (e.g. *Ronnie Pickering*). Position 3 then involves expletive insertion to further pursue a resolution via escalation (*who the fuck's that?*).

(8) A: Come on, who are you then?
B: Ronnie Pickering
A: Who?
B: Ronnie Pickering
A: Who?
B: Ronnie Pickering!
A: Who the fuck's that?
B: Yeah, me!

The work of Hoey et al. (2020) differs from my own in several important ways. Their study is not directly comparing expletive insertion with a lack of expletive insertion. Furthermore, their study does not identify socio-indexical meanings linked with the insertion of a swearword. They do, however, suggest that expletive insertion serves an interactional purpose, which might indicate that it is the source of socio-pragmatic meanings. The same may be true for expletive infixation too.

As suggested in Section 5.2, the exact nature of the social meaning indexed by swearing infixation is likely to vary. I will illustrate this with some examples using *fucking*. Firstly, consider the constructed example in (9). Intuitively, the infixed response in (9-b) seems less sincere than the non-infixed response in (9-a), although this is likely to depend on tone of voice, facial expression etc.

(9) Context: The speaker has just been asked what they thought of the film they just watched.

a. It was fucking hilarious
b. It was hi-fucking-larious

Now consider the constructed example in (10). Although B's responses in (10-a) and (10-b) differ with respect to the presence or absence of infixation, it is not immediately obvious how they would differ with respect to social meaning. Although (10-b) could be considered to be more sarcastic, this would require some additional contextual information (e.g., the box being labelled 'kitchen stuff'). Without that information, however, both (10-a) and (10-b) seem to suggest irritation on behalf of B; the presence or absence of swearing infixation may modulate how 'irritated' the speaker is perceived to be.

(10) Context: A is helping B move house.
A: Where should I put this box?

a. B: Fucking anywhere!
b. B: Any-fucking-where!

Ultimately, as I discuss in Section 5.5.1.3 below, the evaluative dimensions that appear to be particularly relevant to swearing infixation are *sarcasm*, *rudeness*, *funniness* and *happiness*. In Experiment II, therefore, I am limited to drawing conclusions about these four scales when it comes to concluding what people *notice* about swearing infixation in social evaluation. While this type of limitation is a downside to controlled experiments such as matched guise tasks, it does allow me to test for the systematicity of reactions to swearing infixation in a large group of participants and using a large group of constructions without needing to analyse open text responses.

Using the examples discussed above, as well as my own intuitions, I have suggested that infixed swearing constructions should activate different social meanings to minimally different non-infixed constructions. The null hypothesis is that the two types of constructions will not differ with respect to social meaning. As I discussed in Chapter 2, swearing is highly cognitively and socially salient. It may be that, although *fan-fucking-tastic* and *fucking fantastic* differ with respect to infixation, the presence of the swearword *fucking* overrides any other linguistic variation that may be occurring. That is, *fucking* may be so

	Well-formed	**Ill-formed**
Infixed	fan-fucking-tastic	fucked fantastic
Non-infixed	fanta-fucking-stic	fucking fantastic

Table 5.4: Experiment II: Experimental design

salient that it is all that people care about when evaluating a speaker. Due to this possibility, a perception experiment is required to test whether or not people *notice* swearing infixation in social evaluation of a speaker.

5.5.1 Methods

Experiment II used a visual matched guise task (Lambert et al., 1960). Participants were presented with two-line dialogues on screen. In the first line of the dialogue, one speaker provides a piece of information. In the second line of the dialogue, a second speaker reacts to this piece of information. The critical stimuli were contained in the reaction of the second speaker. The critical stimuli followed a 2x2 experimental design. The first two-level factor was Infixation (Yes vs No). The second two-level factor was Form (Well-formed vs Ill-formed). This design is schematized using the example *fantastic* in 5.4.

The well- and ill-formed infixed stimuli were carried forward from Experiment I. The well-formed, non-infixed stimuli consisted of the modifier *fucking* and an adjective as separate words. The ill-formed, non-infixed stimuli were also in this format, but with the modifier *fucking* changed to the adjective *fucked*. The use of these stimuli in full sentences is ungrammatical e.g., **that's fucked fantastic!*. While this is a different type of ill-formedness to that seen in the infixed stimuli, the inclusion of an ungrammatical non-infixed construction allowed me to test whether ill-formed infixed swearing constructions were treated the same as ill-formed non-infixed swearing constructions.

5.5.1.1 Experimental Hypotheses

Experiment II had the following experimental hypotheses:

H1 Infixed swearing constructions will receive different responses across Likert scales compared to non-infixed swearing constructions.

H2 Well-formed infixed swearing constructions will be evaluated less favourably across Likert scales compared to ill-formed infixed swearing constructions.

These hypotheses were admittedly vague. For Hypothesis 1, one might predict that infixed constructions would be evaluated more favourably, due to their poetic and playful nature (Zwicky & Pullum, 1987). On the other hand, this form of word play could appear more marked than the non-infixed equivalent (e.g., *fucking fantastic*). If swearing is

penalised more generally, then it may be that this is exacerbated when swearing appears more marked or deliberate. In any case, it was expected that reactions to infixed and non-infixed constructions would differ. The null hypothesis was that there would be no significant difference in responses to the two types of constructions.

For Hypothesis 2, it was expected that well-formed and ill-formed infixing constructions would also differ in the responses they received. This was ultimately a question of whether or not people are penalized for doing wordplay with swearing without obeying the well-formedness conditions of that wordplay. In simpler terms, does it matter whether people swear correctly? Unlike for Hypothesis 1, there was a clearer expectation about the direction of this effect; if well-formedness matters, it makes sense that ill-formed constructions would be perceived less favourably, as this would indicate that the speaker doesn't have a full mastery of English grammar due to not understanding prosodic well-formedness constraints.

5.5.1.2 Stimuli

The results of Experiment I informed the stimuli selection for Experiment II; it was important that the infixed stimuli sounded as natural as possible. The well-formed versions of all 201 test adjectives were ranked from most acceptable to least acceptable based on the mean response they received from participants in Experiment I. Initially, the top 20 words were selected. All of these items were 'acceptable' over 70% of the time. A large proportion of these words began with the prefix 'un-'. To create a more diverse list of infixed constructions, 4 of these '-un' words were removed. This resulted in 16 test adjectives. The only stress-initial word that was included was *fabulous*.

For each adjective, 4 items were created following the experimental design, that is, one well-formed infixed (e.g. *fan-fucking-tastic*), one ill-formed infixed (e.g. *fanta-fucking-stic*), one well-formed non-infixed (e.g. *fucking fantastic*) and one ill-formed non-infixed (e.g. *fucked fantastic*). These items were then embedded in a declarative statement of the form 'That's X', where X was a test item. These statements were then put into 2-line dialogues. The first line of each dialogue was a declarative statement made by Person A, with the second line of the dialogue containing the statement with the test item. The first line set up a context in which the test adjective would be a natural-sounding response. An example dialogue, in the well-formed non-infixed condition, is included in (11), with the full list of dialogues and adjectives included in Appendix A.2.

(11) Person A: I've been sat watching TV all week
Person B: That's fucking unhealthy!

With all combinations of the 2x2 research design exhausted, a total of 64 test dialogues

were constructed. These were divided into 4 lists of 16 dialogues, counterbalanced for condition and test adjective, such that each list contained exactly 1 occurrence of each test adjective and exactly 4 stimuli in each of the 4 experimental conditions. In addition, these lists contained 20 distractor dialogues. Distractors followed the same format as test dialogues, containing 2 declarative statements, the second of which was of the form 'That's X'. For distractors, X was a mixture of swearing (e.g. *shit*, *crap*) and non-swearing adjectives (e.g. *amazing*, *terrible*). One of the 4 lists is included in Appendix A.3.

5.5.1.3 Selecting the scales

The Likert scales for Experiment II were selected based on the results of a separate task. In this task, 100 participants were asked to provide four adjectives that they would associate with a person who uttered a particular phrase. Participants were asked to evaluate the two minimally different constructions in (12) and (13). Participants were presented with the phrases one at a time with order randomised between participants. The task was aimed at eliciting explicit differences in perception between the two constructions, hence why all participants were shown both utterances. This process was motivated by the need for scales to reflect the social meanings relevant to speakers (see Section 3.3.3 of Chapter 3).

(12) That's fucking fantastic

(13) That's fan-fucking-tastic

A large number of unique words were provided for each utterance. The non-infixed construction (12) elicited 226 words, while the infixed construction (13) elicited 245. The tables in 5.8 report the most frequent words; for brevity, a minimum of 5 mentions was used as a cut-off point.

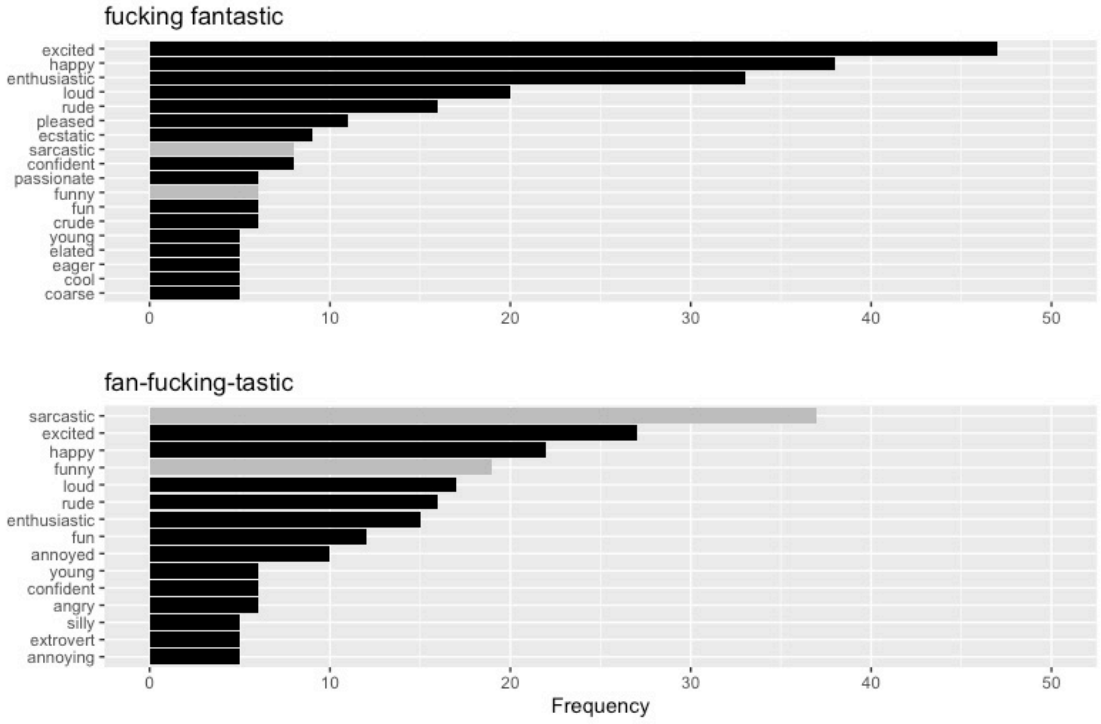

Figure 5.8: Experiment II: Words associated with 'That's fucking fantastic' and 'That's fan-fucking-tastic'

As Figure 5.8 shows, the adjectives *funny* and *sarcastic* (highlighted in grey) had very different frequencies for the two constructions. The majority of the other adjectives had broadly similar associations with the two constructions. For the experiment, scales for *funny*, *sarcastic*, *happy* and *rude* were chosen, with their antonyms *unfunny*, *sincere*, *sad* and *polite* chosen for the other ends of those scales. *Happy* was chosen to cover both *happy* and *excited*. *Rude* was chosen partly for its high frequency in the task, but also because it is an adjective that is frequently associated with swearing more generally due to its breaking of societal taboos and offensiveness. A scale for *rude* would explore whether the use of infixed swearing could directly mitigate or exacerbate the offensiveness typically associated with swearwords.

5.5.1.4 Procedure

After giving informed consent, participants were given the instructions in (14).

(14) *Hi there! Welcome to the experiment. In this task, you're going to be presented with a series of 2-line dialogues between two people: Person A and Person B. In each, Person A is going to make a statement and Person B is going to respond. For each dialogue, we want you to consider the response given by Person B. You will be presented with a set of scales from 1-6 depicting characteristics (e.g. cool-uncool, happy-sad). We want you to draw an inference about Person B,*

> *based on their response, using these scales. For example, if you thought that their response made them sound cool, you would choose a number closer to the label that said 'cool'. We will start with a couple of practice examples so that you can get used to the scales. Click the 'continue' button below to start the practice trials.*

Once participants had read the instructions, they moved on to a set of practice trials. Each participant completed the same two practice trials. On each practice trial, the participant was presented with a dialogue in the centre-top of the screen. Below the dialogue were 4 Likert scales depicting the characteristics chosen in Section 5.5.1.3: *funny*, *sarcastic*, *happy* and *rude*. The direction of the scales was kept constant across participants (e.g., *funny* was always at the same end of the scale). The presentation order of the scales was randomised between participants. Participants were required to input a response for all four scales before they could move to the next trial. The two practice trials contained dialogues unrelated to the test stimuli.

Once participants completed the two practice trials, they were given another short set of instructions to remind them of the procedure before moving on to the test trials. The procedure for the test trials matched the procedure for the practice trials. Each participant was assigned to one of the 4 test lists. As such, each participant completed a total of 16 test trials and 20 distractor trials, with order randomised between participants. At one random point during the test trials, participants completed a catch trial to ensure that they were paying attention to the task. Instead of a dialogue, the catch trial presented participants with an instruction to select 6 on all of the scales. Any participant that failed this catch trial would be removed from the analysis. In addition, a break screen appeared after every 12th trial, instructing participants to look away from the screen for 10 seconds before they could continue with the task.

After completing all of the test trials, participants completed a survey on swearing consisting of the following three questions. Their responses were coded as Swear Frequency, Infix Frequency and Swear Appropriacy respectively.

1. How frequently do you swear?

2. How frequently do you use infixed swearwords? e.g. fan-fucking-tastic

3. How often do you think swearing is appropriate?

For each question, participants were presented with a 5-point Likert scale from *never* to *always*. Participants then completed a demographic survey eliciting information on their age, gender, level of education and yearly household income (in blocks of 10,000).

5.5.1.5 Participants

139 participants were recruited and paid £1.05 via Prolific Academic (2022).

5.5.2 Results

Figure 5.9 plots the mean responses for each adjective in each condition.

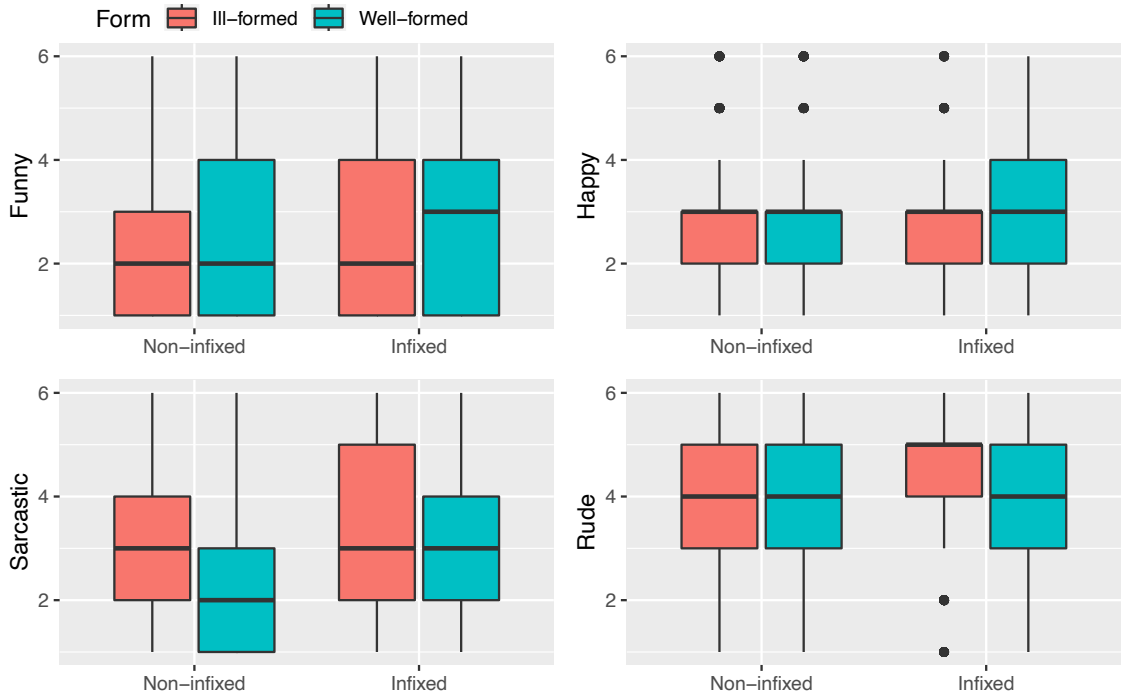

Figure 5.9: Experiment II: Boxplot - Scale response Infixation * Form

Results were analysed in R using a linear mixed effect regression model using the *lme4* package (Bates et al., 2015). Four separate models were created for each response scale, with *funny*, *happy*, *sarcastic* and *rude* the respective dependent variables (1-6). The full models included the sum-coded categorical predictors Infixed (No, Yes), Form (Well-formed, Ill-formed), as well as their interaction, Participant Gender and Participant Level of Education. The continuous predictors were Participant Age, Participant Income, Swear Frequency, Infix Frequency, Swear Appropriacy Rating, and Adjective log Frequency from the SUBTLEX UK database (Van Heuven et al., 2014). Variables that did not significantly improve model fit based on chi-square comparisons of the sums of the squares of the residuals were removed. Where variables outside of the experimental manipulations (i.e. Infixation and Form) improved model fit, they were also tested in interactions with the experimental manipulations. Random intercepts were included for Participant and Adjective. A random slope was originally included for Participant over Adjective, but this was dropped due to issues with model convergence.

The final models for the scales *sarcastic*, *funny* and *rude* are presented in Table 5.5. There were no significant effects for the scale *happy*, so this model is not presented. For *sarcastic*, there were significant main effects for Infixed: Yes (B = 0.58, $p < 0.001$), Formedness: Ill-formed (B = 0.25, $p < 0.001$) and Participant Age (B = -0.01, p = 0.028). For *funny*, there were significant main effects for Infixed: Yes (B = 0.2, $p < 0.001$), Formedness: Ill-formed (B = -0.14, p = 0.01), Participant Age (B = -0.2, p = 0.001), Adjective Log Freq (B = 0.25, p = 0.003) and Participant Income (B = 0.25, $p < 0.001$). For *rude*, there were significant main effects for Infixed: Yes (B = 0.18, $p < 0.001$), Formedness: Ill-formed (B = 0.18, $p < 0.001$) and Infix Freq (B = -0.27, p = 0.003).

Table 5.5: Experiment II: Model summary - Response ∼ Infixed + Formedness + Participant Age + Adjective Log Freq + Participant Income + Infix Freq + (1|Participant) + (1|Adjective)

	Sarcastic			**Funny**			**Rude**		
Predictors	*Estimates*	*Statistic*	*p*	*Estimates*	*Statistic*	*p*	*Estimates*	*Statistic*	*p*
(Intercept)	3.16	6.81	**<0.001**	1.59	3.47	**0.001**	4.62	9.66	**<0.001**
Infixed: Yes	0.29	10.62	**<0.001**	0.10	3.80	**<0.001**	0.09	3.90	**<0.001**
Formedness: Ill-formed	0.12	4.54	**<0.001**	-0.07	-2.58	**0.010**	0.09	3.97	**<0.001**
Participant Age	-0.02	-2.53	**0.011**	-0.02	-2.69	**0.007**	0.01	1.08	0.278
Adjective LOG Freq	0.12	1.34	0.180	0.25	2.96	**0.003**	-0.12	-1.51	0.131
Participant Income	0.06	1.89	0.059	0.07	2.51	**0.012**	0.05	1.63	0.104
Infix Frequency	-0.14	-1.64	0.100	0.23	2.68	**0.007**	-0.25	-2.71	**0.007**
N	120 $_{\text{ProlificID}}$			120 $_{\text{ProlificID}}$			120 $_{\text{ProlificID}}$		
	16 $_{\text{adjective}}$			16 $_{\text{adjective}}$			16 $_{\text{adjective}}$		
Observations	1880			1874			1864		

5.5.3 Interim discussion

The results of Experiment II can be summarised as follows. Across participants and adjectives, the 'speakers' were considered significantly funnier, ruder and more sarcastic when they used an infixed construction compared to when they used a non-infixed construction. There were several additional main effects that suggested that, across all four experimental conditions, participant age and income, participants' own use of swearing infixation and the log frequency of the adjective all influenced responses, predominantly on the *funny* scale. While these effects are interesting, none of these predictors interacted with the main experimental manipulations. As a result, they will not be discussed further.

The experimental hypotheses for Experiment II are repeated below:

H1 Infixed swearing constructions will receive different responses across Likert scales compared to non-infixed swearing constructions.

H2 Well-formed infixed swearing constructions will be evaluated less favourably across Likert scales compared to ill-formed infixed swearing constructions.

The result for the predictor Infixed provides some support for the first hypothesis. On the one hand, the presence of infixation had no effect on the scale of *happy*/*sad*. This may have been an artefact of the scale selection task detailed in Section 5.5.1.3, in which participants only saw the constructions *fan-fucking-tastic* and *fucking fantastic*; it may be that the scale *happy*/*sad* is not as relevant to swearing infixation as was originally predicted. Furthermore, many of the adjectives used in Experiment II could be thought of as either *happy* (e.g., *fantastic*, *fabulous*) or *sad* (e.g., *unpleasant*, *indecent*); participants may have chosen their responses based on the valence of the adjective itself, meaning that the presence or absence of infixation was irrelevant.

The presence of infixation did have a significant effect on responses for the other three scales, however. Significant main effects for the predictor Infixed were found on the other three scales, such that infixed constructions were considered more *sarcastic*, *funny* and *rude* than non-infixed constructions. While it perhaps makes intuitive sense that *sarcastic* and *funny* showed the same trend in the data, there is no obvious reason why *rude* showed the same trend as *funny*. To examine this further, I conducted further exploratory analysis on the data.

Principal Component Analysis (PCA) was carried out using the *psych* package (Revelle, 2021) in R. PCA is used to check for the collinearity of dependent variables. Using a loading threshold of 0.5, this analysis suggested that the three scales loaded onto two principal components. This is illustrated in Figure 5.10, with the first component along the y axis and the second component along the x axis. The first component suggests a significant positive correlation between *rude* (0.68) and *sarcastic* (0.87). The second component suggests inverse correlation between *rude* (-0.61) and *funny* (0.84).

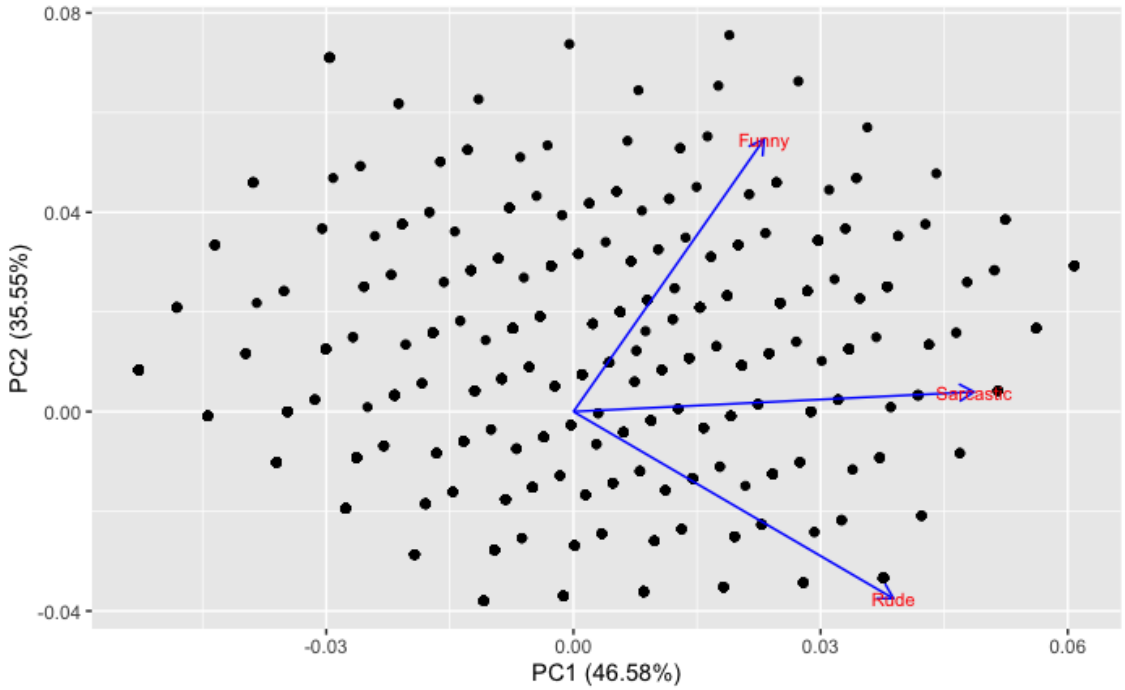

Figure 5.10: Experiment II: Principal Component Analysis: plotted using *autoplot* in the *ggfortify* package

The PCA suggests that there were two (potentially overlapping) groups of participants in the data. For one group, their perceptions of sarcasm and rudeness co-varied; that is, if they perceived a 'speaker' to be sarcastic, they also perceived them to be rude. For a second group, their perceptions of rudeness and funniness were inversely correlated; that is, if they perceived a 'speaker' to be funny, they perceived them to be polite. While the statistical models suggest that, on the aggregate across responses and with participant and word factored in using random intercepts, infixed constructions received higher ratings on all three scales, this follow-up analysis suggests that the results are more complicated than this.

In light of this, the effects for *sarcastic* and *rude*, on the one hand, and *funny*, on the other, require different explanations. Regarding the first effect, I hypothesized in Section 5.5.1.1 that, given that swearing is typically penalized more generally for being rude, it may be considered even ruder if swearing is perceived to be more deliberate. While uttering the construction *fucking fantastic* could be explained away as a slip of the tongue, a speaker who uses an infixed construction has deliberately integrated a swearword into another word in a sentence. Such deliberate swearing might have been perceived as rude by a significant portion of participants for a significant portion of items. The co-variation of *sarcastic* with *rude* suggests that *sarcastic* may not have been considered a positive trait.

The effect for *funny* requires a different explanation. I hypothesized that another pos-

sible reaction to an infixed swearing construction was to consider it poetic and playful, per Zwicky and Pullum's (1987) characterisation of expressive morphology. To take this one step further, it is possible that the structural markedness of swearing infixation is the underlying factor in such constructions being perceived as funny. To explain this further, I will reflect on work by Dingemanse and Thompson (2020) on *playful iconicity*.

Dingemanse and Thompson (2020) use large datasets with lexical ratings to test whether perceived *iconicity* - i.e., the degree to which words sound like what they mean - and perceived *funniness*, covary. For example, words like *zigzag* and *flop* are frequently rated high in iconicity. Using an algorithm trained on a large corpus of natural language text, the authors estimated funniness and iconicity ratings for a total of 70202 English words. They also tested for other lexical characteristics that might predict perceived funniness. Using the estimated values, the authors first found that estimated iconicity is the strongest predictor of perceived funniness in a model that also included log lexical frequency and mean lexical decision times.

Another important factor in their data was structural markedness. Log letter frequency, used as a proxy for markedness, emerged as a strong predictor of perceived funniness, following previous findings (Westbury & Hollis, 2019). When conducting a linguistic analysis of the combined upper ten percentiles of iconicity and funniness ratings (N = 80), Dingemanse and Thompson (2020) identified three recurring cues of structural markedness: complex onsets (e.g., *flap*, *drizzle*), complex codas (e.g., *oink*, *whirl*) and the expressive verbal diminutive suffix '-le' (e.g., *tingle*, *wobble*); a combined measure of these three cues was an even stronger predictor of perceived funniness than log letter frequency. The authors suggest that their findings indicate that "structural markedness can function as a metacommunicative cue inviting playful and performative interpretations" (p. 218).

This may be the dynamic at play in Experiment II of this chapter. Much like the diminutive suffix '-le', infixation conveys something expressive, rather than descriptive. Furthermore, as infixation is not used anywhere else in English, it would be fair to say that it is a less frequent structure than the non-infixed constructions used in Experiment II e.g., *fucking fantastic*. This would suggest that an infixed swearing construction is more structurally marked than a non-infixed swearing construction. If so, the results of Dingemanse and Thompson (2020) would suggest that this increased structural markedness could be explaining the increase in perceived funniness for infixed swearing constructions in this experiment (see also Rastall, 2004 on the playful use of reduplication).

Turning to Hypothesis 2, the findings of Experiment II provide supporting evidence for this hypothesis. Firstly, interaction effects between Formedness and Infixed did not improve model fit for any of the scales, meaning any effects for Formedness relate to all test stimuli, not just the infixed items. For both infixed and non-infixed constructions, ill-formed items were rated as significantly ruder, more sarcastic and less funny than

well-formed constructions. Regarding infixation, this would suggest that the increase in perceive funniness that using an infixed construction gets you is conditional on that infix being well-formed; if it is ill-formed, the benefit of using infixation is minimal. Similarly, while people are penalised in terms of rudeness (and sarcasm, assuming that sarcasm is a negative trait based on their covariance) for an infixed swearing construction, this penalty is lessened if they at least use a well-formed construction.

Despite the lack of significant interaction effects in any of the models, however, the plot in Figure 5.9 suggests that there may be some interaction between Formedness and Infixed that has not quite reached the threshold for statistical significance. As you can see, in each plot, the difference between the ill-formed and well-formed items appears to be different in each of the Infixed conditions. For *funny*, while the non-infixed constructions are roughly similar, the well-formed infixed construction appears to score marginally higher than the ill-formed infixed construction; the same is true for the *happy* scale. The reverse is true for the *rude* scale; the non-infixed constructions again appear to score similar, but this time it is the ill-formed infixed constructions that appear to score higher than the well-formed infixed constructions. Finally, on the *sarcastic* scale, while the infixed constructions appear to have scored similar, the difference between well- and ill-formed non-infixed constructions appears to be more pronounced, with the ill-formed non-infixed constructions perceived to be more sarcastic than their well-formed counterparts. Future modelling of this data might therefore reveal some more interesting trends.

5.6 Chapter discussion and conclusions

Returning once again to the overall aims of this thesis - to examine whether people are sensitive to language-internal factors in the social evaluation of swearing - the findings of this chapter provide significant insight from the perspective of morphology. I have shown that infixation influences social evaluation of the speaker as a function of their swearing in relation to the traits *sarcastic*, *rude* and *funny*. Notably, unlike the previous chapter which dealt with social meanings that were predominantly representative of a person's social identity (e.g., *intelligent*, *working-class* etc), these social meanings are really socio-pragmatic meanings; the chosen scales arose during a separate task aimed at objectively generating the relevant evaluative dimensions. While it is possible that swearing infixation also affects perceptions of the types of social traits that were tested in Chapter 4, this was not tested here. Rather, this chapter presents evidence that swearing infixation affects the emergence of social meanings that are specifically relevant to how a person is behaving in discourse.

As well as providing evidence of the social significance of swearing infixation, both Experiments I and II provide an insight into the level of awareness that people have of

the linguistic constraints on swearing infixation. Following the generalization posed by McCarthy (1982) for expletive infixation, I provided evidence via Experiment I that native speakers of English can generalize these rules to novel constructions which they are unlikely to have heard in an infixed state before. This suggests that people have an implicit knowledge of the prosodic constraints on swearing infixation. The results furthermore indicate that a number of other linguistic factors influence the perceived 'acceptability' of an infixed construction. I then provided evidence in Experiment II that people's implicit knowledge of the linguistic constraints on swearing infixation can influence the conscious process of social evaluation. Just as people are penalised for disfluencies more generally, the same is true for swearing. In essence, if you are going to swear and get some kind of positive feedback as a result (e.g., be perceived as funny), you have to do so *correctly*.

In Chapter 7, I will expand further on what this result means for swearing research and for the overarching aims of this thesis, as well as suggesting other forms of expressive morphology that could be a potential source of social meaning. Regarding morphological variation more broadly, the findings of this chapter suggest that expressive morphology may be a reasonable exception to the notion that morphological variation does not fit the mould of sociolinguistic variation in the Labovian sense (Labov, 1972). With meaning kept constant (with the exception of not-at-issue expressive meaning), variation in expressive morphology concerns only a change in form. The findings of this chapter therefore open the door to more work in this vein.

Chapter 6

Swearing and the socio-semantic variable

6.1 Introduction

I will now move on to the third and final domain of language in which I expect variation in swearing to have consequences for social evaluation of the speaker, namely semantics and pragmatics. As discussed in Section 2.3.3 of Chapter 2, swearwords, as expressives, have interesting semantic properties that set them apart from neutral words as a homogenous group (Potts, 2007). Within the category of expressives, however, there is heterogeneity in the meaning contribution they make to a sentence (Kratzer, 1999; McCready, 2010; Gutzmann, 2011). While swearwords can be semantically circumscribed as a discrete category, not all uses of the same swearword will be intended to mean, or be interpreted to mean, the same thing. Previous work on intensification - the main function of swearwords (McEnery, 2004; Love, 2017) - has shown that the same form can have different meanings, both semantic and social, depending on the linguistic environment in which it is used (Beltrama & Staum Casasanto, 2017). In a similar vein, this chapter focuses on semantic/pragmatic variation in the linguistic environment of swearwords that may have consequences for how they affect social evaluation.

The rest of this chapter is structured as follows. Section 6.2 introduces the concept of socio-semantic variation. In Section 6.3, I introduce the two variables, namely animacy and gradability. In line with the methodological logic of this thesis, I then present two consecutive, related experiments that examine different aspects of these two variables. In Section 6.4, I detail a self-paced reading task that tests whether people *perceive* the two variables with respect to swearing. If people have an internal representation of these variables, we might expect this to be reflected in an automatic process like reading. In Section 6.5, I detail a visual matched-guise task that tests whether people *notice* the variables in

the social evaluation of swearing. Finally, in Section 6.6, I discuss the findings of this chapter and give my conclusions.

6.2 Socio-semantic variation

The previous two chapters dealt with a contrast between forms that were semantically equivalent. There is no clear difference in reference expressed by *fucking* and *fuckin* (although see Acton, 2020 on *fishing* and *fishin*) or by *fan-fucking-tastic* and *fucking fantastic*. This type of contrast is typical of work in the Labovian tradition which has often been dominated by comparing two ways of saying "the same thing" (Labov, 1972, p. 271). With truth-conditional meaning kept constant, it is possible to directly compare how two linguistic variables differ in terms of social significance (Weiner & Labov, 1983). This approach to the linguistic variable underpins much of the work in sociolinguistics at the levels of phonetics and phonology.

Not all sociolinguistic variation involves forms that are strictly semantically equivalent, however. Variation above the level of sound is likely to express a difference in meaning, with the exception of purely expressive forms of morphology (see Chapter 5). There is plenty of evidence of structural variation being socially stratified, however (e.g., (Cheshire, 2005; Moore & Podesva, 2009), among others). It has been argued, therefore, that a more appropriate approach to the sociolinguistic variable involves 'functional equivalence' (Sankoff, 1973; Lavandera, 1978; Dines, 1980; Romaine, 1980; Cheshire, 1987), or two ways of "doing the same thing" (Chambers & Trudgill, 1998, p. 50). This loosening of the definition of a sociolinguistic variable has allowed for its extension to variation in language in which there is a change of reference, for example syntactic or semantic variables, but which have the same discourse function (Campbell-Kibler, 2010c).

As Cheshire (2005) makes clear in her study of adolescent use of different noun phrase constructions in London English, syntactic variation differs from phonological variation, in that speakers regularly use syntactic variation to construct discourse. While studies of phonological variation have typically relied on a standard vs non-standard distinction between two variants (e.g., variable (ING)), which lends itself to large multivariate analyses, Cheshire argues for a greater incorporation of the tools of pragmatics in understanding the situated social functions of structured variation.

A relaxation of Labov's (1972) definition of the sociolinguistic variable has also been argued for by linguists seeking to incorporate semantics into studies of sociolinguistic variation. While a semantic contrast would appear to directly contradict the desire for semantic equivalence, the notion of 'functional equivalence' (Lavandera, 1978) allows for the boundaries between semantic and social meaning to become blurred. Increasing attention is now being paid to the role of variation in semantic interpretation as a source

for sociolinguistic meaning-making (Eckert, 2019; Beltrama, 2020; Acton, 2020).

A prime example of work in this area is Beltrama and Staum Casasanto (2017), whose study on the intensifier *totally* exemplifies the utility of considering the functions of semantic variation in dynamically constructing social meaning. While this study was briefly discussed in Chapter 3 regarding its use of a matched guise task, the variable of interest bears further discussion here. The key contrast in their paper is between sentences like (1-a) and (1-b). The two uses of *totally* are functionally equivalent, as they both function as intensifiers.

(1) a. John is totally bald
    b. John is totally tall

The difference between (1-a) and (1-b) is in how the intensifier is integrated with the rest of the sentence. In (1-a), *totally* specifies that its argument, the adjective *bald*, holds of the individual, *John*, to the maximum degree (Kennedy & McNally, 2005). This canonical use of *totally* is referred to by Beltrama and Staum Casasanto as the *lexical* use. This contrasts with the use of *totally* in (1-b). Here, *totally* cannot combine with *tall* in the same way it did with *bald*, because *tall* is associated with an unbounded scale. While it is possible to reach a maximal degree of baldness, there exists no maximal degree of height.

In (1-b), Beltrama and Staum Casasanto suggest that *totally* "targets a higher-level scalar attitude" (p. 160), which in turn signals the speaker's maximal commitment toward the proposition i.e., that John is tall; this is referred to as the *speaker-oriented* use of *totally* (see also Beltrama, 2018). As Beltrama and Staum Casasanto (2017) show experimentally, the speaker-orientated use of *totally* shows a socio-indexical association with higher degrees of solidarity (e.g., *friendly*, *outgoing*) and lower degrees of social status (e.g., *intelligent*, *mature*) compared to the lexical use. The authors link this association with the role of interactional alignment/convergence, a factor not at play in the lexical usage.

There is an important difference in the comparisons made between two sociolinguistic variants at the levels of phonetics and phonology, such as variable (ING), and at the level of semantics. In phonetics and phonology, the typical contrast is between two different forms in the same environment. For example, one might directly compare responses to different phonetic realisations of a word in the same sentence. This is not always the case (see e.g., Callier, 2013 on the sentence position of creaky voice), but is broadly typical in sociophonetics. At the level of semantics, however, the contrast can involve the same form (e.g., *totally*) in different environments (e.g. before either bounded or unbounded adjectives). In essence, it is the linguistic environment itself that becomes the sociolinguistic variable.

This has also been shown in socially meaningful syntactic variation, such as in Bender's

(2005) work on copula absence (see Section 3.4.1 of Chapter 3 for a fuller review). In Bender (2005), pre-nominal copula-absence was shown to be more marked than pre-verbal copula-absence and, as a result, showed a more extreme association with social meanings linked to copula absence more generally. In Beltrama and Staum Casasanto (2017), too, the more marked use of *totally*, the speaker-orientated use, showed a greater association with a particular set of traits. In sum, the same forms can be more or less marked in a particular linguistic environment and, as a result, become a stronger or weaker indicator of social meaning.

While there has been plenty of research on the formal semantics of expressives (Kaplan, 1999; Kratzer, 1999; Potts, 2005, 2007; McCready, 2010; Gutzmann, 2011, 2015), including a small amount of experimental work (Frazier et al., 2015; Donahoo, 2019; Donahoo & Lai, 2020; Gutzmann, 2019), no previous work has examined the socio-semantics of swearwords or any other form of expressive. That is, no previous study has measured the relationship between semantic/pragmatic variation in expressivity and socially meaningful demographic categories, traits or personae, either in production or perception. The experiments discussed in this chapter are, therefore, the first to draw on the possible relationship between swearwords and socially meaningful semantic/pragmatic variation.

In this chapter, I introduce two forms of linguistic variation in swearing which are strong candidates for socio-semantic variables, namely *animacy* and *gradability*. Each variable follows the blueprint suggested above. In the key constructions containing the variable, the form will stay the same. That is, the constructions will always contain a swearword. This form will perform the same function in each construction, namely intensification. The aspect of the construction that will vary will be the linguistic environment in which that swearword is situated. For the *animacy* variable, the animacy properties of the noun immediately following the swearword will vary. For the *gradability* variable, the gradability properties of the adjective immediately following the swearword will vary. In both cases, it is expected that Experiment I will reveal whether or not people have an implicit knowledge of these variables, that is, whether or not they *perceive* them. Should this be the case, Experiment II will reveal whether or not this implicit knowledge affects the social evaluation of a speaker, that is, whether or not they *notice* the variables. First, however, I will build a case for each distinction using observations from the literature on semantics and pragmatics, as well as original empirical evidence of my own in the form of judgments.

6.3 The variables

Unlike in the two previous chapters, this chapter employs two different variables to test perceptions of the same phenomenon, namely the distinction between local and non-local

swearing intensification. The reader will recall from Section 2.3.3 of Chapter 2 that a property of expressives is flexibility of interpretation (Potts, 2005, 2007). This property concerns the fact that expressive modifiers, such as the swearwords *fucking*, *damn* and *bloody* in English, can operate over multiple different constituents of a sentence. Descriptive modifiers, on the other hand, can only operate over their syntactic sisters. This is best encapsulated in minimal pair sentences like those in (2). *Black* in (2-a) can only combine with *cat*, to give the meaning of *black cat*. In (2-b), however, it intuitively seems that *damn* can operate over *cat* or *the cat is in the garden*, either expressing heightened emotion about an individual or about the state of affairs described by the whole sentence, namely the fact that the cat is in the garden. I will refer to cases where a swearing modifier is interpreted as operating over its syntactic sister as examples of *local swearing intensification* or the *local reading/interpretation* of the modifier. I will refer to cases where a swearing modifier is interpreted as operating over the whole sentence as examples of *non-local swearing intensification* or the *non-local reading/interpretation* of the modifier; this is also known elsewhere in the literature as the *global* reading.

(2) a. The black cat is in the garden
b. The damn cat is in the garden

Whether *damn* gets a local or non-local reading in (2-b) is up to the interpretation of the hearer. They are likely to use other forms of information at their disposal in order to inform this interpretation. Perhaps the hearer has pre-existing knowledge that the speaker despises the cat in question, making the local reading more likely. Conversely, maybe the hearer has pre-existing knowledge that the speaker adores their cat and would be unlikely to blame it for being in the garden, making the non-local reading more likely. There may even be other aspects of the linguistic stimulus that influence interpretation. If the speaker puts stress on the word *cat*, then perhaps the local reading would again be more likely.

Relatedly, what I am proposing in this chapter is that certain properties of the syntactic sister of a swearing modifier can affect the interpretation that the modifier receives. My basic claim in this section, and indeed in this chapter, will be that the animacy or gradability properties of the word immediately following an expressive modifier will affect the availability and/or likelihood of a local or non-local reading. I will claim, following Frazier et al. (2015), that when a swearing modifier is followed by an animate noun, a local reading is more likely than when it is followed by an inanimate noun; this has already been demonstrated experimentally by Frazier et al. (2015). Similarly, I will claim that when a swearing modifier is followed by a gradable predicate, both a local reading and a non-local reading are available, whereas only a non-local reading is available when it is followed by a non-gradable predicate; this is a novel observation that I will evidence using linguistic judgments and the existing literature on both gradability and intensification.

In the following two subsections, I will describe the concepts of animacy and gradability in greater detail. I will also use these subsections as the basis for the categorical distinctions I am drawing between the two variants of two sociolinguistic variables. As discussed in the previous section, a socio-semantic variable can involve the use of the same form in different linguistic environments (e.g., *totally* in Beltrama & Staum Casasanto, 2017). This is what I am proposing for animacy and gradability. The change in the linguistic environment, i.e., alternating between animate and inanimate nouns or gradable and non-gradable adjectives, is the sociolinguistic variable.

6.3.1 Variable I: Animacy

The first linguistic variable to be tested in this chapter is *animacy*. My claim concerns examples like the minimally different sentences in (3). In both examples, the expressive *fucking* appears prejacent to a noun. The important difference is the animacy properties of the nouns involved, with *fucking* appearing before the animate noun *kids* in (3-a) and the inanimate noun *paint* in (3-b).

(3) a. The fucking kids stained the carpet
b. The fucking paint stained the carpet

Frazier et al. (2015) suggest that, in a sentence like (3-a), *fucking* would be more likely to receive a local reading, while, in a sentence like (3-b), *fucking* would be more likely to receive a non-local reading. This stems from the fact that, in (3-a), the noun modified by *fucking* is more likely to be thought of as responsible for the state of affairs. As such, a hearer will be more likely to infer that the speaker intends to direct their emotion towards that individual, rather than expressing emotion about the state of affairs in general. In contrast, the *paint* in (3-b) is not a possible culprit; that is, the *paint* is not responsible for putting itself on the carpet.

Following Frazier et al. (2015), I will henceforth refer to this as the CULPRIT HYPOTHESIS. More specifically, the CULPRIT HYPOTHESIS "claims that the negative attitude conveyed by the expressive tends to be construed with respect to an individual/entity that could be held responsible for the eventuality described" (Frazier et al., 2015, p. 299). In the absence of such an individual, as in (3-b), the expressive tends to be construed with respect to the state of affairs, i.e., that the carpet is stained with paint.

6.3.1.1 What is animacy?

Animacy concerns whether or not something is sentient or alive. In some languages, such as Japanese and Navajo, the animate-inanimate distinction is reflected in the grammar, such as in case-marking systems. In English, animacy is marked in the pronouns *I*, *you*

and *he/she*, although not in the third person plural pronoun *they*. There is also corpus evidence that the animacy of a noun can influence word order in English, with speakers preferring to use animate nouns in subject position and inanimate nouns in object position in transitive sentences (Dahl & Fraurud, 1996). Other factors shown to be sensitive to the animacy distinction include Saxon genitive vs of-genitive, double NP vs prepositional dative, active vs passive and pronominal vs full noun reference (see Zaenen et al., 2004).

There is no difference between animate and inanimate nouns in terms of semantic type. Consider the examples *dog* and *stone*. Both are one-place predicates that map individuals to truth-values if and only if they satisfy those properties. There are however selectional restrictions that mean specific verbs can only take noun phrases with certain thematic roles as arguments (Chomsky, 1965). For example, a verb like *amuse* requires a human experiencer, while a verb like *eat* needs an animate agent. Animacy is relevant for our understanding of thematic relations, with Agents and Experiencers argued to always be animate (Jackendoff, 1978, although see Dowty, 1991). Understanding animacy is also an important part of human cognition more generally. Humans rely on the distinction in their causal interpretation of actions, attribution of mental states, and attribution of biological processes (Szewczyk & Schriefers, 2011).

In examples like those in (4), which do not contain verbs with selectional restrictions, the only difference between (4-a) and (4-b), and between (4-c) and (4-d), is the animacy properties of the subjects. As we will see in the following section, these differences have consequences for how the swearwords will be interpreted with respect to the rest of the sentence.

(4) a. The fucking dogs were on the sofa
b. The fucking shoes were on the sofa
c. My bloody friend is stuck in the snow
d. My bloody truck is stuck in the snow

6.3.1.2 The culprit hypothesis

As discussed previously in Section 2.3.3.1 of Chapter 2, the flexibility of expressive modifiers to apply to different constituents of a sentence has been tested experimentally. To briefly re-review, Frazier et al. (2015) showed sentences like those in (5) to participants, with the word *damn* appearing before either the subject (5-a), the object (5-b) or the whole sentence (5-c). For each, participants were asked multiple choice questions like that in (6). The null hypothesis was that responses would categorically match-up with the position of *damn* in the sentence, e.g., for (5-a), participants would always give the answer 'the holiday'; this would contradict the expected flexibility of interpretation of expressive modifiers.

(5) a. The damn holiday is on the weekend.
b. The holiday is on the damn weekend.
c. Damn. The holiday is on the weekend.

(6) Which is the speaker most likely to have a negative attitude toward?
i The holiday
ii The weekend
iii The holiday being on the weekend

Frazier et al.'s results suggest that sentence-modifier (or non-local) readings are available for all three sentence types, with sentence-modifier responses given for subject-modifier sentences on 39% of trials and for object-modifier sentences on 42% of trials. While similar experimental work by Gutzmann (2019) on German suggests that this flexibility is somewhat limited - see the examples of embedded expressives discussed in Section 2.3.3.1 of Chapter 2 - the findings of Frazier et al. (2015) do hold for sentences such as these.

As their results further show, another important factor in predicting how speakers will interpret the negative attitude expressed by *damn* is the animacy properties of the noun it modifies. In their experiment, Frazier et al. contrasted *non-causal* sentences like those in (5) with *causal* sentences like those in (7). Causal sentences contained an animate subject and featured the same experimental manipulations as in (5) i.e., the three different sentence positions for *damn*.

(7) a. My damn neighbor drove over my lawn.
b. My neighbor drove over my damn lawn.
c. Damn. My neighbor drove over my lawn.

Frazier et al. (2015) predicted that sentences with causal relations would seem "easy to interpret in terms of laying blame on the person or entity responsible" (p. 294); in these cases, "the reader may infer that the author has a negative attitude toward the entity responsible for bringing about the unwanted situation". This is the case in the sentences in (7), in which the *neighbor* is plausibly responsible for the unwanted situation of someone having driven of the speaker's lawn. This contrasts with the sentences in (5), in which the *holiday* cannot plausibly be responsible for the unwanted situation of it being on the weekend.

Frazier et al. refer to this prediction as the CULPRIT HYPOTHESIS; the availability of a culprit that is plausibly responsible for the state of affairs would tend to lead to an interpretation where the expressive is construed with respect to that culprit. This prediction was borne out in their results. Across all sentence positions, but particularly when *damn*

was in subject-modifier position, causal sentences were more likely to lead to a subject-level reading of *damn* than non-causal sentences. That is, for example, in sentences like those in (7), participants were more likely to think that the speaker was expressing a negative attitude towards the *neighbor* than they were to think that the speaker of sentences like those in (5) was expressing a negative attitude towards the *holiday*.

In Frazier et al. (2015), in causal sentences, there was always an animate subject and, in most cases, an inanimate object, as in (7). This follows the general trend in English (Dahl & Fraurud, 1996). The CULPRIT HYPOTHESIS should however predict that the minimally different sentences in (4), repeated in (8), would be interpreted differently by readers due to the presence of a causal relation in (8-a) and (8-c), but not in (8-b) and (8-d). More specifically, the presence of an animate subject (e.g., *dogs* or *friend*) should lead to the local reading of a swearing intensifier being more likely compared to when the subject is inanimate (e.g., *shoes* or *truck*).

(8) a. The fucking dogs were on the sofa
b. The fucking shoes were on the sofa
c. My bloody friend is stuck in the snow
d. My bloody truck is stuck in the snow

Frazier et al. (2015) suggest that the more general flexibility effect they observe provides further support for their overarching claim about expressives, namely that they constitute separate speech acts to the asserted content of a sentence. Frazier et al. suggest that expressives can convey a general negative attitude towards a proposition or towards a particular individual. Determining what that attitude is directed towards requires pragmatic inference on behalf of the addressee; both readings are available, but they must determine, based on the available evidence, which reading is most likely.

Exploring the inference in more detail, the higher likelihood of a local reading for the swearing intensifiers in (8-a) and (8-c) than in (8-b) and (8-d) is likely the result of conversational norms and the hearer's prior knowledge of two factors. The first is participants' real-world knowledge of the property of animacy; as discussed in the section above, nouns in English are not marked for animacy. Rather, people have real-world knowledge that words like *kids*, *neighbour* and *policeman* denote individuals that are sentient, while words like *paint*, *lawn* and *ticket* denote individuals that are not.

The second factor is participants' understanding of the role animacy plays in the causal interpretation of actions, an important part of human cognition more generally (Szewczyk & Schriefers, 2011); Frazier et al. (2015) suggest, citing work by Hobbs (1979) and Kehler (2002), that their results perhaps reflect "a pervasive cognitive tendency to postulate causal relations when they might be warranted" (Frazier et al., 2015, p. 300). People will infer a causal relation between an Agent and a state of affairs even when one is not

explicitly stated. For example, in (8-c), they might assume a responsibility on behalf of the friend for getting stuck in the snow, even though this is not semantically entailed.

Combining their knowledge of what constitutes an animate entity with their understanding of how this relates to that entity's role in and responsibility for the state of affairs, people would be more likely to think that, in sentences like (8-a) and (8-c), the speaker is expressing their anger towards the Agent of the sentence for their likely role in causing the state of affairs.

The CULPRIT HYPOTHESIS is pragmatically driven, rather than the result of some truth-conditional difference between minimally different sentences containing animate and inanimate objects that are preceded by swearing intensifiers; it is not the case that sentences like (8-a) and (8-b) compose differently. People read such sentences and draw an inference about what the most likely interpretation of the swearword is based on their understanding of the real world and of causal relations.

This is the explanation given by Frazier et al. (2015) for explaining why a particular reading of an expressive is most likely, given that multiple readings are available. If a speaker "syntactically associates" an expressive with a particular part of the sentence (e.g., the subject in 'the damn dog is on the couch'), the hearer may infer that that the speaker wishes to "imply that the negative attitude expressed is directed toward the referent of the constituent containing the expressive adjective, rather than the whole proposition itself" (Frazier et al., 2015, p. 294). In sentences containing an animate noun, even if the expressive is not syntactically associated with that noun (e.g., in 'the dog is on the damn couch), the hear may infer that the speaker has a negative attitude towards the entity responsible for the unwanted situation. As a result, the negative attitude might "transfer to the person or entity responsible for the situation" (Frazier et al., 2015, p. 294).

In summary, the sociolinguistic variable of *animacy* that I am proposing concerns the animacy properties of the noun following a swearing intensifier. While both local and non-local readings of a swearing intensifier are available when modifying either an animate or inanimate noun, the local reading is more likely with an animate noun due to listeners' understanding of the relationship between language and the causal interpretation of actions.

6.3.2 Variable II: Gradability

The second linguistic variable tested in this chapter is *gradability*. My claim concerns examples like the minimally different sentences in (9). In both examples, the expressive *fucking* appears prejacent to a sentence-predicate. The important difference is the gradability properties of the predicates involved, with *fucking* appearing before the gradable predicate *long* in (9-a) and the non-gradable predicate *wooden* in (9-b).

(9) a. The table is fucking long
b. The table is fucking wooden

In this section I will make three claims about swearing modifiers and gradability which combine to create what I will henceforth refer to this as the SCALE HYPOTHESIS. Firstly, when a swearword modifies a gradable predicate, as in (9-a), it can receive either a local (degree modifier) interpretation, or a non-local interpretation. Secondly, when a swearword modifies a non-gradable predicate, as in (9-b), it can only receive a non-local interpretation. Thirdly, in examples like (9-a), in which a swearword modifies a gradable predicate, the local interpretation seems to be the preferred interpretation. That is, the preferred reading of (9-a) is one in which *fucking long* has the same meaning as *very long*.

6.3.2.1 What is gradability?

A gradable predicate is a predicate that has a lexically encoded scale in its semantics. For example, the adjective *tall* has a lexically encoded scale of *height*. An individual, such as a person or a building, can be *very tall*, *quite tall* or *somewhat tall*. That individual can also be evaluated as *taller* than another individual. This contrasts with non-gradable adjectives such as *dead*. Here there is no lexically encoded scale of *deadness*. Rather, *dead* is either true of an individual or it is not. An individual cannot be *very dead* or *somewhat dead* and one individual cannot be *more dead* than another individual.

A full account of the semantics of gradability is unnecessary for my purposes (see Kennedy & McNally, 2005, Kennedy, 2013 and Burnett, 2017a, among many others). I will instead present some behavioural characteristics of both gradable and non-gradable predicates and their diagnostics.

A well-known property of gradable predicates is their *context sensitivity*; what counts as *tall* in one context will be different in another. Gradable predicates are therefore computed with respect to a standard of comparison. If an individual is *tall*, then that individual exceeds a standard of tallness in a given context (Kennedy & McNally, 2005 and numerous citations therein). The same is not true for non-gradable adjectives; what counts as *dead* is the same in all possible contexts.

Gradable adjectives have been argued to denote measure functions of type $\langle d, et\rangle$ that combine with a null *pos* morpheme to derive a property (Kennedy & McNally, 2005). This *pos* morpheme picks out a standard of comparison on a scale that is associated with this measure function. For example, for the gradable adjective tall, *pos* picks out a standard degree on a scale of *height*. The standard is the lowest degree on that scale which 'stands out' on that scale in the relevant context. That is, a sentence like 'John is tall' is said to be true if and only if John is tall equal to or greater than the contextually relevant

standard degree of height. In contrast, non-gradable adjectives denote functions of type $\langle e, t \rangle$. In a sentence like 'John is dead', *dead* maps the individual *John* to TRUE if and only if John is dead.

The are several diagnostics for gradability. While discussing these I will return to the examples in (9). The first diagnostic is the ability to occur in comparative constructions. In such constructions, gradable adjectives should be felicitous (10-a) and non-gradable adjectives should be infelicitous (10-b)[1].

(10) a. This table is *longer* than that one
b. *This table is more *wooden* than that one

A similar diagnostic is the ability to occur in 'how-'questions. In such questions, gradable adjectives should again be felicitous (11-a), while non-gradable adjectives should be infelicitous (11-b).

(11) a. How *long* is the table?
b. *How *wooden* is the table?

A further diagnostic is the ability to pass the *for*-phrase test. Gradable adjectives should be felicitous in sentences of the form 'This is X for a Y' where X denotes an adjective and Y denotes a noun (12-a), while non-gradable adjectives should be infelicitous (12-b).

(12) a. This is long for a table
b. *This is wooden for a table

Finally, gradable adjectives can be degree modified (13-a) and non-gradable adjectives cannot (13-b). In essence, a degree modifier like *very* increases the amount by which the contextually relevant standard must be exceeded (Morzycki, 2016). The contextual standard must be exceeded to a large degree for *very long* to be true of *the table*, for example. Context determines what counts as *to a large degree*. Non-gradable adjectives cannot be modified by degree modifiers, as this would lead to a type mismatch due to the lack of a lexically encoded scale in their semantics.

(13) a. The table is very long
b. The table is very wooden

In this section I have given a small snapshot of gradability. I will now move on to describing and providing evidence for the SCALE HYPOTHESIS in more detail.

[1]It is worth noting that while *wooden* is non-gradable in the traditional sense, it could be gradable under a mereological reading. Under a mereological reading, (10-b) could be true if a greater number of parts of the table in question were made of wood than the other table (Burnett, 2014).

6.3.2.2 The scale hypothesis

The SCALE HYPOTHESIS that I am proposing consists of three sub-hypotheses, restated below as 1a. 1b. and 1c. respectively.

1a. When a swearword modifies a gradable predicate it can receive either a local (degree modifier) interpretation, or a non-local interpretation.

1b. When a swearword modifies a non-gradable predicate it can only receive a non-local interpretation.

1c. When a swearword modifies a gradable predicate, the preferred interpretation is the local (degree modifier) interpretation.

I will start with evidence for 1a., that is, the availability of both readings for examples like (14-a).

(14) a. The table is fucking long
b. The table is fucking wooden

The availability of a local degree-modifier reading for *fucking* is clear in examples like (15) and (16), in which speaker A's assertion about Trump or the table can be denied by B while maintaining that it is still true that Trump is old or that the table is long, while not being *fucking old* or *fucking long*. This suggests that a swearing modifier like *fucking* can increase the amount by which the contextually relevant standard on the scale associated with an adjective (e.g., age or length) must be exceeded. That is, it can behave like a degree modifier.

(15) A: Trump is fucking old!
B: No he's not. He's only 75. He's old, but he's not fucking old.

(Hazel Pearson, pc)

(16) A: That table is fucking long!
B: No its not. I've seen lots of tables much longer than that. It's long but not fucking long.

This can also be seen in the example conversation in (17), in which A and B are discussing the same table throughout. This again highlights the availability of a local reading for swearing modifiers.

(17) A: Is the table long?
B: Yes it is

A: Is it fucking long?
B: No not fucking long, just long

The availability of the non-local reading for sentences like (14-a) can be seen in the following examples. Firstly, in the example in (18), in response to A's statement about Trump, B can signal their agreement with A's emotional commitment to the assertion that Trump is old, perhaps agreeing that this is a noteworthy piece of information about him. At the same time, they can assert that he is not *very old*, creating a contrast between *very old* and *fucking old*. B's agreement only targets the non-local meaning of *fucking*.

(18) A: Trump is fucking old!
B: Yes he is! He's only 75, so not *very* old. But fuck yeah, he's old

(Hazel Pearson, pc)

Another example of this is the lack of a degree modifier reading for *fucking* in the *AP as AP can be* constructions discussed by Potts et al. (2009). These constructions are argued to have an identity condition, whereby each AP must match. This condition explains the felicity of (19-a) and the infelicity of (19-b). The inclusion of an expressive appears to be a unique exception to this matching condition however, shown by (19-c). There is no degree modifier reading available here, even though *fucking* is modifying a gradable adjective. It is not possible to replace the expressive with *very*, as in (19-d). In (19-c) therefore, only the non-local reading appears to be available. In (19-c), *fucking* just expresses the speaker's emotional commitment to their assertion about Sue.

(19) a. Sue is as funny as funny can be
b. #Sue is as funny as humorous can be
c. Sue is as fucking funny as funny can be
d. #Sue is as very funny as funny can be

Another example in which a degree modifier reading appears to be absent is in questions of the form 'how AP is X?'. While *very long* cannot be used in such questions (20-b), *fucking long* can (20-a). This again suggests that *fucking* can have a non-local reading. Again, *fucking* appears to express the emotional commitment of the speaker to their question.

(20) a. How fucking long is the table?
b. *How very long is the table?

Taken together, these examples provide evidence to support 1a. of the SCALE HYPOTHESIS by showing that, when modifying a gradable predicate, a swearword can have either a local or non-local interpretation.

I will now move on to evidence for 1b. of the SCALE HYPOTHESIS which states that swearwords that modify non-gradable predicates can only receive a non-local interpretation. This is clear in examples like (21) and (22). In (21), it is possible for B to deny the content of A's utterance, but they can only target the assertion that 'Elvis is dead'. They cannot disagree with A's statement and simultaneously assert that Elvis is dead. Their assertion that 'he's dead, but he's not fucking dead' is a contradiction. The same is true in (22); it is infelicitous for B to disagree with A and then assert 'it's wooden'. This suggests that the use of *fucking* in (21) and (22) cannot have a local reading. It can only have a non-local reading, expressing the speaker's emotional commitment to their assertion.

(21) A: Elvis is fucking dead!
B: No he's not # He's dead, but he's not fucking dead!

(22) A: The table is fucking wooden!
B: No it's not # It's wooden, but it's not fucking wooden!

(Hazel Pearson, pc)

The important difference between sentences containing e.g., *fucking long* and *fucking wooden* is that the former appears to have two available readings, while the latter only has one. Ideally, we want to assume a unified treatment of an expressive modifier like *fucking*, rather than have two different lexical entries. This is possible if we assume that all instances of expressive modification are completely separate from the at-issue content of a sentence and that the degree modifier meaning of *fucking* arises from pragmatic reasoning. Even in the non-local reading of *fucking*, however, the modifier is still targeting some kind of scale. Rather than targeting a lexically encoded scale, it targets an attitudinal scale at the speech act level. In the following section, I go into more detail about different types of scales. I will also provide evidence for 1c. of the SCALE HYPOTHESIS.

6.3.2.3 Scales in intensification

In their paper on conveying emphasis for intensity, Beltrama and Trotzke (2019) identify a number of lexical and syntactic strategies for intensification. These include the examples in (23). Although different in other ways, both of the modifiers in (23) "force a contraction of the denotation of the modified predicate" (p. 2), obtained by ranking elements along an ordered dimension and eliminating those at the lower end. For example, *extremely tall* is a subset of *tall* which only includes elements that are ranked highly on the scale of tallness. Similarly, the set denoted by *SALAD-salad* contains a subset those elements included in the set denoted by *salad* which are most prototypically salad-like.

(23) a. Mark is **extremely** tall.

b. I'll make the tuna salad, you make the **SALAD**-salad.

Other modifiers target an ordering with respect to a speaker's attitude, rather than targeting a scale in the denotation of the predicate. In (24), the use of *totally* is of the speaker-orientated variety discussed in previous work in relation to social meaning (Beltrama & Staum Casasanto, 2017). Beltrama and Trotzke (2019) suggest that pragmatic *totally* targets an attitudinal scale, heightening the speaker's commitment towards the proposition (see also Beltrama, 2018).

(24) You should **totally** click on that link.

Beltrama and Trotzke (2019) compare the ways in which different types of scales are targeted by an intensifier using the examples in (25) and (26) (p. 4).

(25) Scale in the denotation

a. Mark is tall, but not extremely tall

b. A: Mark is extremely tall
B: No! He is just tall

(26) Attitudinal scale

a. #You shouldn't *totally* click on that link! (Intended: 'I'm not certain you should click')

b. A: You should *totally* click on that link
B: No! # I should click on that link, but you can't be certain about giving this advice.

In (25), *extremely* targets a scale that involves an individual and their ordering on a scale; it can be used in the scope of a negator (25-a) and it can be challenged with denials (25-b). *Totally*, on the other hand, can apply to proposition-level or speech act-level objects, as it does in (26-b); this meaning cannot be used in the scope of a negator (26-a) or challenged by denials (26-b). Beltrama and Trotzke identify similarities between pragmatic *totally* and the expressive *damn*, in that both signal a high degree of a particular attitude towards a proposition, with the latter conveying the emotional involvement of the speaker.

There are two sources of variation here that Beltrama and Trotzke identify: variation in the nature of the target ordering (i.e., lexical, context-based, or attitudinal) and variation in the nature of the target linguistic elements (i.e., individuals, propositions, or speech acts); they suggest that this variation can be modulated by semantic factors such as the gradability properties of a modified predicate. It has already been suggested that the same form can target different types of orderings and different types of elements, for example the lexical and pragmatic uses of *totally* (Beltrama, 2018). The same has been suggested

for *so* (Irwin, 2014) and the *-ass* suffix (Irwin, 2015) in varieties of English, as well as the Italian suffix *-issimo* (Beltrama & Bochnak, 2015) and the German intensifier *sau* (Gutzmann, 2015).

Importantly when it comes to *totally*, both the lexical and pragmatic readings are available when it is paired with a bounded adjective like *bald*, as in (27-a). Under one reading, *totally* targets the lexically encoded scale of baldness; this is the most obvious reading. The pragmatic reading of *totally* is also available in (27-a), however. That is, it is possible that the speaker is using *totally* to maximally convey their commitment to the proposition that John is bald. Admittedly, this reading is made clearer if the sentence has an accompanying context to emphasize its availability, perhaps as in (27-b). Similarly, the example in (28) with *finished* also emphasizes the availability of the speaker-oriented meaning of *totally*, with the lexical meaning still available.

(27) a. John is totally bald
b. A: I just don't believe that John is bald
B: John is totally bald! It's obvious!

(28) (Context: A has just walked in and seen B's painting)
A: That doesn't look finished
B: It's totally finished, what are you talking about?!

Using pragmatic reasoning over alternatives, a person hearing a sentence like (27-a) is likely to infer that the speaker could have used an alternative sentence to most effectively convey the pragmatic meaning, had that been their intention. The pragmatic reasoning here follows Grice's Cooperative Principal and the maxim of manner (Grice, 1975), whereby speakers are assumed by listeners to be communicating clearly and without ambiguity.

For example, the sentence in (29-a) would also have been available to the speaker. Arguably, had the speaker intended to use the pragmatic version of *totally*, the sentence in (29-a) would have been of greater utility. By moving *totally* further away from the bounded adjective *bald*, the lexical meaning of *totally* is no longer available. As such, this option would have avoided any possible ambiguity. The hearer, aware of the availability of alternative utterances like (29-a), might therefore infer that the speaker was well-motivated in choosing (27-a) over (29-a) for conveying their desired meaning. The result of this Manner implicature (Grice, 1975) is that the lexical meaning of *totally* is the preferred reading in (27-a). The same would be true regarding B's response in (28) and the availability of (29-b) as an alternative structure.

(29) a. John totally is bald
b. The painting totally is finished

While the difference between sentences containing e.g., *totally bald* and *totally tall* is a semantic one, therefore, pragmatics is playing a role in the interpretation of the former. That is, despite both lexical and pragmatic meanings being available, the preferred reading of (27-a) is the lexical meaning due to the hearer's pragmatic reasoning over alternative utterances and their meanings.

I would like to propose something similar for swearing intensifiers like *fucking*. Unlike *totally*, *fucking* always makes the same semantic contribution to a sentence; whether it precedes a gradable or non-gradable adjective, it always targets an attitunal scale, intensifying the degree to which the speaker is emotionally committed to their assertion. Due to the availability of a lexically encoded scale in the semantics of a gradable adjective, however, and the availability of multiple other sentence structures for communicating the desired meaning, *fucking* can appear to target an individual and their ordering on a particular scale (e.g., length). I demonstrated the availability of this reading in examples like (30-a) in Section 6.3.2.2 above. The context provided in (30-b) emphasizes the availability of the non-local reading in a sentence with a gradable adjective.

(30) a. The table is fucking long
b. A: I don't remember the table being long
B: Of course the table is fucking long, we used it for Christmas dinner
c. The table fucking is long
d. Fuck! The table is long

Without this additional context however, the degree modifier reading is the most prominent. As with *totally*, this is due to pragmatic reasoning over alternative utterances that could have been used to unambiguously communicate the desired non-local meaning, but were not. For example, if the speaker had wanted to convey the non-local meaning, they could have used (30-c), or indeed (30-d), albeit with a different derivation of *fuck* (this wouldn't be the case with *damn*). Again, by moving the modifier further away from the adjective, the degree modifier meaning is absent. For example, in (31), B cannot felicitously deny the content of A's statement and simultaneously assert that the table is long. This contrasts with (16), repeated in (32), in which B can deny A's assertion while maintaining that it is true that the table is long.

(31) A: That table fucking is long
B: No it's not, # it's long, but it fucking isn't long

(32) A: That table is fucking long!
B: No its not. I've seen lots of tables much longer than that. It's long but not fucking long.

The sentence in (30-a) is ambiguous between a local and a non-local reading. The sentence in (30-c) unambiguously expresses the non-local reading, as does the sentence in (30-d). Given the Maxim of Manner (Grice, 1975), if a speaker had wanted to convey the non-local interpretation, they should have used (30-c) or (30-c). As they did not, they must not have wanted to convey the non-local interpretation, but rather, the local interpretation, when they used (30-a).

1c. of the SCALE HYPOTHESIS is motivated by this reasoning. The local interpretation of a swearing modifier is the preferred interpretation for sentences in which it modifies a gradable predicate. It is the preferred interpretation due to the hearer's pragmatic reasoning over alternative utterances, such as (30-c), and their meanings. The non-local interpretation of such sentences requires a special context in order to become the preferred interpretation, such as the one in (30-b). Although I do not have the type of experimental data found by Frazier et al. (2015) for the CULPRIT HYPOTHESIS, the examples presented in this section do suggest that the local interpretation of a swearing modifier is preferred to the non-local interpretation when the modifier precedes a gradable predicate.

In summary, the sociolinguistic variable of *gradability* that I am proposing concerns the gradability properties of the predicate modified by a swearing modifier. Both local and non-local interpretations of a swearword are available when modifying a gradable predicate. On the local interpretation, the swearword appears to target an individual and their ordering on a particular scale; this is the preferred interpretation. On the non-local interpretation, the swearword appears to target a proposition and its ordering on an attitudinal scale. When a swearword modifies a non-gradable predicate, however, only the non-local interpretation is available; again, on this interpretation, the swearword appears to target a proposition and its ordering on an attitudinal scale.

6.4 Experiment I: Perceiving semantic variation swearing

The previous section drew two categorical distinctions. The first was between swearwords that precede animate nouns and those that precede inanimate nouns. The second was between swearwords that precede gradable predicates and those that precede non-gradable predicates. I will now turn my attention to examining the extent to which people have an implicit knowledge of these distinctions. That is, do people *perceive* the difference between e.g., *fucking kids* and *fucking paint*, or *fucking tall* and *fucking wooden* in minimally different sentences?

If people have an internal representation of these contrasts that shows up in an automatic process like reading, then the contrast may also affect a more conscious process

like social evaluation. Following the logical steps outlined in Chapter 3, before testing whether people *notice* either of these contrasts in social evaluation, it is important to first establish whether they perceive them at all.

6.4.1 Methods

I tested people's implicit knowledge of these distinctions using a self-paced reading task. I reviewed the use of this task in Section 3.3.3 of Chapter 3, but certain details bear repeating. A self-paced reading task involves presenting participants with sentences on a screen, one word at a time; participants control the sentence presentation, pressing a button to reveal another word, with the time between button presses recorded by the experimenters. Longer periods of time between button presses are understood to reflect increased processing costs. The incremental nature of the experiment allows researchers to identify particular portions of the sentence that might be causing greater processing difficulty (see Jegerski, 2014 for a review of the method).

Using a self-paced reading task with these specific variables allowed for the detection of on-line processing differences between local and non-local swearing intensification. The task combined two sets of stimuli to test for the effects of both animacy and gradability on reading speeds. It was expected that local intensification would be triggered when swearing intensifiers were immediately prejacent to animate nouns and gradable adjectives, whereas non-local intensification would be triggered when swearing intensifiers were immediately prejacent to inanimate nouns and non-gradable adjectives. The dependent variable was the reading times in milliseconds for the word immediately following the swearword and every word thereafter till the end of the sentence, as well as the cumulative reading time for all of these words.

6.4.1.1 Experimental hypotheses

There were two related experimental hypotheses for this experiment. They were as follows:

H1 Sentences with swearwords that precede inanimate nouns will be read slower than sentences with swearwords that precede animate nouns.

H2 Sentences with swearwords that precede non-gradable predicates will be read slower than sentences with swearwords that precede gradable predicates.

The logic behind these hypotheses is fairly simple and is indirectly supported by previous work on modifiers that can be ambiguous between local and non-local readings (Dörre et al., 2018). The logic is that when a swearword modifier receives a local interpretation, the reader will likely only consider the syntactic sister of that swearword to

understand it's meaning. That is, upon encountering a sentence like (33), the modifier *fucking* is likely to be read as only expressing a heightened emotion towards the underlined portions of the sentence, that is, the syntactic sister of the modifier, the modified noun. As a result, once the meaning of that word has been integrated with the modifier, the reader can continue reading the sentence at normal speed.

(33) The fucking dog/kids was/were on the sofa (local interpretation)

Similarly, when encountering a sentence like (34), if the reader interprets the swearword locally, it will be read as only expressing heightened emotion towards the underlined portions of the sentence, namely the modified adjective. Again, once the meaning of that word has been integrated with the modifier, the reader can continue reading the sentence at normal speed.

(34) The table was fucking smooth/cheap and covered in scratches (local interpretation)

I expected this to contrast with when readers encountered sentences like (35) and (36). In a sentence like (35), it was predicted that the non-local interpretation of *fucking* would be most likely, per the CULPRIT HYPOTHESIS (Frazier et al., 2015). In this case, *fucking* should be read as expressing heightened emotion towards all the underlined portions of the sentence. As a result, the meanings of each part of the sentence must be integrated in order for the reader to understand the meaning of *fucking*. After encountering the modifier *fucking*, the reader has to refer back to the words that preceded it before paying greater attention than usual to the words that follow it.

(35) The fucking shoes/toys were on the sofa (non-local interpretation)

Similarly, when encountering a sentence like (36), the only available interpretation should be the non-local interpretation, per the SCALE HYPOTHESIS that I posited in Section 6.3.2.2. Again, the result should be that *fucking* is read as expressing a heightened emotion towards all the underlined portions of the sentence, requiring the reader to retrieve the words they read before the modifier, and concentrate on the words thereafter, in order to understand the meaning of *fucking*. It is for this reason that I expected reading to take longer for sentences with inanimate objects and non-gradable adjectives.

(36) The table was fucking hexagonal/Swiss and covered in scratches (non-local interpretation)

It is worth remembering that, in both cases, the non-local interpretation is always available. For example, when encountering the sentences in (33) and (34), the reader might

generate the non-local interpretation of *fucking* in which it expresses a heightened emotion towards all the underlined portions of the sentence. Similarly, when encountering a sentence like (35), although the CULPRIT HYPOTHESIS suggests that the non-local interpretation is most likely, the local interpretation is also available. If a reader generates this reading, then the *fucking* would be read as expressing heightened emotion towards the modified noun.

The hypotheses H1 and H2 reflect predicted tendencies in the reading behaviours of participants. In the case of H1, the CULPRIT HYPOTHESIS suggests that participants would tend towards the local interpretation of a swearing modifier when it modifies an animate noun and towards the non-local interpretation when it modifies an inanimate noun. In the case of H2, the *scale hypothesis* suggests that participants would only interpret the swearing modifier non-locally if it modifies a non-gradable adjective, whereas both the local and non-local interpretation are available when it modifies a gradable adjective. A preference for the local interpretation in the latter case should lead, on average, to faster reading times compared to sentences containing non-gradable adjectives.

Although neither of these distinctions has been previously addressed in this way in the literature, the work of Dörre et al. (2018) on German modal particles (MPs) provides some indication of where we might expect the slowdown effect to occur (see Section 3.3.3 of Chapter 3 for more detail on Dörre et al., 2018). In their self-paced reading task, Dörre et al. directly compared sentences containing the same MP (one of ten), but with either the local, at-issue meaning (37-a) or the non-local, not-at-issue meaning (37-b) primed. While it was expected that, in my experiment, the different readings would be primed by the word immediately following the modifier, their study involved another sentence before the critical sentence which unambiguously primed one of the two readings.

(37) a. Sie hat zwar sehr viel Geld abgeholt, doch sie soll
She has indeed very much money withdrawn, but she should
bloß das Kleid kaufen.
BLOSS the dress buy
'Altough she has withdrawn a lot of money, she should only buy the dress.'

b. Sie soll lieber etwas mehr Geld abheben, denn sie soll
She should rather something more money withdraw, since she should
bloß das Kleid kaufen.
BLOSS the dress buy
'She should withdraw more money because she should really buy the dress.'[2]

When the local, at-issue meaning of *bloß* was primed, as in (37-a), participants had

[2]Both glosses are my own translation of the individual words. Only the original German and the rough translation of the whole sentence come from the original authors.

faster reading times on the +2 word (*Kleid*) and the +3 word (*kaufen*) compared to minimally different sentences where the non-local, not-at-issue meaning of *bloß* was primed, as in (37-b). The slowdown did not occur immediately after the MP, but rather, in the second and third words after the MP. Given the similarities between the distinction tested in Dörre et al. (2018) and the animacy and gradability distinctions being tested in this chapter - with their focus on local and non-local readings of the same word - we might expect the slowdown predicted in the experimental hypotheses to occur around the +2 and +3 regions after the swearword in sentences like (36) and (35) compared to sentences like (34) and (33).

Further indirect support for the experimental hypotheses comes from the work of Donahoo (2019) on the processing difference between taboo and non-taboo modifiers. This evidence is restricted to the reading times of the word immediately following the modifier, with no available data for the words thereafter; it is therefore not suggestive of the region in which a slowdown is to be expected. This work does, however, provide some evidence of a reading slowdown associated with non-local readings compared to local readings of modifiers. Participants were significantly faster at reading the word *car* in a sentence like (38-a) than in a sentence like (38-b). That is, the word *damn*, which is ambiguous between local and non-local readings, led to a slowdown compared to *old*, which can only be read as a local modifier.

(38) a. The old car broke down yesterday
b. The damn car broke down yesterday

Finally, another reason to expect a slowdown in non-local readings of swearing modifiers comes from the broader literature on the processing cost associated with non-local dependencies (Gibson, 2000; Lewis & Vasishth, 2005). In such sentences, for example in (39-a), the dependent (e.g., *what*) cannot be interpreted until the reader has parsed the head of the dependency. An increase in the distance between the dependent and the head, for example in (39-b) compared to (39-a), typically leads to a processing cost, reflected in longer reading times at the head of the dependency (Gibson, 2000; Lewis & Vasishth, 2005).

(39) a. Someone asked what the man did last summer
b. Someone asked what the man who loves football did last summer

Under Gibson's (2000) Dependency Locality Theory (DLT), this processing cost is linked to an increase in integration and storage costs. To successfully understand the sentence, a parser must store the structure in working memory, including any incomplete dependencies, and integrate each word that is encountered into the structure built thus far. In

(39-a), when the reader gets to the word *did*, they must retrieve the word *what* from their working memory in order to integrate it into the existing structure. The same is true in (39-b), but the reader has to hold the dependency in working memory for a longer period of time; they also have to integrate a greater number of other words into the existing structure while keeping track of the dependency until they get to the head.

While DLT and my own theory for the predicted processing cost of sentences like those in (40-a) and (40-b) are different in many ways, the underlying mechanisms responsible for processing these sentences are the same. They both concern a greater amount of processing resources linked to *integration* and *storage*. In both (40-a) and (40-b), once the non-local reading becomes clear, the reader has to retrieve the preceding words from working memory and integrate them and all following words with the swearword in order to understand it's meaning. In (40-c) and (40-d), this was not predicted to be the case; the reader must only integrate the swearword with the word immediately following it. My prediction was that the need to integrate the swearword with a greater number of other words in (40-a) and (40-b) would result in a processing cost, leading to a reading slowdown for those words following a non-gradable predicate/inanimate object.

(40) a. The fucking shoes/toys were on the sofa
b. The table was fucking hexagonal/Swiss and covered in scratches
c. The fucking dog/kids was/were on the sofa
d. The table was fucking smooth/cheap and covered in scratches

6.4.1.2 Stimuli

There were two separate sets of stimuli in this experiment: the animacy stimuli and the gradability stimuli. Each set of stimuli followed a different experimental design. The animacy stimuli set followed a 3x2 research design (see Table 6.1). The 3-level factor was Modifier (Bloody, Fucking, Damn). The 2-level factor was Noun Type (Inanimate, Animate). For the animacy stimuli, the decision was made to use three different swearwords. Frazier et al. (2015) have already established, albeit not using on-line methods, that readers are sensitive to the critical difference in the Noun Type factor. Their experiment only tested the swearword *damn*, a fairly mild swearword in British English. By including *fucking* and *bloody*, the current experiment tests whether the distinction is still a significant one when the swearword is either stronger, as with *fucking*, or distinctly British, as with *bloody* (Dewaele, 2015).

	Inanimate	**Animate**
Bloody	*bloody wind*	*bloody bear*
Fucking	*fucking wind*	*fucking bear*
Damn	*damn wind*	*damn bear*

Table 6.1: Experiment I: Animacy stimuli - 3x2 research design

The gradability stimuli set followed a 2x2 design (see Table 6.4.1.2). The first 2-level factor was Modifier (Fucking, Probably), with *probably* used as a non-swear baseline condition. The second 2-level factor was Adjective Type (Gradable, Non-gradable). Unlike with the animacy distinction, the gradability distinction is one that had not previously been tested experimentally. Rather than complicate the data by including multiple swear-words therefore, the decision was made to only use *fucking*; the results for *fucking* will be compared to those for a non-swear modifier (*probably*). The selection of *probably* will be justified in the following sub-section.

	Gradable	**Non-gradable**
Fucking	*fucking sweet*	*fucking Swiss*
Probably	*probably sweet*	*probably Swiss*

Table 6.2: Experiment I: Gradability stimuli - 2x2 research design

6.4.1.2.1 Animacy stimuli While the animacy stimuli were based partly on those in Frazier et al. (2015) - with respect to the sentence structure - new items were required. This was due to the additional constraints of a self-paced reading task, with sentences needing to be as uniform as possible. As such, the animate and inanimate nouns had to felicitously appear in the same sentence position. The full list of noun pairs is provided in Appendix B.1. The full list of sentence pairs is provided in Appendix B.2. In total there were 15 sentence pairs, containing 30 different nouns. Nouns were list-matched using group-wise t.tests in R for log lexical frequency and Bigram frequency from the SUBTLEX database (Van Heuven et al., 2014) (both $p > 0.05$).

To control for differences in collocation between the selected nouns and modifiers, test sentences were coded for bigram frequency using the SUBTLEX UK Bigram Database Van Heuven et al. (2014). For bigram frequency, the values for each MODIFIER-NOUN pair were compared by group. Matching for further linguistic factors was limited by the need for both words to sound natural in the same carrier sentence. For example, for the sentence 'I just saw that the fucking dogs were on the sofa again', the bigram frequency was included for *fucking* and *dogs*. Groupwise t-tests in R (R Core Team,

2018) showed there to be no significant difference between the animate and inanimate groups in MODIFIER-NOUN bigram frequency ($p > 0.2$).

The modifier in each sentence was either *fucking*, *damn* or *bloody*. The total number of test stimuli for animacy was 90, exhausting every combination of modifier and noun. These stimuli were divided into 3 sets of 30. Each set contained exactly one sentence for each noun (i.e. a sentence for each noun in every pair), randomised and counterbalanced for modifier, such that every set contained exactly 10 sentences with each modifier.

A non-swear control condition was not used for the animacy stimuli. While previous work by Donahoo (2019) used a swear vs non-swear distinction in nominal modification, this was not appropriate for the current study due to the constraints of a self-paced reading task. In the gradability stimuli, it was possible to use the same non-swear modifier across all sentences. For the animacy stimuli however, this was not possible. For example, it would be very difficult to find one adjective, or even a set of adjectives such as colour adjectives, that could felicitously be used to modify all of the possible objects (e.g. *wind*, *dog*, *heat*, *car*, *daughter* etc). Even if a suitable non-swear adjective was found for each pair of nouns, introducing more variability into the stimuli may have masked other interesting effects from emerging.

6.4.1.2.2 Gradability stimuli Stimuli for the non-gradable category were taken from Burnett (2014) and Ziegler and Pylkkänen (2016), with any additional stimuli being variations on the stimuli therein (e.g. *hexagonal* $\sim$ *circular*, *square*). Each non-gradable adjective was paired with a gradable counterpart. The gradable adjectives were taken from the English Lexicon Project (Balota et al., 2007) and were pairwise matched for character length, number of syllables and part-of-speech dominance. List-wise matching was also completed for the following additional factors for which groupwise t-tests were conducted to ensure no significant difference existed between the two groups: log lexical frequency using the SUBTLEX UK database (Van Heuven et al., 2014), number of morphemes, mean lexical decision accuracy and valency using measures from Warriner et al. (2013) (all $p > 0.1$); the two groups did differ significantly in arousal based on the same measures however ($p = 0.042$). Matching was restricted by the need to have both adjectives appear realistically in the same sentence context. Gradable adjectives associated with closed scales were not included. Closed scale adjectives such as *dry*, *straight* and *clean* are associated with maximal endpoints on their respective scales (Kennedy & McNally, 2005). In total, 18 adjectival pairs were selected for the test stimuli (see Appendix B.3).

In the test condition, the modifier was *fucking*. *Fuck*, and its derivations, is the most frequently used swearword in British English, with *fucking* the most frequent derivation and 'emphatic intensification' the most common function (McEnery, 2004; Love, 2017).

Given that swearwords are highly arousing (Janschewitz, 2008) - a factor shown to affect attentional processing (A. K. Anderson, 2005; Aquino & Arnell, 2007, among others) - using a highly frequent swearword like *fucking* helps reduce the further effects of frequency or surprisal on processing. It is also the case that variation in arousal and valency across different swearwords could have complicated the analysis; using a single swearword helps overcome this complication.

In the control condition, the modifier was *probably*. *Probably* is a modal adverb that takes the proposition expressed by a sentence as its argument (Bellert, 1977; see also Bach, 1999). As Bellert (1977) suggests, every utterance containing *probably* can be paraphrased as in (41).

(41) It is probably true that S

Probably therefore involves sentence-level (non-local) modification. In this way, it mirrors the non-local reading given to *fucking*; where it differs is that it relates to epistemic modality rather than intensification. With gradable adjectives, therefore, *probably* should lead to a reading in which it is probable that an entity has reached the required standard of that adjective's associated scale (e.g. *height*, *weight* etc) for that adjective (e.g., *tall*, *heavy*) to be true of that entity (see Kennedy & McNally, 2005).

Test stimuli consisted of declarative sentences containing one of the test adjectives. For each adjective pair, a sentence was constructed of the type 'I realised/learnt that the [NOUN] was [MODIFIER] [ADJECTIVE] and that it [VP]'. *Realise* and *learn* were used as embedding verbs to make *fucking* and *probably* equally as felicitous as modifiers. The [NOUN] in each case was an inanimate object; this reduced the possibility that *fucking* would be interpreted as a verb. The following context containing a VP was of the form 'and it was/had [ADJECTIVE]/[NOUN]' (e.g. 'and it was small' or 'and it had stripes'). The same test sentence was used for each adjective in a pair and for each modifier condition. This resulted in 72 test sentences (see Appendix B.4).

To control for differences in collocation between the selected nouns and adjectives and modifiers and adjectives, test sentences were coded for bigram frequency using the SUBTLEX UK Bigram Database Van Heuven et al. (2014). For example, for the sentence 'I realised that the bookcase was fucking/probably creaky...', the bigram frequency was included for *bookcase* and *creaky*, *fucking* and *creaky* and *probably* and *creaky*. Groupwise t-tests in R (R Core Team, 2018) showed there to be no significant difference between the gradable and non-gradable groups in NOUN-ADJECTIVE, fucking-ADJECTIVE or probably-ADJECTIVE bigram frequencies (all $p > 0.3$).

The 72 test sentences were divided into four counterbalanced test blocks, each containing test 36 stimuli, such that each set contained exactly on occurrence of each of the 36 adjectives, with 18 in each modifier condition and 18 in each gradability condition.

Each set contained exactly one occurrence of each of the 36 test adjectives.

6.4.1.3 Procedure

The self-paced reading task was constructed in Gorilla (Anwyl-Irvine et al., 2018). Participants were assigned to one of twelve stimuli sets, exhausting every combination of the three sets of animacy stimuli and the four sets of gradability stimuli. Each of the twelve sets also included a further fifteen filler sentences, consisting of five garden path sentences, five grammatical sentences containing other swearwords and five ungrammatical sentences.

Each test sentence was presented to participants one word at a time in the centre of the screen, with all words masked by black boxes either side of the shown word, thus hiding the length of individual masked words. Participants used the space bar on their keyboards to advance to the next word in the sentence. Once they had read the final word, pressing the space bar ended that trial. Reaction times between button presses were recorded. A 1000ms blank screen followed by 2000ms fixation cross was shown between each trial to refocus attention to the centre of the screen. On approximately one third of test trials, the trial was immediately followed by a 2-Alternative Forced Choice question about the acceptability of the previously read sentence, which participants answered using their arrow keys to answer 'yes' or 'no'. The ungrammatical filler sentences served as catch trials; each of these sentences was followed by a question about acceptability. Participants were informed about these questions on the understanding that they would not be paid for their participation if they answered them incorrectly. If a participant scored lower than 70% on these trials (i.e., all trials that included an acceptability judgment), their data would not be included in the analysis.

Before starting the test trials, participants were given instructions on the task. They also completed two practice trials. Practice trials included sentences from outside of the test stimuli. Each practice trial was followed by a 2-Alternative Forced Choice question about the acceptability of the previously read sentence. Participants were required to answer both questions correctly to continue to the test trials; participants were given three attempts at correctly completing the practice trials. Once test trials finished, participants completed two short questionnaires. The first elicited participants' gender, age and yearly household income (in blocks of 10,000). The second included questions regarding their own swearing behaviour and attitudes towards swearing; participants provided responses to the following questions, with answers from 1 (Never) to 5 (Very frequently):

1. How frequently do you swear?

2. How frequently do you use the word 'fuck'?

3. How often do you think swearing is appropriate?

6.4.1.4 Participants

The two studies principally used to motivate the experimental hypotheses in this experiment both had relatively small numbers of participants (Dörre et al., 2018; Donahoo, 2019). To maximise the possibility of finding a significant result, while simultaneously ensuring that this did not result from an overpowered experiment, statistical power analysis was employed. Psychological studies typically employ statistical power analysis in order to estimate the required number of participants to reach a reasonable level of power (typically 0.8) (Cohen, 1992). This requires an estimated effect size taken from previous studies with similar motivations, preferably after a meta-analysis of multiple studies. While empirical support for this study is provided most closely by Dörre et al. (2018), using the effect size from a single study risks inaccurately calculating statistical power. Furthermore, effect sizes taken from studies with small sample sizes (in this case, N = 60) are themselves often inaccurate (Leon, Davis, & Kraemer, 2011). An alternative approach is to conduct a smaller pilot version of the main experiment, with a view to conducting statistical power analysis using the observed effect size to estimate the required N; this has also been warned against in the literature, however, due to issues around the accuracy of effect sizes in small populations and the further issue of follow-up bias (Albers & Lakens, 2018). The follow-up bias concerns the tendency for only pilot studies with significant effects to be followed up, despite null results being a perfectly valid and often-times revealing result.

The number of participants used in the current study was instead calculated using a *sequential analysis* following Lakens (2014). This approach involves pre-planned data analysis at set intervals during data collection. The aim of sequential analysis is to minimize the need for overly large sample sizes where a smaller sample would achieve the requisite statistical power. For example, a researcher may stop a study after every 100th participant. At each interval, the desired statistical test is conducted on all observations collected up to that point. Provided that the study achieves a significant result with an effect size greater than the *smallest effect size of interest* (SESOI) - often any effect size different from 0 - data collection can cease.

Due to financial constraints, the current study had a maximum N of 600 participants; note that this is 10 times the N used in Dörre et al. (2018). 6 intervals of analysis were intended to be used, with one after every 100th participant. To control for Type 1 errors, a linear spending function was used (see Jennison & Turnbull, 1999). Spending functions calculate the cumulative Type I error spent up until the time of an observation and specify how much alpha to use at that particular interval of analysis; this value is very small at the first interval and gradually increases. The values are cumulative, such that the total of

the alpha values equals the original alpha value (e.g., 0.05, the standard alpha threshold in the social sciences). For example, when specifying an alpha value of 0.05 with a single interval, the alpha value at that interval would typically be 0.025.

For the current study, any effect above 0 was treated as the SESOI. Based on 6 intervals, using an original alpha value of 0.05, the alpha values were calculated using the *GroupSeq* package in R (Pahl, 2018); these are provided in Table 6.3. At each interval, effects had to reach these significance thresholds with a non-zero effect size for data collection to cease. Effect sizes for reaction time tasks are typically around $d = 0.1$ (Brysbaert & Stevens, 2018). The main effect size of interest from Dörre et al. (2018) was d = 0.181. As such, if the effect size at an interval was greater than $d = 0.2$ (defined by Cohen, 1992 as a small effect size), data collection would continue until $d < 0.2$ and the p-value was sufficiently small for that interval, or until the maximum N had been reached.

Number of participants	Nominal alpha
100	1.5794e-06
200	0.0006853154
300	0.0048877018
400	0.0108000698
500	0.0154159903
600	0.0182093433

Table 6.3: Experiment I: Nominal alpha values calculated in *GroupSeq* using Power Family: $alpha * t^{\phi}$

The requirements for Cohen's d and the alpha value were tested for at each interval for both sets of stimuli. It was determined that data collection would cease if, for either set of stimuli, the effect size was smaller than $d = 0.2$ and the p-value was lower than the alpha value specified for that interval, per the values in Table 6.3.

Participants were native speakers of British English based in the UK. They were recruited, screened and paid £2.00 for their participation via Prolific Academic (2022). Participants gave informed consent; they were warned that they would be reading sentences containing swearwords and that they could abort the task at any time.

6.4.2 Results

Data collection was ultimately stopped at 300 participants (176 Female, 123 Male, 1 Transmasculine), with a mean age of 37.54, because the criteria stated in the section above had been fulfilled for the animacy stimuli. That is, for the animacy stimuli, an effect was found which was smaller than $d = 0.2$ and which was significant to a threshold of $p = 0.00488...$, the nominal alpha for 300 participants. One participant was excluded because their mean by-word reaction time for test words was 2 standard deviations above the group

mean. 2.3% of observations were removed, as they were two standard deviations above or below the mean for that word in that position.

I will now detail the statistical analyses performed on each set of stimuli. For each, linear mixed effect regression analyses were completed. The specific main effect predictors, interactions and random effects used in each model are specified in each subsection. While the adjusted alpha value was used for the purpose of ending data collection, all effects significant to the original alpha threshold of 0.05 are reported. Models for each stimuli set were completed using both by-word and summed reading times. For each model, the original full model is first stated. The final reported models exclude any predictors that did not significantly improve model fit based on chi-square comparisons of the sums of the squares of the residuals; this was done using the *anova* function in R.

6.4.2.1 Animacy results

For the animacy stimuli, separate models were produced using the reading times in milliseconds at the Modifier, Noun, +1, +2, +3 and +4 positions as the dependent variables; Table 6.4 shows how these labels correspond to an example test sentence. Reading times at the +5 were not analysed, as these were much larger and varied much more significantly. This likely reflects the fact that participants were taking a moment to think about the sentence before continuing to a possible acceptability judgment question. As such, these responses are not as automatic as the reading time measures earlier in the sentence.

Inanimate	I	just	saw	that	the	fucking	dogs	were	on	the	sofa	again
Animate	I	just	saw	that	the	fucking	shoes	were	on	the	sofa	again
Labels						Modifier	Noun	+1	+2	+3	+4	+5

Table 6.4: Experiment I: Animacy stimuli - sentence labels

The full model was as follows. The main effect predictors Noun Type (*animate*, *inanimate*) and Modifier (*bloody*, *damn*, *fucking*), as well as their interaction, were included as main effect predictors; Noun Type was sum-coded and Modifier was treatment coded, with *bloody* as the reference level. Responses to the demographic survey and the swearing survey were also included as main effect predictors. Participant was included as a random intercept in all models, with a random intercept for Word included in all models except for the Modifier model, as this contrast is already captured by the Modifier predictor; a random slope for Participant over word was initially included, but was dropped due to convergence issues. Predictors and interaction terms that did not significantly improve model fit based on chi-square comparisons of the sums of the squares of the residuals were excluded. The interaction term was dropped for all but the +3 model, as it failed to improve model fit in all other models. .

Tables 6.5 and 6.6 summarise the results of the final models, with Cohen's *d* (Cohen, 1992) calculated for the experimental manipulations using the *lme.dscore* function in the *EMAtools* package (Kleiman, 2017). There was a significant main effect for Modifier: Fucking at the Modifier position ($\beta = 8.22$, p = 0.004, d = 0.0629). There was a significant main effect for Noun Type: Animate ($\beta = -4.04$, $p < 0.001$, d = -0.0886) at the +2 position; this was the effect that allowed data collection to be stopped, as the effect was significant to the adjusted alpha threshold at that interval. There were significant interaction effects at the +3 position for Noun Type: Animate and Modifier: Damn ($\beta = -5.79$, p = 0.006, d = -0.0599) and for Noun Type: Animate and Modifier: Fucking ($\beta = -5.05$, p = 0.04, d = -0.0446). There were additional significant main effects, across all sentence positions, for Participant Age ($3.5 < \beta < 3.78$, all $p < 0.001$) and 'Swearing = appropriate' ($-43.26 < \beta < -25.5$, all $p < 0.005$).

Table 6.5: Experiment I: Model summary: Modifier, Object, +1 - Reaction Time ~ Modifier + Noun Type + Participant Age + Swearing = appropriate + (1|Participant) + (1|Word)

	Modifier			**Object**			**+1**		
Predictors	*Estimates*	*CI*	*p*	*Estimates*	*CI*	*p*	*Estimates*	*CI*	*p*
(Intercept)	377.69	306.60 – 448.79	**<0.001**	365.70	301.17 – 430.22	**<0.001**	378.94	315.09 – 442.78	**<0.001**
Noun Type: Animate	0.13	-2.14 – 2.40	0.910	-2.92	-8.86 – 3.01	0.334	-2.14	-4.32 – 0.05	0.055
Modifier: Damn	-3.01	-8.58 – 2.57	0.290	2.13	-3.07 – 7.33	0.422	1.77	-3.57 – 7.11	0.516
Modifier: Fucking	8.22	2.68 – 13.77	**0.004**	3.52	-1.66 – 8.71	0.183	0.23	-5.11 – 5.57	0.932
Participant Age	3.50	2.44 – 4.57	**<0.001**	3.51	2.55 – 4.47	**<0.001**	3.50	2.57 – 4.44	**<0.001**
Swearing = appropriate	-31.15	-50.61 – -11.69	**0.002**	-25.50	-43.10 – -7.90	**0.005**	-27.07	-44.17 – -9.96	**0.002**
N	299 $_{\text{Participant}}$			299 $_{\text{Participant}}$			299 $_{\text{Participant}}$		
				30 $_{\text{Response}}$			11 $_{\text{Response}}$		
Observations	8828			8683			8708		

Table 6.6: Experiment I: Model summary: +2, +3, +4 - Reaction Time ~ Modifier * Noun Type + Participant Age + Swearing = appropriate + (1|Participant) + (1|Word)

	+2			**+3**			**+4**		
Predictors	*Estimates*	*CI*	*p*	*Estimates*	*CI*	*p*	*Estimates*	*CI*	*p*
(Intercept)	362.08	303.31 – 420.86	**<0.001**	368.35	308.72 – 427.98	**<0.001**	373.87	317.80 – 429.95	**<0.001**
Noun Type: Animate	-4.04	-5.99 – -2.09	**<0.001**	1.82	-1.61 – 5.24	0.298	0.59	-1.32 – 2.49	0.545
Modifier: Damn	0.55	-4.23 – 5.32	0.822	1.22	-3.61 – 6.06	0.621	-2.03	-6.70 – 2.64	0.394
Modifier: Fucking	0.04	-4.73 – 4.81	0.988	-0.26	-5.09 – 4.57	0.916	2.37	-2.29 – 7.04	0.319
Participant Age	3.57	2.70 – 4.45	**<0.001**	3.73	2.86 – 4.61	**<0.001**	3.78	2.97 – 4.60	**<0.001**
Swearing = appropriate	-27.45	-43.45 – -11.44	**0.001**	-28.50	-44.55 – -12.45	**0.001**	-28.37	-43.26 – -13.47	**<0.001**
Animate * Damn				-6.79	-11.62 – -1.95	**0.006**			
Aninmate * Fucking				-5.05	-9.88 – -0.22	**0.040**			
N	299 $_{\text{Participant}}$			299 $_{\text{Participant}}$			299 $_{\text{Participant}}$		
	13 $_{\text{Response}}$			11 $_{\text{Response}}$			14 $_{\text{Response}}$		
Observations	8746			8757			8690		

Figures 6.1-6.3 plot the mean reaction time and two times standard error, both in milliseconds, for the three modifiers in each condition at each position in the sentence.

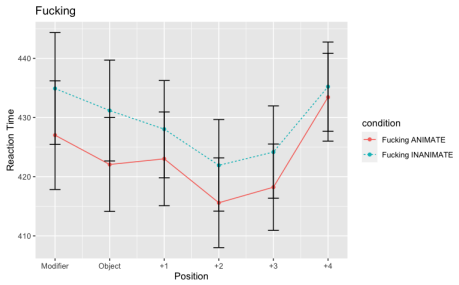

Figure 6.1: Experiment I: Animacy results for *fucking*

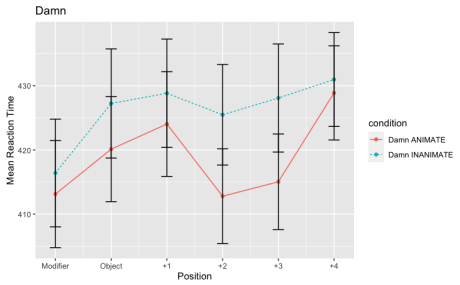

Figure 6.2: Experiment I: Animacy results for *damn*

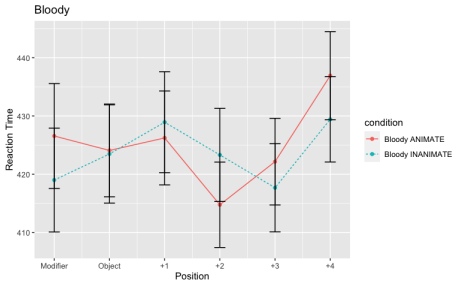

Figure 6.3: Experiment I: Animacy results for *bloody*

A further model was also constructed for the summed reaction time from the +1 to the +4 position using the same predictors as above. As a number of observations had been removed in data preparation (see above), reaction times were only totalled on trials where every data point was available. That is, if the reaction time for one word in a trial was missing, all of the data for that trial was removed; 8.37% of observations were removed in this manner.

All main effect predictors apart from Modifier and Noun Type were dropped after model comparison. Table 6.7 summarises the model. There were no significant main effects for either predictor (both $p > 0.6$).

Table 6.7: Experiment I: Model summary: +1, +2, +3, +4 - Summed Reaction Time $\sim$ Modifier + Noun Type + (1|Participant) + (1|Object)

	Summed Reaction Time		
Predictors	*Estimates*	*CI*	*p*
(Intercept)	2421.09	2273.88 – 2568.29	**<0.001**
Modifier: Damn	-50.96	-196.94 – 95.02	0.494
Modifier: Fucking	-27.24	-167.20 – 112.72	0.703
Noun Type: Animate	22.56	-62.00 – 107.12	0.601
$N_{Participant}$	163		
N_{word}	30		
Observations	460		

6.4.2.1.1 Exploratory analysis The animacy stimuli were created using a binary distinction between animate and inanimate objects. The status of animacy in natural language is more complicated than this, however. For example, many languages make at least the minimal distinction between human, non-human animate and inanimate. This is commonly known as the animacy hierarchy (Silverstein, 1976). Some languages further distinguish between higher and lower animals (Yamamoto, 1999). In English, the animacy hierarchy can affect other parts of the sentence, such as the choice between the Saxon genitive and the *of*-genitive (Zaenen et al., 2004). In their annotation project on English, Zaenen et al. (2004) distinguished between up to ten different categories of animacy, including *animal*, *place*, *vehicle*, and *organisation*. While the results reported above already provide support for one of my experimental hypotheses, it is worth checking whether this effect is being driven by a specific subset of either the animate or inanimate objects.

The stimuli were recoded using the schema from Zaenen et al. (2004). In the animate category, the objects were recategorized as either *human* (11) or *animal* (4). In the inanimate category, the objects were recategorized as *concrete* (8), such as *shoes* or *key*, *non-concrete* (5), such as *wind* or *heat*, or *vehicle* (2), such as *truck* or *car*. The category of *vehicle* was motivated by the fact that vehicles are frequently treated as living beings in certain linguistic contexts (Zaenen et al., 2004), such as the use of gendered pronouns, which are usually only used with animate objects. The same models as in the previous section were constructed, with the predictor Noun Type replaced by a new predictor called Animacy Category (shortened to Ani_Cat), with *human* as the baseline. These models were then compared to the previous models using chi-square comparisons of the sums of the squares of the residuals. Only the fits of the +1 and +4 models were significantly improved by the Animacy Category predictor variable. These models are summarised in Table 6.8; I will not report the Cohen's *d* for the exploratory analysis.

Table 6.8: Experiment I: Model summary: +1, +4 - Reaction Time ∼ Modifier + Animacy Category + Participant Age + Swearing = appropriate + (1|Participant) + (1|Object)

	+1			+4		
Predictors	*Estimates*	*CI*	*p*	*Estimates*	*CI*	*p*
(Intercept)	373.00	308.96 – 437.03	**<0.001**	374.26	318.21 – 430.31	**<0.001**
Ani_Cat: Animal	12.26	3.38 – 21.14	**0.007**	-0.09	-8.76 – 8.58	0.984
Ani_Cat: Concrete	9.11	2.76 – 15.46	**0.005**	2.82	-2.80 – 8.44	0.325
Ani_Cat: Non-concrete	9.07	1.39 – 16.74	**0.021**	-1.50	-8.11 – 5.10	0.655
Ani_Cat: Vehicle	-2.67	-12.57 – 7.24	0.597	-16.73	-26.85 – -6.60	**0.001**
Modifier: Damn	1.82	-3.51 – 7.16	0.503	-2.01	-6.67 – 2.66	0.399
Modifier: Bloody	0.21	-5.12 – 5.55	0.938	2.38	-2.29 – 7.04	0.318
Participant Age	3.50	2.57 – 4.44	**<0.001**	3.78	2.97 – 4.60	**<0.001**
Swearing = appropriate	-27.07	-44.18 – -9.97	**0.002**	-28.37	-43.26 – -13.48	**<0.001**
N	299 $_{\text{Participant}}$			299 $_{\text{Participant}}$		
	11 $_{\text{Response}}$			14 $_{\text{Response}}$		
Observations	8708			8690		

Both models had significant effects for Animacy Category in addition to the previously reported effects for Participant Age and Swearing = appropriate. In the +1 model, there were significant main effects, all $p < 0.05$, for Animacy Category: Animal ($\beta = 12.26$), Animacy Category: Concrete ($\beta = 9.11$) and Animacy Category: Non-concrete ($\beta = 9.07$). In the +4 model, there was a significant main effect for Animacy Category: Vehicle ($\beta = -16.73$, $p = 0.001$).

These results can be summarized as follows. Firstly, the effects reported in the previous section for the +2 and +3 models, with the binary predictor Noun Type (Inanimate, Animate), are still the best way of accounting for the variation in the data at these positions in the test sentences; this follows from model comparison. That is, the animate/inanimate distinction used to construct the stimuli originally is meaningful in explaining participants' reading times at these positions.

For the +1 and +4 models however, the variation is best explained using Animacy Category as a predictor. At the +1 position, i.e., the first word after the object, the fastest reading times were in the *human* condition, with the *vehicle* condition not significantly different. The +1 word was read significantly slower in *animal*, *concrete* and *non-concrete* conditions. At the +4 position, i.e., the fourth word after the object, reading times were only significantly different in the *vehicle* condition, where responses were significantly faster than in the baseline *human* condition.

6.4.2.2 Gradability results

For the gradability stimuli, separate models were produced using the reading times in milliseconds at the Modifier, Adjective, +1, +2, +3 and +4 positions as the dependent variables; Table 6.9 shows how these labels correspond to an example test sentence. Reading times at the +5 were again not analysed due to their significant length and variation (see Figure 6.4).

The full model was as follows. Adjective Type (gradable, non-gradable) and Modifier (swear, non-swear) were included as categorical predictors, as well as their interaction. Responses to the demographic survey and the swearing survey were also included as main effect predictors. Participant was included as a random intercept in all models; a random slope for Participant over Word was initially included where possible, but was dropped due to issues with model convergence. Word was only included as a random intercept in the Adjective and +5 models, as it was only at these positions that Word varied significantly. Modifier (*probably*, *fucking*) and Adjective Type (*gradable*, *non-gradable*) were sum coded. Predictors and interaction terms that did not significantly improve model fit based on chi-square comparisons of the sums of the squares of the residuals were excluded.

Non-gradable	I	realised	that	the	bookcase	was	fucking	wooden	and	that	it	had	grooves
Gradable	I	realised	that	the	bookcase	was	fucking	creaky	and	that	it	had	grooves
Labels							Modifier	Adjective	+1	+2	+3	+4	+5

Table 6.9: Experiment I: Gradability stimuli - sentence labels

The final models only included Modifier and Adjective type as main effect predictors. Tables 6.10 and 6.11 summarise these models. Cohen's *d* (Cohen, 1992) was calculated as above. Across all positions, there were no significant main effects for Modifier. There was a significant main effect for Adjective Type: Non-gradable at the +1 position ($\beta = -4.71$, p = 0.03, d = -0.0322).

Table 6.10: Experiment I: Model summary: Modifier, Adjective and +1 - Reaction Time $\sim$ Modifier + Adjective Type + (1|Participant) + (1|Word)

	Modifier			**Adjective**			**+1**		
Predictors	*Estimates*	*CI*	*p*	*Estimates*	*CI*	*p*	*Estimates*	*CI*	*p*
(Intercept)	429.57	413.91 – 445.24	**<0.001**	495.10	468.60 – 521.60	**<0.001**	454.71	441.29 – 468.14	**<0.001**
Modifier: Fucking	3.00	-0.95 – 6.95	0.136	4.33	-2.47 – 11.12	0.212	-0.13	-4.38 – 4.12	0.954
Adjective Type: Non-gradable	0.62	-3.33 – 4.57	0.758	6.36	-3.51 – 16.24	0.207	-4.71	-8.96 – -0.46	**0.030**
N	299 Participant			299 Participant 37 Adjective			299 Participant		
Observations	10581			10485			10443		

Table 6.11: Experiment I: Model summary: +2, +3, +4 - Reaction Time $\sim$ Modifier + Adjective Type + (1|Participant) + (1|Word)

	+2			**+3**			**+4**		
Predictors	*Estimates*	*CI*	*p*	*Estimates*	*CI*	*p*	*Estimates*	*CI*	*p*
(Intercept)	421.63	408.27 – 434.99	**<0.001**	412.48	400.38 – 424.58	**<0.001**	410.32	398.48 – 422.16	**<0.001**
Modifier: Fucking	1.09	-2.59 – 4.76	0.562	-1.68	-5.09 – 1.73	0.333	-0.62	-3.75 – 2.52	0.700
Adjective Type: Non-gradable	-2.29	-5.97 – 1.38	0.222	0.11	-3.30 – 3.52	0.949	0.63	-2.51 – 3.76	0.696
N	299 Participant			299 Participant			299 Participant		
Observations	10598			10483			10570		

Figure 6.4 plots the mean reaction time and two times standard error, both in milliseconds, for each condition at each position in the sentence.

A further model was constructed for summed reaction times of all words following the critical adjective. Reaction times were only totalled on trials where every data point was available. The full model includes Modifier and Adjective Type (both sum-coded) as main effect predictors, as well as their interaction. The interaction term was dropped, as it did not improve model fit. Random intercepts were included for Participant and

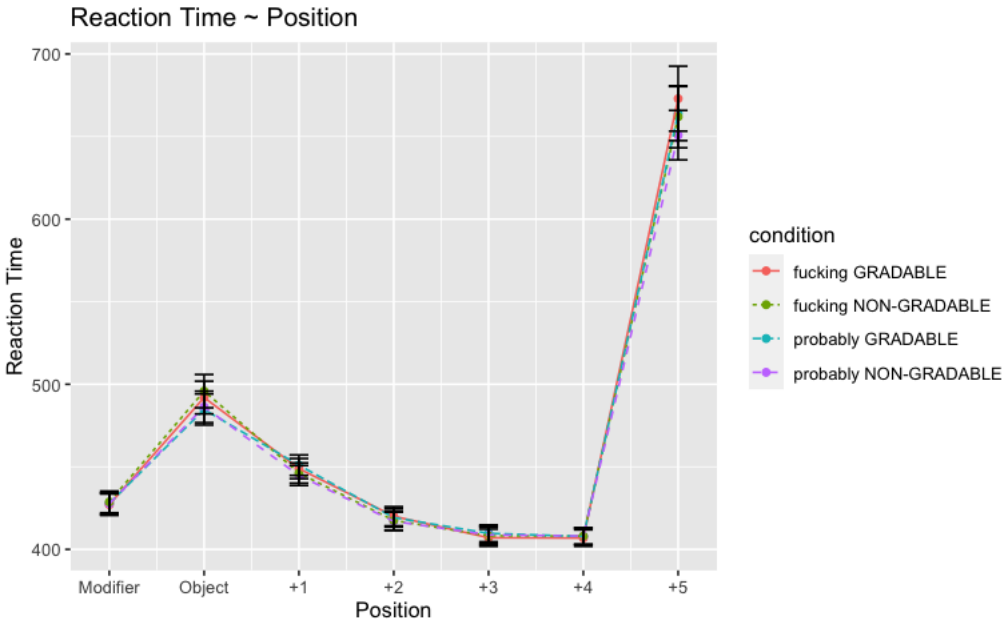

Figure 6.4: Experiment I: Gradability results

Adjective. Table 6.12 summarises the model. There were no significant main effects for either predictor (both $p > 0.2$).

Table 6.12: Experiment I: Model summary: +1, +2, +3, +4 - Summed Reaction Time ~ Modifier + Adjective Type + (1|Participant) + (1|Adjective)

	Summed Reaction Time		
Predictors	*Estimates*	*CI*	*p*
(Intercept)	2324.65	2253.39 – 2395.91	**<0.001**
Modifier: Fucking	1.27	-8.39 – 10.92	0.797
Adjective Type: Non-gradable	10.51	-7.27 – 28.30	0.246
$N_{Participant}$	298		
N_{word}	36		
Observations	9648		

6.4.2.2.1 Exploratory analysis Further exploratory analysis was conducted to examine responses for particular subsets of items and participants. The full dataset was first subsetted by adjectival pair into 18 separate datasets. The same model specified in the above analysis was constructed for each pair at each sentence position. Participant was included as a random intercept in all models. Out of the 18 adjectival pairs, a significant interaction effect for Modifier: Fucking * Adjective Type: Non-gradable was found for the following 5 pairs: *inconsistent*/*geographical*, *memorable*/*hexagonal*, *cheap*/*Swiss* (all positive estimates), *ancient*/*Chinese* and *ornate*/*Russian* (both negative estimates).

The full dataset was then subsetted by Trial Number. Due to the repetitive nature of

the task, it is possible that the effects of the manipulations were attenuated towards the end of the task. The full-data set was reduced to the first 50% of observations per participant. The same models specified above were run on this new dataset. No further significant effects were found in this dataset.

While the expected effects were not found for gradability, a number of effects were found for the animacy stimuli. In the exploratory analysis of the gradability data, participant behaviour on animacy stimuli was used to filter participants. The gradability data for participants whose mean reaction time for inanimate objects was greater than their mean reaction time for animate objects, indicating a slowdown in the expected direction, were included in a separate dataset (N = 179). The same models were run again on this new dataset. No further significant effects were found in this dataset.

6.4.3 Interim discussion

The hypotheses for Experiment I are repeated below:

H1 Sentences with swearwords that precede inanimate objects will be read slower than sentences with swearwords that precede animate objects.

H2 Sentences with swearwords that precede non-gradable predicates will be read slower than sentences with swearwords that precede gradable predicates.

The results provide support for one of the experimental hypotheses stated in Section 6.4.1.1, but not the other. Support is provided for the first experimental hypothesis regarding the property of *animacy*. Sentences with inanimate objects were read slower than sentences with inanimate objects. This effect was strongest at the second word after the object; it was only significant at the third word after the object when the modifier was *fucking* or *damn*, but not *bloody* (see the underlined portions of the example sentence in (42)). This might reflect the fact that, unlike *fucking* or *damn*, *bloody* can be ambiguous between expressive and descriptive meanings. For example, upon encountering a sentence containing the phrase *the bloody shoes*, some participants may have taken this to mean that the shoes were covered in blood. In these cases, *bloody* would have a local reading, which would explain the smaller slowdown effect compared to *fucking* and *damn*.

(42) I just saw that the fucking shoes/dogs were <u>on the</u> sofa again

As the further exploratory analysis in Section 6.4.2.1.1 suggests, however, the picture might be a bit more complex than this. While humans and animals were grouped together in the *animate* category, they had different effects on reading speeds at the +1 position, with animals leading to a significant slowdown compared to humans. At the +1 position,

animals appeared to pattern with the *inanimate* objects. Conversely, the vehicles, of which there were only two, did not differ significantly from the humans in their effect on reading speed, suggesting that they did not pattern with the other *inanimate objects*. Furthermore, further on in the sentences at the +4 position, the vehicles led to significantly faster reading times than did the *human* baseline condition. None of the other animacy categories differed from the baseline in their effect on reading speeds at this position. That is, only the vehicles patterned differently to the rest of the animacy categories. These results might suggest that vehicles also prompt local reading of expressive modifiers.

The results for the gradabilty stimuli do not provide support for the second experimental hypothesis. That is, sentences with non-gradable predicates were not read significantly slower than sentences with gradable predicates. The only significant effect found for *gradability* was a significant main effect, such that, across all sentences at the +1 position, reading times were slower for gradable adjectives. Importantly, these are not interaction effects. That is, these effects are not specific to *fucking*, but rather, both *fucking* and *probably*.

Given that *probably* was predicted to have a non-local reading with both gradable and non-gradable predicates, it is difficult to make much of these main effects; the effect that was found also occurred earlier in the sentence than was expected. One could speculatively hypothesise that, in the context of a sentence like (43-a), perhaps *probably* is coerced into having a degree modifier reading. This seems very unlikely however; there doesn't appear to be a reading of (43-a) that approximates *very tall*. If anything, *probably* would act as a downgrader. Even if this had been the case, returning to the observed main effect for gradability, this would mean that the local readings, i.e., sentences with gradable predicates, had led to a reading slowdown at the +1 position compared to those with non-local readings, i.e., sentences with non-gradable predicates. This would be contrary to what was predicted for these stimuli.

(43) a. The door is probably tall
b. The door is probably wooden

It is difficult to know exactly how *probably* was interpreted during the experiment however. It is possible that the observed effect is simply anomalous, or perhaps it reflects some other element of the stimuli. Rather than attempt to draw conclusions from the effect that has been observed, for which an explanation is unclear, I will instead treat this as a null result.

This experiment has provided partial support for the proposal that people *perceive* semantic/pragmatic variation in swearing. The results found for the animacy stimuli suggest that people have some implicit knowledge of the contrast between local and non-local swearing intensification based on the animacy distinction. It may therefore be possi-

ble that people also *notice* this difference when socially evaluating a speaker. Having established the animacy contrast in swearing intensification in Section 6.3.1 and found evidence for its effect on processing in the current section, the use of this variable in a sociolinguistic perception task is now well-motivated.

The same is not true for the gradability contrast established in Section 6.3.2. That is, I have not found evidence that the distinction between local and non-local swearing modification based on the gradability properties of the modified predicate is reflected in online processing. It does not appear that people *perceive* this distinction. There are several possible reasons. It may be that people do have an implicit knowledge of the distinction, but that a flaw in the experiment meant that it was not able to be captured; perhaps a self-paced reading task was insufficiently sensitive or maybe there were other confounding variables in the stimuli that prevented differences as a function of gradability from emerging.

Alternatively, it may be that it is in fact possible to generate a local reading of an expressive modifier that precedes a non-gradable adjective, albeit not a degree modifier reading. Rather, the modifier could be read as expressing heightened emotion towards the property denoted by the adjective. For example, in a sentence like (44), *fucking* could be interpreted as expressing heightened emotion towards *Swissness*, that is, the property of being Swiss. If it is true that expressive modifiers that precede non-gradable adjectives can receive a local reading, we would not in fact expect a reading slowdown for sentences containing such constructions relative to sentences in which expressive modifiers precede gradable predicates.

(44) The table was fucking Swiss and it was covered in scratches

Without doing further experiments, however, I cannot currently conclude from these results that people perceive the variable. Given that *perceiving* is a pre-cursor to *noticing* (Schmidt, 1990; Squires, 2016), it would therefore not be logical to test the effect of gradability on social perceptions of swearing. As such, in the following section on *noticing* socio-semantic variation in swearing, I did not test *gradability* as a factor in a matched-guise task. Rather, it is only *animacy* that will continue to be discussed.

6.5 Experiment II: Noticing socio-semantic variation in swearing

The experiment discussed in the previous section has established that people *perceive* the difference between local and non-local swearing intensification. This was done by testing perceptions of sentences containing swearwords that preceded either animate or

inanimate objects. As the participants were sensitive to this distinction while performing a relatively automatic process, namely reading, it suggests that they had created some internal representation of the distinction in their minds. It has not yet been established, however, whether people also *notice* the distinction. That is, is this internal representation indexically linked to particular social meanings and stances? In this section, I will describe the results of a matched-guise task (MGT) which suggest that it is.

6.5.1 Methods

Experiment II used a visual MGT. Participants were presented with stimuli made up of two consecutive sentences from one speaker, one of which contained the experimental manipulations of Swear (*fucking* vs *damn*) and Condition (*animate* vs *inanimate*). In light of *bloody* having the smallest effect in Experiment I, it was dropped from Experiment II, with only *fucking* and *damn* retained; the experimental design for Experiment II is fully schematized using examples in Table 6.13. Participants read a series of these sentences and rated the speaker on Likert scales generated through a pre-experiment task.

	Inanimate	**Animate**
Fucking	fucking paint	fucking kids
Damn	damn paint	damn kids

Table 6.13: Experiment II: Experimental design

It is important to note, before continuing, that the exploratory analysis reported in Section 6.4.2.1.1 on the influence of the animacy hierarchy on reading times was conducted after Experiment II had already been completed. As such, the stimuli were not designed with this exploratory analysis in mind. Given that the animate/inanimate distinction originally used to design stimuli Experiment I did affect reading speeds in two positions, regardless of which animacy category the objects came from, the use of this distinction in Experiment II is nonetheless well-motivated.

6.5.1.1 Experimental hypothesis

The experimental hypothesis for this experiment was as follows:

H1 Speakers uttering sentences in which a swearword is modifying an animate noun will be evaluated differently on the chosen Likert scales than speakers uttering sentences in which a swearword is modifying an inanimate noun.

This hypothesis was admittedly vague. This reflects the fact that there was no clear prediction to be made regarding which type of construction would be judged more or less favourably by participants. On the one hand, one might predict that sentences in the *animate* condition would be judged less favourably. Animate objects, unlike inanimate objects, have feelings. As a result, as the target of the swearword in local intensification, there is an individual that might take offence to the swearing and/or the attitude expressed by the swearing, in addition to the person hearing the sentence. On the other hand, given that, in the *inanimate* condition, the speaker was expected to be interpreted as expressing emotion about a state of affairs rather than to express their anger towards a specific culprit, their swearing may be evaluated as gratuitous. Why bother swearing if you can't hold a particular individual responsible for their actions? In the absence of a clear evidence in support of either prediction, I simply hypothesized that evaluations of sentences in the two experimental conditions would differ.

6.5.1.2 Stimuli

The stimuli for Experiment II were predominantly similar to those used in Experiment I to test the effect of animacy on sentence processing. The 15 existing animacy stimuli were modified minimally for Experiment II. As the constraints on stimuli construction in an MGT are less strict than for a self-paced reading task, the stimuli were changed to create greater variety between each sentence. This was done to make the sentences sound more natural and to distract participants from the experimental manipulations. Every test sentence was preceded by another sentence, setting up a context in which the test sentence would sound like a natural follow-up. A full list of the 16 sentences in the animate condition is included in Appendix B.5, while the same sentences in the inanimate condition are included in Appendix B.6.

To allow for counterbalancing and randomisation, an additional sentence was included; this sentence is provided as an example in (45). This brought the number of test sentences up to sixteen. Exhausting every combination of the experimental manipulations (Condition and Swear) created a total of 64 test stimuli. These stimuli were divided into 4 counterbalanced lists of 16 stimuli. Each list contained exactly one instance of each test sentence and four sentences in each experimental condition, that is, 4 x *fucking* ANIMATE, 4 x *fucking* INANIMATE, 4 x *damn* ANIMATE and 4 x *damn* INANIMATE.

(45) I've been stuck at home all day. The MOD *cold/plummer* made things really uncomfortable.

In addition to the 16 test stimuli, each list contained the same 15 filler items. Filler items also contained two sentences of a similar length to test sentences. The filler items came

in three categories, one containing other swearwords (e.g. *shit*, *dickhead*), one containing positive sentences and one containing negative sentences. Finally, an additional catch trial was included to allow for the removal of participants who were not fully concentrating. The sentence in the catch trial read 'THIS IS A TEST. Please put all of the scales to the far right'. Failure to comply with this instruction resulted in a failure of the catch trial.

6.5.1.3 Selecting the scales

As in Experiment II of Chapter 5, the scales for this task were constructed using a separate task with a different group of participants. 101 participants completed a survey hosted on Qualtrics (Qualtrics, 2013). In the survey, participants were presented with two sets of sentences containing shorter versions of three test sentences with the intensifier *fucking*, listed in (46) and (47).

(46) a. The fucking chef burnt my steak all over
b. My fucking neighbour tore down my tree
c. The fucking kids stained the carpet

(47) a. The fucking heat burnt my steak all over
b. My fucking snowstorm tore down my tree
c. The fucking paint stained the carpet

Set 1 (46) contained sentences with animate nouns, while Set 2 (47) contained sentences with inanimate nouns. For each set, participants were asked to provide four adjectives to describe a possible speaker of those sentences. All participants saw both sets, one set at a time, with order of presentation randomised between participants.

After providing four adjectives for each set, participants were asked two follow-up questions. First, participants were presented with both sets simultaneously and asked 'Do you see a difference between the two groups of sentences? If so, what is it?'. They were then asked 'In hindsight, would you associate different adjectives with the two groups? If so, what would you change?'. These questions gave participants a chance to reflect on their choices, with the two sets directly contrasted.

Figure 6.5 plots the frequency of each unique adjective, grouped by whether the sentences contained animate or inanimate objects and limited to the fifteen most-selected adjectives for each set, using the *ggplot2* package (Wickham, 2016). Two observations were removed, as they were repeated adjectives by the same participant for the same set. Results for the two sets were broadly similar, with subtle differences highlighted in colour. While *angry*, *rude* and *aggressive* were highly frequent for both sets, the frequencies for these adjectives were higher for the set with animate nouns. Similarly, while *annoyed* was highly frequent for both sets, it was marginally more frequent for the inanimate set.

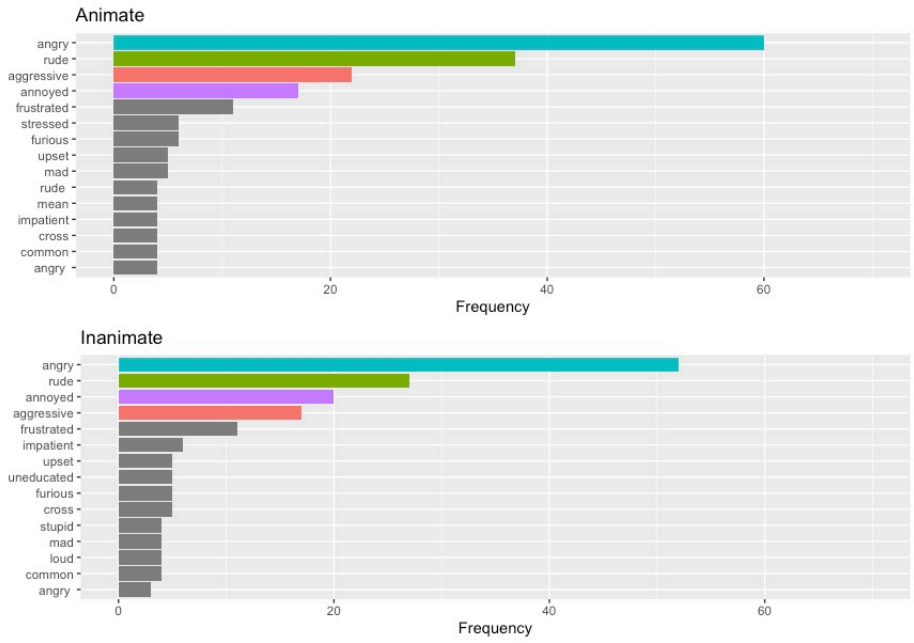

Figure 6.5: Experiment II: Scale selection adjective frequencies

In the follow-up questions, the vast majority of participants indicated that they noticed a difference between the two sets, correctly identifying the presence of an *individual* or *person* to blame in the animate sets. In response to the second question, while many participants stated that they would not change their answers in retrospect, a number of participants did. A selection of these are listed below in (48). A number of responses suggested that the sentences with animate subjects would be associated with a speaker that is more *annoyed*, *menacing*, *mean*, *rude*, *aggressive* and *arrogant*, in-line with the marginal differences illustrated in Figure 6.5.

(48)
a. i would change my answers calling the person rude when blaming factors out of someones control
b. Speaker 1 is more arrogant and blames other people
c. They have the same kind of enraged feeling but maybe the second is not as strong as the first.
d. I would change the adjective rude so it is not associated with the second
e. the first set is directing blame at a person whereas the second set is annoyed at the way things have happened but isn't blaming anyone so it is less aggressive
f. first group is definitely rude since it refers to a person, second would just be improper
g. In hindsight the first sentences are quite comical whereas the second are more menacing because they are directed at people.
h. I would no longer associate mean adjective as you cannot be mean to inani-

mate objects.

As in Chapter 5, this process is not completely objective and requires a degree of human input and consideration. The differences in adjective frequencies across the two sets is small. Taken together with the responses to follow-up questions, however, it appears that participants associated the sentences with animate objects with a greater degree of anger, rudeness and aggression and a slightly lesser degree of annoyance. I therefore included these four characteristics as test scales. In addition, two dummy scales were included, namely *tall* and *attractive*.

Regarding the experimental hypothesis in Section 6.5.1.1, while the results of this scale selection task suggest that swearing intensification of animate objects may be perceived as ruder, angrier and more aggressive, the differences between groups in the task were not statistically significant following posthoc z-tests in R (all $p > 0.05$); the results of this task therefore only suggest which attributes are most likely to be relevant to participants when reading test stimuli.

6.5.1.4 Procedure

This task was constructed and hosted on Qualtrics (2013) and distributed online via Prolific Academic (Prolific, 2022). After giving informed consent, participants were given the instructions in (49).

(49) *In this task, you are going to be reading a series of sentences - there will be about 30 in total. For each sentence, we would like you to imagine the type of person you think would say this sentence. With this person mind, we want you to evaluate them using a series of scales. An example scale is included below. For example, if you think the person that would say this sentence would be very funny, you might put a response closer to the word 'Funny'.*

Participants were assigned to one of the four experimental lists. The order of trials was fully randomised between participants. On each trial, participants were presented with the test stimuli, at the top of the screen, and the six test scales. The order of presentation of the scales was randomised between trials. Participants were required to provide a response on all scales in order to continue to the next trial.

After completing all 32 trials, participants completed two brief surveys. The first asked participants about their own swearing behaviour and attitude towards swearing - specific to this task, they were asked how frequently they used the word *fuck* and the word *damn*. They then completed a demographic survey eliciting their gender, age, yearly household income (in blocks of 10,000), highest level of education and self-defined social class. Participants also completed a post-task debrief, before these surveys, to find out

whether they could guess what the aim of the survey was; this was done using a text entry box.

6.5.1.5 Participants

150 British English speakers were initially recruited and paid £1.17 for their participation; 4 participants failed catch trials, bringing the total number of participants down to 146.

6.5.2 Results

Results for the 146 remaining participants were analysed in R using linear mixed effect regression models using the *lme4* package (Bates et al., 2015). The dependent variable in each model was the response on each of the test Likert scales. The full model was as follows. The categorical predictors were Condition (Inanimate, Animate) and Swear (Fucking, Damn), both sum-coded, and their interaction, as well as the participant demographics age, gender, social class, level of education and income and participant responses to the post-experiment surveys. Random intercepts were included for Participant and Noun; a random slope for Participant over Noun was initially included, but this was dropped due to convergence issues. Predictor variables that did not significantly improve model fit based on chi-square comparisons of the sums of the squares of the residuals were dropped. The only remaining predictors for all models were Condition and Swear, with the interaction term only kept for *rude* and *angry*.

Table 6.14 summarises the models for *rude* and *angry*. For *rude*, there were significant main effects for Condition: Animate ($\beta = 0.18$, p < 0.01) and Swear: Damn ($\beta = -0.36$, $p < 0.001$). There was also a significant interaction effect for Condition: Animate * Swear: Damn ($\beta = 0.05$, $p < 0.05$). For *angry*, there was a significant main effect for Swear: Damn ($\beta = -0.27$, $p < 0.001$) and a significant interaction effect for Condition: Animate * Swear: Damn ($\beta = 0.05$, $p < 0.05$). The means and two-times standard error by Condition and Swear are plotted for *rude* and *angry* in Figure 6.6.

Table 6.14: Experiment II: Model summary for *rude* and *angry*: Response ~ Condition * Swear (1|Participant) + (1|$Noun$)

	rude				angry			
Predictors	*Estimates*	*std. Beta*	*Statistic*	*p*	*Estimates*	*std. Beta*	*Statistic*	*p*
(Intercept)	5.74 ***	-0.00	62.83	**<0.001**	6.30 ***	-0.01	72.44	**<0.001**
Condition: Animate	0.18 **	0.13	2.80	**0.005**	0.03	0.03	0.50	0.617
Swear: Damn	-0.36 ***	-0.26	-16.97	**<0.001**	-0.27 ***	-0.22	-13.70	**<0.001**
Animate * Damn	0.05 *	0.04	2.45	**0.014**	0.05 *	0.04	2.37	**0.018**
N	146 ResponseId				146 ResponseId			
	32 noun				32 noun			
Observations	2336				2336			

** p<0.05 ** p<0.01 *** p<0.001*

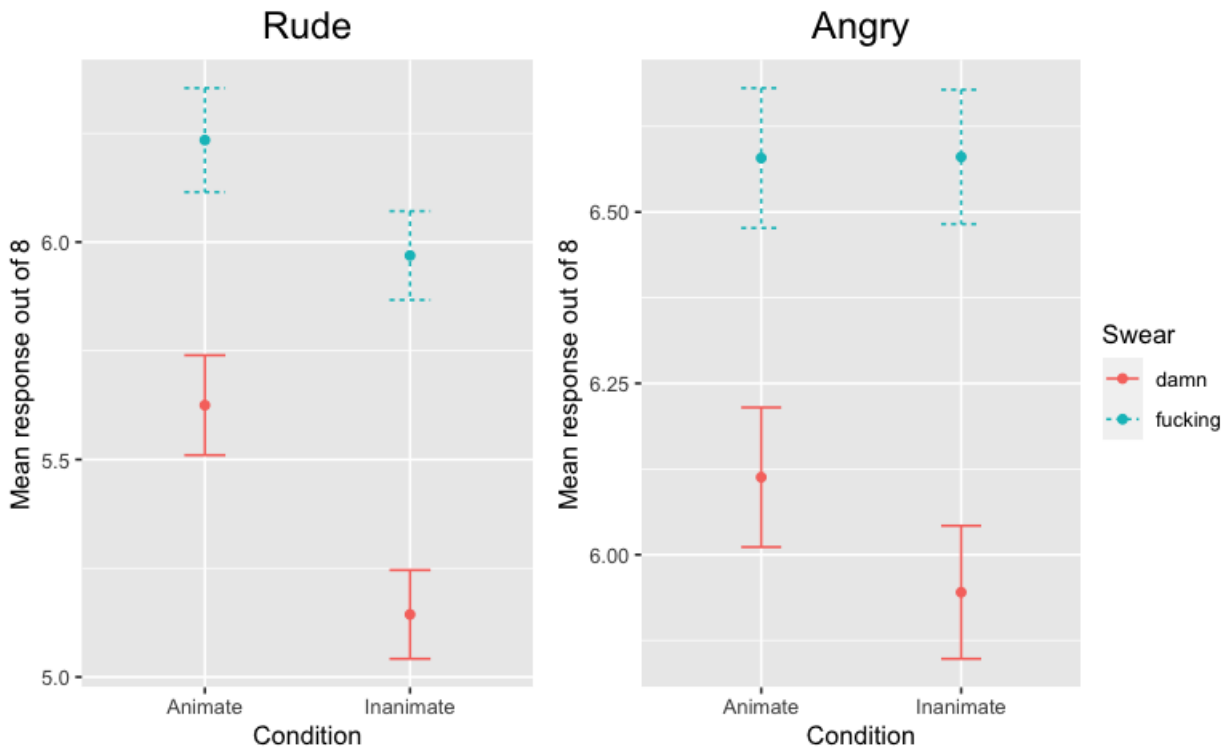

Figure 6.6: Experiment II: Mean response for *angry* and *rude* by Condition and Swear

Table 6.15 summarises the models for *annoyed* and *aggressive*. For *annoyed*, there was a significant main effect for Swear: Damn ($\beta = -0.18$, $p < 0.001$). For *aggressive*, there was a significant main effect for Swear: Damn ($\beta = -0.37$, $p < 0.001$). There were no significant effects for Condition in either model. The means and two-times standard error by Condition and Swear are plotted for *annoyed* and *aggressive* in Figure 6.7.

Table 6.15: Experiment II: Model summary for *annoyed* and *aggressive*: Response Condition + Swear (1|Participant) + (1|$Noun$)

	annoyed				aggressive			
Predictors	*Estimates*	*std. Beta*	*Statistic*	*p*	*Estimates*	*std. Beta*	*Statistic*	*p*
(Intercept)	6.67 ***	-0.00	77.49	**<0.001**	5.83 ***	-0.00	72.45	**<0.001**
Condition: Animate	-0.00	-0.00	-0.05	0.957	0.09	0.07	1.64	0.101
Swear: Damn	-0.18 ***	-0.15	-9.21	**<0.001**	-0.37 ***	-0.28	-17.66	**<0.001**
N	146 ResponseId				146 ResponseId			
	32 noun				32 noun			
Observations	2336				2336			

** p<0.05 ** p<0.01 *** p<0.001*

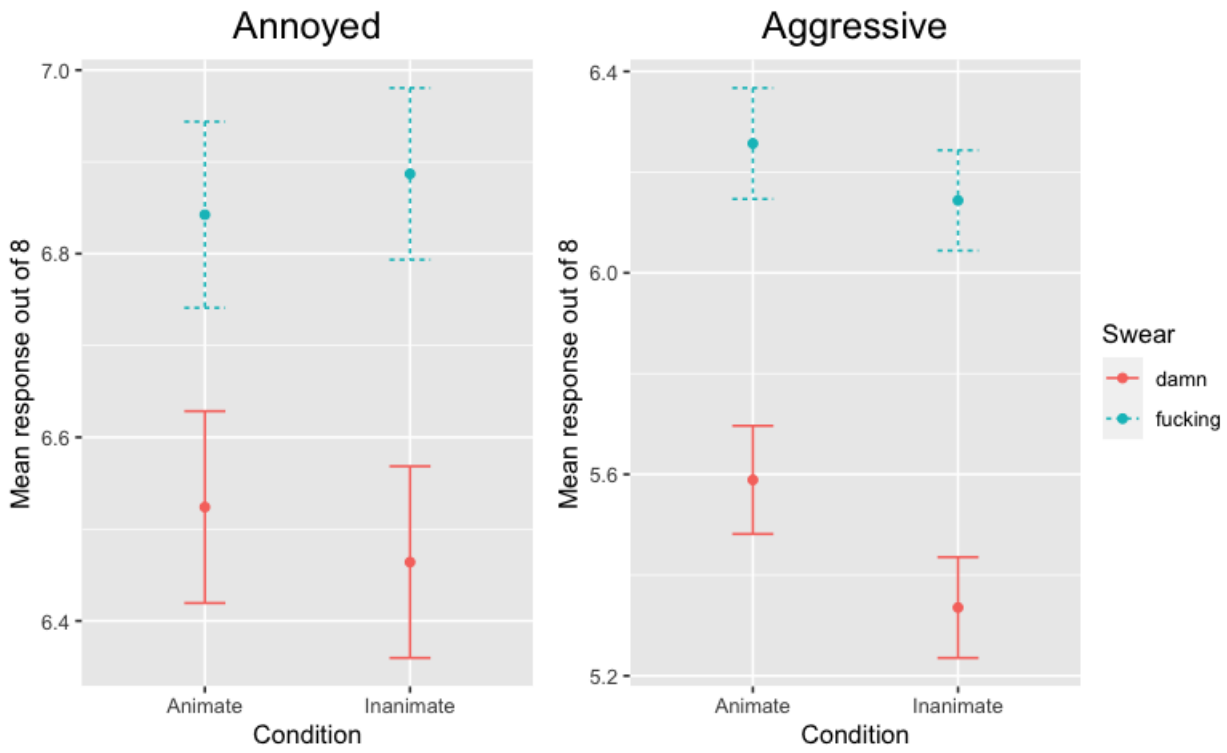

Figure 6.7: Experiment II: Mean response for *aggressive* and *annoyed* by Condition and Swear

6.5.3 Interim discussion

The hypothesis for Experiment II is repeated below:

H1 Speakers uttering sentences in which a swearword is modifying an animate noun will be evaluated differently on the chosen Likert scales than speakers uttering sentences in which a swearword is modifying an inanimate noun.

Support is provided for the experimental hypothesis on two of the four test Likert scales. When the 'speaker' uttered a swearword immediately prejacent to an animate noun (e.g., *kids*, *chef*, *dog* etc), they were evaluated as significantly ruder and, in the case of *damn* but not *fucking*, angrier, compared to when the 'speaker' uttered a swearword immediately prejacent to an inanimate noun (e.g., *paint*, *heat*, *shoes* etc). The plot in Figure 6.7 seems to suggest that responses for the *aggressive* scale were similar to those for *angry*, but this is misleading. Although a model for *aggressive* with an interaction term resulted in an interaction effect that approximated significance ($p = 0.078$), the inclusion of this interaction term did not significantly improve the fit of the model, which also included by-participant and by-item random intercepts; the main effect predictor of Condition for *aggressive* only had a p-value of 0.101

That *damn* and *fucking* received significantly different responses across participants is unsurprising. Compared to *fucking*, *damn* is frequently rated as less taboo, less arousing and less offensive (Janschewitz, 2008; Beers Fägersten, 2012). It therefore makes sense that the use of *fucking* was rated as ruder, angrier, more aggressive and more annoyed.

In interpreting the effect for *rude*, it is worth thinking about how *rude* can be conceptualised. On the one hand, it could relate to how rude the uttering of a particular word is to the person hearing it. Simply because words like *fucking* and *damn* are taboo - because they relate to the societally taboo topics of sex and religious profanity - uttering them in any context might be considered rude. This is similar to the *immediacy* property of expressives (Potts, 2007; see Section 2.3.3.1 of Chapter 2).

Another possible conceptualisation of *rude* concerns the entity to which the swearword appears to be directed; that is, participants may have also considered how *rude* the speaker was being to some other individual referred to in the sentence, rather than just how rude they were being to the hearer of the sentence. In the opening sections of this chapter, I detailed the CULPRIT HYPOTHESIS (Frazier et al., 2015), which suggests that when a swearword modifies an animate noun, the most likely reading is one under which the swearword applies directly to that noun. In contrast, when a swearword modifies an inanimate noun, the more likely reading is under which the swearword applies to the whole state of affairs described by the sentence.

If participants did interpret *rude* in this second way, then this would be a logical way of explaining the effects observed for *rude*. In the Animate condition, *damn* and *fucking* would have been read as being directed towards the animate noun. This, in turn, would have been evaluated as being *rude* towards the individual denoted by that noun. For example, in (50-a), it would be seen as *rude* to the daughter to refer to her using *fucking*. Importantly, because the *daughter* is animate, she can feel some emotion about being referred to in this way; for example, she could feel hurt or offended.

(50) a. Our fucking daughter smelt of something really bad
b. Our fucking carpet smelt of something really bad

In contrast, there is no individual to offend in (50-b). In the Inanimate condition, the *damn* and *fucking* would most likely be read as being directed towards the state of affairs. In (50-b), this would be 'the carpet smelling of something really bad'. There is also an available reading where the *fucking* applies directly to the carpet, as it did to the daughter in (50-a). Importantly, neither the state of affairs nor the carpet can feel some emotion towards the use of a swearword; that is, they cannot feel hurt or offended.

It may be that both interpretations of *rude* were used by participants. Importantly however, it was only in the Animate condition that both interpretations were available. That is, only in those cases can *fuck* or *damn* be viewed as *rude* towards both the referent and the hearer; in the Inanimate condition, the use of a swearword can only be *rude* towards the hearer. This is a plausible explanation for the observed effect.

For *angry*, the same logic may hold for *damn*. In the Animate condition, the use of *damn* may have indicated the speaker's anger towards the animate referent. In the Inanimate condition, the anger would be directed either towards the state of affairs or, potentially, the inanimate referent. For *fucking* however, the lack of a significant effect for Condition may indicate a ceiling effect. That is, uttering *fucking* may already be so indicative of anger that responses immediately moved towards the top end of the scale. As a result, responses on this scale may not have been sensitive to the more subtle variation in animacy. The same may not have been true for *damn*, to which reactions were more variable.

That *annoyed* and *aggressive* did not follow the same trend suggests that participants thought of these as distinct concepts from anger - this perhaps justifies their selection for the scales. Intuitively, *annoyed* seems like a weaker emotion than *anger*. The difference between *angry* and *aggressive* concerns the action taken by the speaker. The Oxford English Dictionary defines *aggressive* as 'feelings of anger or antipathy resulting in hostile or violent behaviour'. Being *aggressive* might include *angry*, with the addition of some concrete behaviour as a result of that anger. It may be that, while participants considered speakers to be angrier when using *damn* before animate objects, this did not lead them to expect the speaker to be violent towards that object. This would explain the lack of a significant effect for *aggressive*.

With regard to the effects that I observe for *rude* and *angry*, it is worth noting two caveats, both regarding the experimental stimuli. The first concerns the nature of the individuals in the animate condition. A number of these individuals were family members (e.g., *daughter*, *son*, *kids*, *dad*) or people working in roles that might have lead readers to sympathise with them (e.g., *maid*, *teacher*, *nurse*). Had the individuals been people

in positions of power over the speaker (e.g., *boss* or *landlord*) or worked in roles with which people do not tend to sympathise (e.g., *taxman*, *billionaire*), it might not have been seen as so rude to express heightened emotion towards them. The second concerns the nature of the 'unwanted situation' for which that individual is believed to be responsible. Some of these are relatively minor inconveniences such as lying on a bed, being in the kitchen or making noise. If the individuals had caused something more serious to occur, such as stealing something or hurting someone, the swearing may have been considered more warranted and, therefore, less rude. Both of these are speculative, but are worthy of consideration in light of the explanation I have offered for the results of Experiment II.

On a final note, in the post-experiment debrief, the majority of participants suggested that the experiment had been designed to test perceptions of swearwords, with a significant minority stating suggesting that the experiment was about height and attractiveness as it relates to other social traits; this latter group were clearly influenced by the dummy scales. Only one participant displayed some conscious awareness of the experimental manipulations, stating: "Are you testing the perceived relative severity of swearing about people (of varying degrees of closeness) and about natural occurrences?". This suggests that, with one exception, the results of the experiment reflect implicit rather than explicit knowledge.

6.6 Chapter discussion and conclusions

Returning once again to the overall aims of this thesis - to examine whether people are sensitive to language-internal factors in the social evaluation of swearing - the findings of this chapter provide significant insight from the perspective of semantics and pragmatics. I have shown that animacy is a linguistic variable that can influence social perception of a speaker as a function of swearing. I did this in a logical fashion following the methodological steps outlined in Chapter 3.

Following the empirical observations that have been made about the flexibility of interpretation of expressive modifiers (Potts, 2005, 2007), I used previous work by Frazier et al. (2015) and my own original empirical observations to draw categorical distinctions between two sets of constructions, with the common observation that local interpretations of swearwords were available, and more likely, with animate objects and gradable adjectives, compared to with inanimate objects and non-gradable adjectives. For non-gradable adjectives, the only available interpretation was predicted to be the non-local interpretation, while the non-local interpretation was predicted to be the most likely interpretation for inanimate objects.

With these distinctions sketched out, I tested these constructions in a self-paced reading task to see whether there was a processing cost associated with the non-local readings

of swearing intensifiers. This turned out to be the case for the animacy stimuli, for which there was a significant reading slowdown at two and three words after the critical noun, but not for the gradability stimuli. Exploratory analysis also suggested a more complicated pattern in the animacy data at other points in the sentence when the categories of animate and inanimate nouns were broken into smaller groups following Zaenen et al.'s (2004) animacy schema.

As well as motivating the use of the animacy stimuli in a social meaning experiment, these results also have consequences for the literature on sentence processing. They provide further evidence for the CULPRIT HYPOTHESIS proposed by Frazier et al. (2015), suggesting that the internal representation people have of the likelihood of there being a possible culprit about which a person is swearing affects their processing of a sentence. While some swearing modifiers can indeed apply flexibly across a sentence, allowing for multiple different interpretations, people have an awareness, potentially implicit, that some interpretations are more likely in certain contexts. This also furthers the finding of Donahoo (2019) for the processing of expressive adjectives, showing that the linguistic context in which the expressive adjective is used can also affect processing.

For the gradability stimuli on the other hand, no slowdown was found. While this could suggest that the empirical observations I made in Section 6.3.2 are false, I believe that this more likely stems from issues with the experimental design and/or the stimuli. Without positive evidence for a processing cost associated with swearing intensifiers that modify non-gradable compared to gradable adjectives, however, I had no evidence to suggest that people *perceived* a difference between the two constructions. Without evidence of perception, I had no clear motivation to test whether people *noticed* that same difference.

For constructions with animate and inanimate nouns, however, I did find evidence of a difference in perception. With that in mind, I used a visual matched-guise task to test whether people would also *notice* the difference in social evaluation of the speaker. Despite significant variability in rest of the sentence - with regards to tone and affect - the results suggest that the participants did *notice* the difference. In terms of social indexicality (Silverstein, 2003; Eckert, 2008), the constructions in which swearing intensifiers modified animate nouns were associated with the socio-pragmatic meanings *rude* and *angry*.

While the existing literature suggests that swearing is broadly associated with rudeness and the expression of extreme emotions, this result suggests that this association is modulated by a hearer's interpretation of what the target of that emotion is. Returning to some of the points made in Section 2.3.3 of Chapter 2 on semantics, while expressives are homogenous in ways that differentiate them from neutral words, they are also heterogeneous when it comes to usage. Although the key difference in the animacy constructions

is pragmatically driven, rather than being a difference in semantic composition, the point still stands that the linguistic context in which a swearword is used can affect the social meanings it is perceived to index.

Chapter 7

Conclusion: Swearing and the sociolinguistic variable

7.1 Introduction

In this chapter I discuss the findings of the preceding three chapters. I first summarise my findings using indexical fields (Eckert, 2008). The aim of this is to show that 'swearing', rather than representing one abstract concept associated with social meanings that are precise or fixed, is actually associated with a constellation of ideologically related meanings. As I have shown in the preceding three chapters, the activation of specific meanings depends not only on the context in which a swearword us used, as has already been shown (Martin, 1997; T. Jay & Janschewitz, 2008; Howell & Giuliano, 2011; Jacobi, 2014), but also on the way in which a swearword is pronounced (Chapter 4), the way in which a swearword is combined with other words (Chapter 5), and the meaning contribution a swearword is understood to be making in a particular sentence (Chapter 6). That is, as well as depending on language-external factors, the social perception of swearwords also depends on language-internal factors. In this chapter, I will make this finding more explicit using indexical fields.

In this chapter, while I am drawing on the indexical fields used in previous work in sociolinguistics, many of which have differed significantly from one another, I have chosen to draw indexical fields that maximise understanding for each specific finding. Each representation of an indexical field that I have provided could have been drawn in a different way or in a different shape based on the elements of the relationship between swearing and a particular sociolinguistic variable that one wanted to emphasize. The ones I have chosen to present are, in my opinion, the best possible illustrations of the findings I have made throughout this thesis.

The rest of this chapter proceeds as follows. In Section 7.2, I recap the concept of the

indexical field (Eckert, 2008) by reviewing some example indexical fields from elsewhere in sociolinguistics. In Section 7.3, I present original indexical fields that I have drawn in order to contextualise some previous studies on swearing and social meaning. In Section 7.4, I present original indexical fields to contextualise my own results for swearing infixation (7.4.1), swearing modification of animate and inanimate objects (7.4.2), and swearing and variable (ING) (7.4.3). In Section 7.5, I discuss the consequences of these findings for swearing research. In Section 7.6, I discuss the limitations of my findings, before suggesting directions for future research in Section 7.7. I then wrap up my discussion in Section 7.8.

7.2 The indexical field: a recap

In Chapter 2, I introduced Eckert's (2008) conceptualisation of the *indexical field* as a way of modelling the collection of social meanings that can be activated by the situated usage of a linguistic form. This concept is dominant among sociolinguists working in the variationist tradition. It allows sociolinguists to account for the context-dependent emergence of related social meanings, without committing them to stating fixed meanings for variables. Figure 7.1 is an example of an indexical field taken from Eckert (2008), based on the results of Campbell-Kibler (2007) for variable (ING).

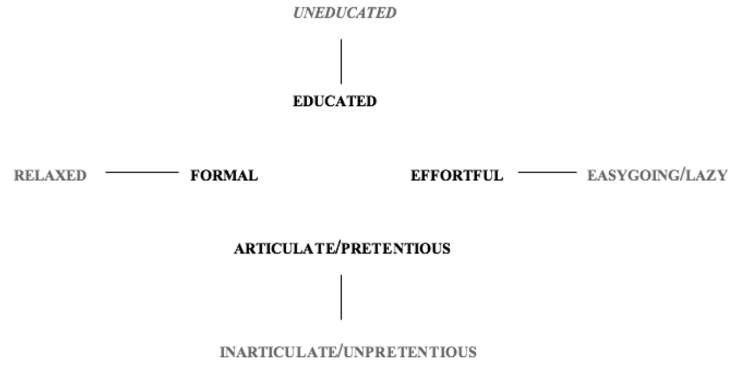

Figure 7.1: Indexical field of (ING. Black = meanings for the velar variant, gray = meanings for the apical variant). From Eckert (2008).

In her matched-guise studies, Campbell-Kibler (2005, 2007, 2010a) found that listeners broadly associated the velar [ɪŋ] variant with being educated, articulate/pretentious and formal, and, in opposition to the alveolar [ɪn] variant, effortful. In contrast, she found that listeners associated the alveolar variant with being uneducated, inarticulate/unpretentious, relaxed and easygoing/lazy. Importantly, as Eckert (2008) emphasizes, one cannot assume that all of these social meanings are activated by every use of a particular variant. For example, in any particular context, it may be that using the alveolar variant leads a

person to be perceived as formal and pretentious, but not educated. As Campbell-Kibler (2007, 2010a) shows, the emergence of particular social meanings also depends on the expectations that the listener has about the speaker and on the co-presence of other socially meaningful linguistic variables (Campbell-Kibler, 2011). This is why the social meanings in an indexical field are *potential* meanings.

Eckert (2008) sketches out another example for a potential indexical field, this time for /t/-release in American English, based on the findings of Bucholtz (2001), Benor (2001), Podesva, Roberts, and Campbell-Kibler (2002) and Podesva (2008) (see Figure 7.2). This indexical field includes a mix of stances, permanent qualities and social types. Stances involve in-the-moment social meanings that a person may momentarily index, but the repeated indexing of which could lead to a person being consistently associated with these qualities. The social types, such as *nerd girl* or *gay diva*, may be indirectly indexed by speakers via the other qualities in the field and the association of these social meanings with the socially salient qualities of such personae or stereotypes. For example, in Section 2.4.3.1 of Chapter 2, I discussed the findings of Eckert (1989) in a Detroit high school, whereby the *burnouts* group frequently used non-standard variants to index particular traits that they associated with people from the urban centre of Detroit, such as being tough and street smart, rather than directly aligning themselves with those people.

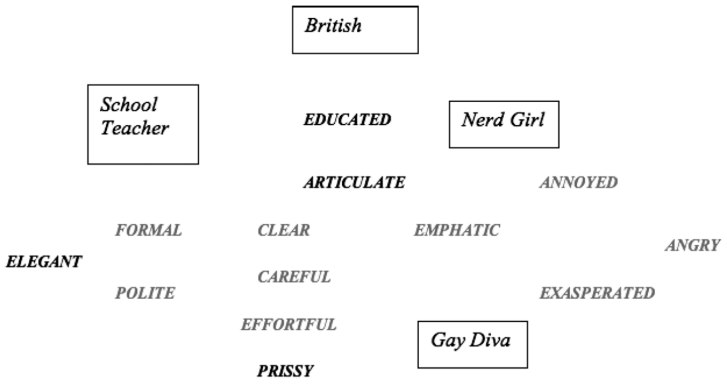

Figure 7.2: Indexical field of /t/ release. Boxes = social types, black = permanent qualities, gray = stances. From Eckert (2008).

Indexical fields have previously been used to model the social meanings of lexical variation with similar expressive meaning to swearing. For example, in their study of the Brazilian Portuguese slur *favelado*, which roughly translates as 'slum-dweller', Beaton and Washington (2015) provide the indexical field in Figure 7.3, with the positive and negative evaluations of *favelado* separated by a dotted line. In a diversion from Eckert (2008), Beaton and Washington (2015) suggest that, unlike phonetic variation, lexical variation "function[s] primarily to index things in the world rather than the attitude or identity of the speaker" (p. 16). What their indexical field aims to capture are the qualities

the speaker sees in a referent, rather than qualities that are typically associated with the speaker. The negative use indexes a variety of social meanings, visible on the left side of Figure 7.3. While *favelado* has typically been negatively associated with the types of people that typically inhabited the *favelas*, it has become reclaimed in a particular situated context, in relation to the Rio de Janeiro soccer team *Flamengo*. In response to opposition fans using *favelado* as an insult, the fans of Flamengo have since embraced the term in their own chants. In doing so, despite few fans actually living in the *favelas*, the fans are temporarily aligning themselves with the perceived characteristics of 'slum-dwellers', such as being street-smart and badass; both of these meanings are ideologically related to the negative associations with *favelado*, but in the context of the football team they are ameliorated. In this way, the slur can be used to index qualities of the speaker as well, such as membership to a particular group and qualities with which they wish to align themselves.

Figure 7.3: Indexical field for *favelado*. Lower case = permanent qualities; upper case = social types; dotted line separates positive and negative qualities and social types. From Beaton (2015).

Christie (2013) applies the concept of the indexical field to swearing in her analysis of media reactions to swearing (see Section 2.4.3.2 of Chapter 2). The swearing of footballer Wayne Rooney, for example, was linked to the social roles 'football player' and 'industrial worker', the latter of which is a social class stereotype linked to particular traits. He was simultaneously evaluated as emotional, angry, immoral and responsible for swearing in front of people who, according to some media outlets, ought not to hear swearing e.g., women and children. Similarly, Christie (2013) shows that the swearing of actress Gwyneth Paltrow was deemed a failed attempt at seeming "edgy and cool", with the media perceiving her actual social identity to be incongruent with the "bad girl" persona she was attempting to index. Her use of swearing was deemed inauthentic and inappropriate due it is lack of spontaneity. The same was argued for Stephen Fry's swearing at the British Academy Film and Television Awards by Stapleton (2020), with swearing again linked with a group to which Fry does not belong (i.e., young people).

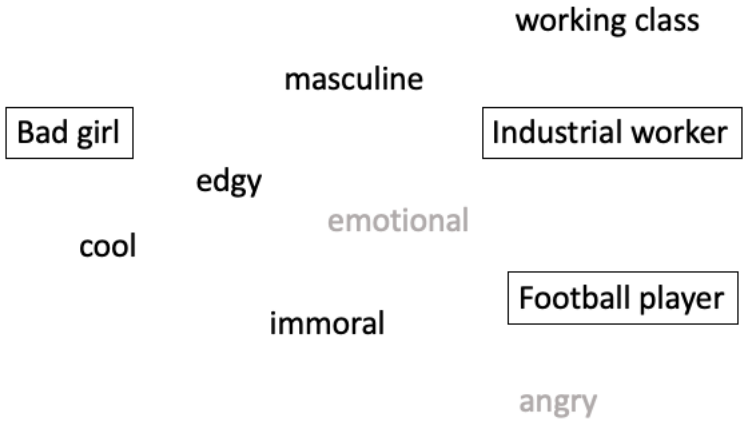

Figure 7.4: Indexical field of 'swearing': Boxes = social types, black = permanent qualities, grey = stances. Based on the findings of Christie (2013).

Just as Eckert (2008) did for the work of Campbell-Kibler (2007) on variable (ING), I have drawn a potential indexical field for 'swearing', broadly construed, based on the findings of Christie (2013). This indexical field includes the in-the-moment stances (e.g., *angry*, *emotional*), more permanent qualities (e.g., *working-class*, *masculine*) and social types (e.g., *industrial worker*) that appear to be socially relevant dimensions in the metadiscourse of swearing that Christie analyses. Many of these social meanings are ideologically related, but their mutual activation it not guaranteed. Just as with /t/-release, situated uses of swearing are likely activate only a subset of these meanings, depending on a variety of speaker- and context-specific factors.

I will focus for a moment on the 'industrial worker' social type. The quote from which this social type has been taken is provided in (1). In her analysis, Christie (2013) explains that this social type is related to socio-demographic group memberships, such as *working-class* and *men*; in my version, to reflect the way in which people perform gender, I have changed men to *masculine*.

(1) The FA (football association) did not act, despite the incident being highlighted, because they understand the football field to be a workplace and industrial language is part of its currency (The Daily Mail, 2011)

Tied up with this are other evaluations of Rooney's swearing as *angry* and, in one case, *violent*. These represent situated stances that the media perceived Rooney to be indexing. This association contributes to the indexing of the 'industrial worker', however, via the association of violence with both masculinity and working-class groups; among urban male adolescents, violence often plays a pivotal social function, through which you can gain respect for defending yourself (E. Anderson, 1997). Rooney's swearing variously indexed a collection of these social meanings, as evidenced by their presence in the media's

metadiscourse, in both their positive and negative evaluations of his behaviour.

It is for similar reasons that Gwyneth Paltrow's swearing was viewed negatively. While the reaction to Rooney's swearing revolved around identifying which parts of his identity were the cause (e.g., his profession, his social class, his masculinity, etc.), reactions to Paltrow reflect her inability, in the eyes of the media, to index some of the social meanings associated with swearing (Christie, 2013). Just as Campbell-Kibler (2010a) showed that listeners' expectations about a speaker affect their perception of variable (ING), people have expectations about the types of people who typically swear and, based on the media reaction to Paltrow, we can assume that the social characteristics they normally associate her with are in opposition to those in Figure 7.4. Notably, Paltrow was using *cunt*, while Rooney used *fuck*; it may be that, had she used *fuck*, she would have been more successful in having her swearing perceived as more authentic. The two words also differ with respect to valence (Janschewitz, 2008; Warriner et al., 2013); *fuck* is arguably more frequently used in positive evaluations, whereas *cunt*, with some culture-specific exceptions, is typically only used in negative evaluations.

An indexical field is a useful way of visualising the social meanings associated with swearing, as will become clear in the following section. We know that swearing is variably linked to particular stances and interpersonal functions. We also know that swearing is variably associated with particular traits and characteristics. Furthermore, we know that swearing is variably considered typical of particular social types and personae. In addition, we know that these stances and functions, traits and characteristics, and social types and personae, are ideologically linked to one another. While a precise indexical field of 'swearing' as an abstract concept would be unimaginably large, drawing smaller indexical fields to represent the ways in which people have been shown to use and been perceived for using swearwords can help to make these links clearer. This helps when trying to explain why particular types of swearing are evaluated in the way that they are.

7.3 The social meaning in swearing: contextualizing previous findings

The social meanings included in the indexical field in Figure 7.4 are just the tip of the iceberg when it comes to the complex social profile of swearing. In Chapter 2, I reviewed a variety of studies that have attempted to link the use of swearing with particular social functions in production and evaluative dimensions in perception. In production, swearing has been variously linked to indexing *solidarity* (Daly et al., 2004; Baruch & Jenkins, 2007; McLeod, 2011), *humour* (Stapleton, 2010; T. Jay, 2009b; Seizer, 2011), *intimacy* and *trust* (Stapleton, 2003), *honesty* (Feldman et al., 2017; cf. R. E. de Vries et al., 2018),

informality (Mercury, 1995; Bayard & Krishnayya, 2001; Stapleton, 2010), particular forms of *masculinity* (De Klerk, 1997; Eckert & McConnell-Ginet, 2013; Lawson & Milani, 2015), particularly working-class *femininities* (S. E. Hughes, 1992) and *working class* speech more generally (Cheshire, 1982; Romaine, 1999; Stapleton, 2010), as well as displaying usage patterning across macro-social categories including age, gender and social class (McEnery, 2004; McEnery & Xiao, 2004; Love, 2017, 2021)

In perception, swearing has been linked to a range of evaluative dimensions, including lower ratings of overall *impression*, *trustworthiness* and *intelligence* (DeFrank & Kahlbaugh, 2019), *professional capability* (Paradise et al., 1980; Patton et al., 2017) and *effectiveness* (Johnson, 2012) and increased ratings of *believability* (Rassin & Heijden, 2005), *offensiveness* (Beers Fägersten, 2012), *persuasiveness* (Cavazza & Guidetti, 2014; cf. Scherer & Sagarin, 2006) and *intensity* (Patton et al., 2017). Additionally, several studies have shown the relevance of these evaluative dimensions to depend on language-external factors such as the gender of the speaker (Martin, 1997; Howell & Giuliano, 2011; DeFrank & Kahlbaugh, 2019), their ethnicity (Jacobi, 2014), and their social status and the formality of the setting (T. Jay & Janschewitz, 2008; Johnson & Lewis, 2010).

To illustrate the social complexity of swearing and how some of these meanings likely interrelate, I have drawn several further indexical fields based on some of these studies. Rather than attempt to include each and every social meaning that swearing has been shown to index, I will instead focus on smaller subsets of these meanings. In doing so, I hope to illustrate how situated uses of swearwords and the interpersonal stances they index are tied to larger associations with social traits and identities, as well as more concrete social types. This will, in turn, unite the disparate findings on social perceptions of swearing. I will then move on to how my own findings figure into the indexical field of swearing in Section 7.4.

Building on Christie's (2013) example of how Wayne Rooney's swearing was linked to the *industrial workplace*, I have modelled the indexical field of swearing (in English) in the workplace based on the findings of Daly et al. (2004), Baruch and Jenkins (2007) and McLeod (2011) in Figure 7.5 (see Section 2.4.1 of Chapter 2 for more details on these studies). This time, I have used grey font to indicate both stances and the interpersonal functions of swearing (e.g., *whingeing*). In each study, swearing appears to index a similar social type - a factory worker, a long-term worker or a *tradie* - each of which revolves around similar permanent qualities and the use of swearing in particular types of speech acts. People are able to index their in-group membership to one of these social types by aligning themselves particular qualities. If you are able to withstand the use of strong swearwords from others, showing you are *resilient*, and if you align your swearing with a *male* way of speaking by avoiding swearing in front of women (Baruch & Jenkins, 2007), you might be accepted into the in-group. If you use swearing to perform the different

interpersonal functions of whingeing and complaining, you can mark *solidarity* with the in-group (Daly et al., 2004). If you use very strong swearwords like *cunt* in excess for the purpose of *humour*, you can differentiate yourself from the rest of society and align yourself with the *tradies* (McLeod, 2011).

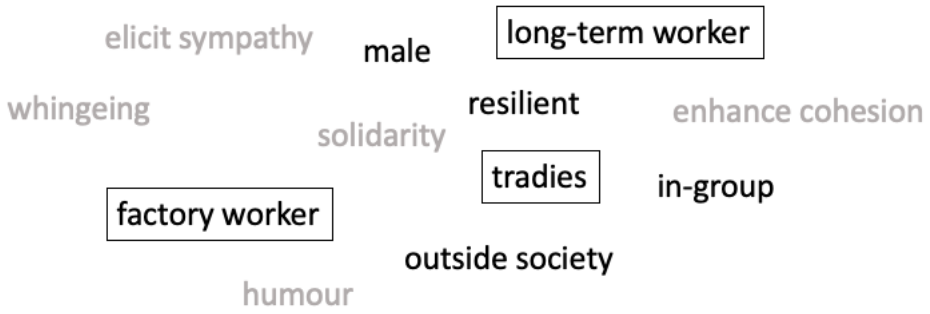

Figure 7.5: Indexical field of 'swearing' in the workplace: Boxes = social types, black = permanent qualities, grey = stances/interpersonal functions. Based on the findings of Daly (2004), Baruch (2007) and McLeod (2011).

The point of combining the findings from these three different studies is to try and build an understanding of how the individual uses of swearwords, each with their own conversational purpose, can accrue meanings that align the users with one social identity and differentiate them from another. The *style* of speaking used in these workplaces, which involves a lot of swearing, reproduces the social structure of its workers. There are those who swear (e.g., the factory workers, the men on the production line, the tradies) and those who don't (e.g., the female workers, upper management, the rest of society). In each case, the stances and interpersonal functions of swearing are linked to these particular social types via an association with more permanent traits and characteristics.

I have modelled another example of an indexical field in Figure 7.6 by synthesizing the findings of Coates (1986), S. E. Hughes (1992), De Klerk (1997), Martin (1997), Stapleton (2003), Howell and Giuliano (2011) and Beers Fägersten (2012) in relation to swearing and gender in English (see Section 2.4.1 of Chapter 2 for more details). Again, only a subset of the possible social meanings is used to illustrate my point. On the one hand, swearing is associated with *masculine* traits (Martin, 1997), including *toughness* and the use of *aggressive* or *forceful* stances (Coates, 1986; De Klerk, 1997). There is also evidence that swearing by women (Beers Fägersten, 2012) and in front of women (Howell & Giuliano, 2011) is not evaluated positively. In the case of Howell and Giuliano (2011), a male coach is seen as more effective when using swearing, provided the team they are coaching is male, aligning swearing with masculinity specifically. Perception work by DeFrank and Kahlbaugh (2019) also suggests that swearing is better received in same sex conversations. These various findings point towards swearing have a socio-indexical

association with masculinity, and potentially specifically masculine social types.

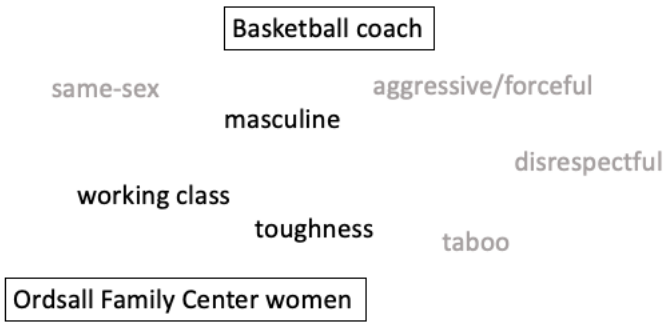

Figure 7.6: Indexical field of 'swearing' and gender: Boxes = social types, black = permanent qualities, grey = stances/interpersonal functions. Based on various findings on swearing and gender.

Some of the same stances and qualities that lead swearing to be associated with men and masculinity also constitute a specifically female social type in the Ordsall Family Centre in Salford in the work of S. E. Hughes (1992), however. The women in the centre are indexing their *toughness* by swearing, building on its association with *working class* speakers to index a locally relevant social type. They are exhibiting a branch of "female speech" that fits their own social network structure. Just as the male-to-male swearing of the basketball coach in Howell and Giuliano (2011) is effective in performing social bonding with the male team, the women of the Ordsall Family Centre move away from the standard norms of other women that would isolate them from the community, preferring to do their social bonding with one another through a locally meaningful swearing practice (S. E. Hughes, 1992). The same sorts of stances and qualities that can lead swearing to be associated with masculinity can, therefore, also come to be associated with more locally relevant feminine social types.

Again, the point of combining these findings into one indexical field is to move beyond thinking about swearing in terms of more basic statements of the form 'X swears a lot' or 'Swearing makes you sound Y', where X is a particular social group or type and Y is a social trait. Even if non-linguists think about swearing in this way - via what Agha (2003) calls *ethnometapragmatic* stereotypes - we as sociolinguists understand that these associations do not happen by accident. Swearwords serve an interpersonal function between individuals. These functions and the use of swearwords in them come to be associated with particular speaker qualities over time. These qualities are then associated with more concrete, socially locatable personae. A single person can activate any number of the ideologically related social meanings associated with swearing. As I have shown in the previous three chapters, and as I will discuss in the following section, the activation of these social meanings in the mind of the perceiver depends not only on the language-

external factors identified by previous researchers (T. Jay & Janschewitz, 2008; Howell & Giuliano, 2011; Jacobi, 2014; Cavazza & Guidetti, 2014), but also on language-internal factors. These factors also need to be accounted for in the indexical field of 'swearing'.

7.4 The social meaning in swearing: contextualizing linguistic variation

The indexical fields for swearing drawn in Section 7.3 rely on the binary distinction between *swearing* and *not swearing*, with the fields depicting the social meanings associated with the former. In studies of sociolinguistic variation, in which each variant has been shown to activate different collections of social meanings, sociolinguists draw their indexical fields by depicting all the possible variants, with social meanings associated with a variant drawn closer to that label. An example of this from Schleef and Flynn's (2015) work on perceptions of variable (ING) in Manchester is included in Figure 7.7. Labels inside dotted lines represent social meanings associated with one particular variant, while more gradient differences are represented outside these lines, with a smaller distance between a social meaning label and a variant indicative of a great socio-indexical association.

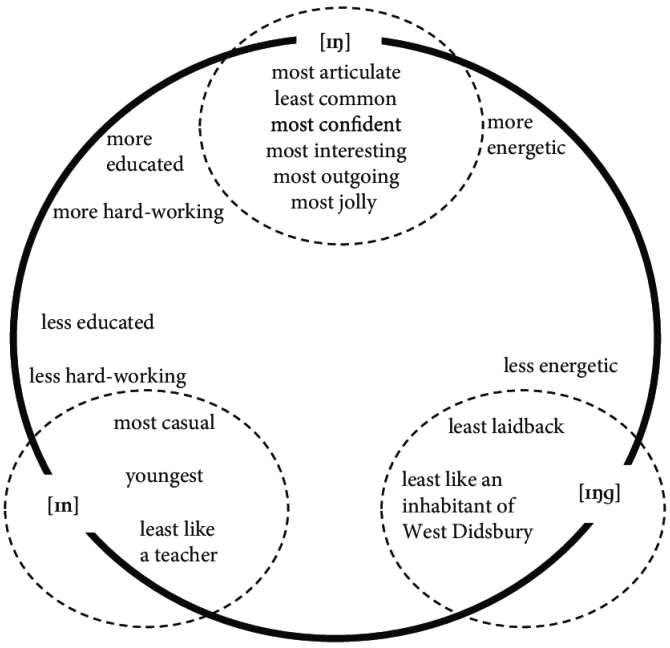

Figure 7.7: Indexical field of variable (ING) in Manchester: Boxes = social types, black = permanent qualities, grey = stances/interpersonal functions. From Schleef and Flynn (2015).

The indexical fields I have drawn for swearing, based on the findings in this thesis, more closely approximate Figure 7.7, as I am comparing different variations of swearing.

They differ in a number of ways, however. My indexical fields only ever include two variants, compared to the three included above. Mine do not have dotted lines as all of the social meanings concern binary comparisons. Mine also feature fewer social meanings, as I tested fewer evaluative dimensions than e.g., Schleef and Flynn (2015). In addition to the social meanings identified as being strongly associated with one form or construction over another, I have also included some of the social meanings and stances that are associated with swearing more broadly, placing these in the centre of the field in proximity to a label denoting the swearword(s) in question. The closer the labels are to a particular variant, the stronger the association.

I will start with the indexical field for variation in infixation (see Chapter 5), followed by variation in animacy (see 6), before finally coming on to variation in (ING) 4; I am ending with the indexical field for variable (ING) because this is somewhat more complicated.

7.4.1 The indexical field of swearing infixation

Figure 7.8 provides a representation of an indexical field for swearing infixation. The two variants - infixed and non-infixed constructions - are represented in small circles at either end of the field, with a box representing FUCKING in the middle. While other swearwords and swearing euphemisms can be used in infixed constructions (e.g., *fan-bloody-tastic*, *un-frigging-believable*), only *fucking* was tested in my experiments. The indexical field in Figure 7.8 is therefore better thought of as an indexical field for *fucking infixation* specifically.

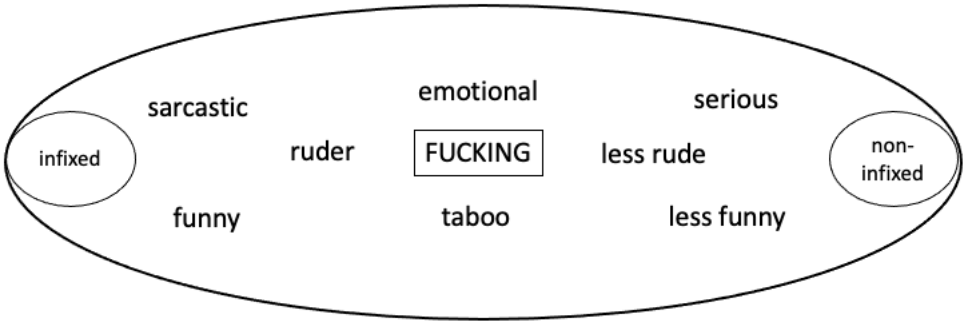

Figure 7.8: Indexical field of *fucking infixation*.

Very close to the representation of FUCKING are the meanings *emotional* and *taboo*, both of which are commonly understood criteria of swearing more generally (Ljung, 2010). The labels *ruder* and *less rude* are reasonably close to FUCKING for similar reasons; rudeness is a core meaning attached to swearing, via its breaking of societal norm.

The role of infixation, compared to a lack thereof, appears to be an enhancement of this meaning. That is, it would be fair to say that rudeness is a consistently relevant evaluative dimension in swearing, and the addition of expressive morphology, such as infixation, pushes this dimension into further prominence.

The labels *sarcastic* and *funny* are further away from the core label of FUCKING . These evaluative dimensions came up in a pre-experiment task that included both infixed and non-infixed constructions. While humour can undoubtedly be indexed by swearing more generally (Stapleton, 2010), this is something that seems much more context-dependent. My results also suggest that the degree of humour associated with swearing varies across different constructions, with infixed swearing perceived as funnier than non-infixed swearing.

I would argue that the same is true of the label *sarcastic*. Swearing and sarcasm share some properties. Both convey something not-at-issue about the speaker's perspective. Both are sometimes grouped together under the label *verbal aggression* (see e.g., Myers & Knox, 1999). I would argue that sarcasm isn't a meaning that is central to swearing, however, but rather, is an evaluative dimension made relevant in particularly marked occurrences of swearing (see Section 4.6 for more discussion on the markedness of different constructions).

There are numerous possible linguistic cues to sarcasm (Attardo, Eisterhold, Hay, & Poggi, 2003), particularly phonological cues, including exaggerated intonational patterns, singsong melody, falsetto, exaggerated stress, monotonous intonation, and separation by long pauses between the words (Haiman, 1998). There are also some examples of morphemes that function as markers of sarcasm, including the use of *quote*, as in (2). In the case of swearing, my results would suggest that infixation is also a marker of sarcasm, with non-infixed constructions judged as more sincere by comparison.

(2) Your quote-unquote principles are nothing but snobbism.

(Haiman, 1998, p. 47)

Importantly, as with any indexical field, not all of the social meanings indexed by swearing infixation will be activated by any one usage for all that hear it. The meanings in an indexical field are *potential* meanings (Eckert, 2008), not meanings that are guaranteed to emerge every time. In Experiment II of the chapter on infixation (Chapter 5), the results of principal component analysis suggested that, while perceptions of rudeness and sarcasm co-varied, ratings for *funny* and *rude* were inversely correlated. Although it is possible to be both funny and rude simultaneously - for example through mocking or crude humour - this doesn't appear to be the case for swearing infixation. For at least some people, in some sentences, infixed swearing is funny. For others, it is rude.

7.4.2 The indexical field of animate and inanimate swearing modification

Figure 7.9 provides representations of indexical fields for animate and inanimate swearing modification. As a reminder, the examples in (3) depict animate (3-a) and inanimate (3-b) swearing modification, respectively. Separate fields for each swearword are presented. This is because they differed significantly in how they influenced participant responses on Likert scales, both independently and in their interaction with the animacy condition.

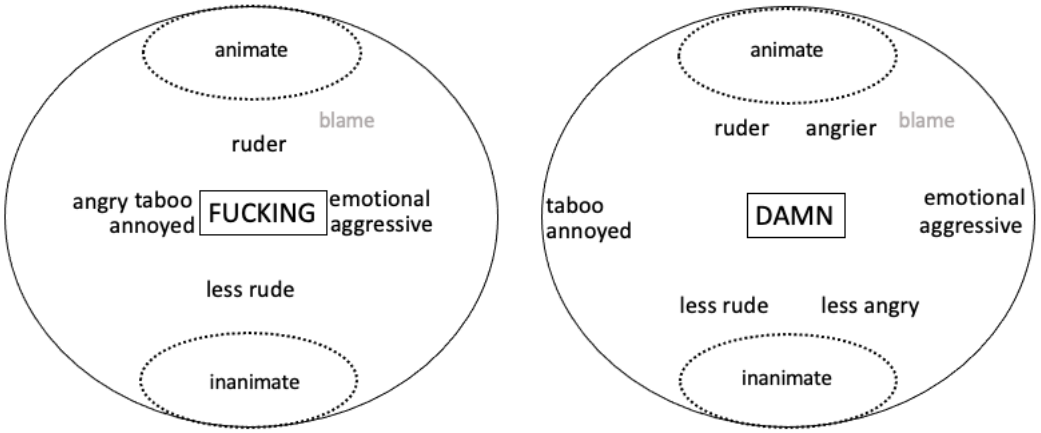

Figure 7.9: Indexical fields of animate and inanimate swearing modification for *fucking* and *damn*.

(3) a. The damn kids stained the carpet
b. The damn paint stained the carpet

The two indexical fields differ in a number of ways to reflect the different potentials that animate and inanimate constructions have for indexing this set of social meanings. Firstly, the labels *taboo*, *emotional*, *annoyed* and *aggressive* are closer to FUCKING than they are to DAMN. Previous work has demonstrated that *damn* is considered less offensive (Beers Fägersten, 2012) and has a lower average value of arousal (Janschewitz, 2008) than *fucking*. Furthermore, in Experiment II of Chapter 6, *fucking* was shown to make 'speakers' appear significantly more aggressive and annoyed, independent of the type of swearing construction they used.

Secondly, in the indexical field for *damn*, the labels *ruder* and *angrier* are closer to the *animate* label, while in the indexical field for *fucking*, only the label *ruder* is closer to the *animate* label. This again reflects the results of the experiment, in which ratings on the scale of *angry* for *damn* were modulated by the animacy condition, but not for *fucking*.

In both indexical fields, in grey font, I have included the label *blame*. Like the indexical fields above, I have used grey font here to represent an interpersonal function. Blame

is what I believe led to higher ratings of rudeness for animate constructions than inanimate constructions, for both *fucking* and *damn*. As Frazier et al. (2015) suggest in their study, and as I explicate in greater detail in Sections 6.3.1 and 6.5.3 of Chapter 6, when faced with a swearword that precedes an animate noun, people are more likely to infer that the speaker is holding that entity responsible for the unwanted state of affairs than if the swearword were to precede an inanimate noun. Blaming someone using a taboo word was perceived to be ruder and, in the case of *damn*, angrier than using a taboo word to express heightened emotion at a general state of affairs. This is because the entity blamed in the animate case is capable of feeling offended by the swearword, while the entity in the inanimate case is not. The interpersonal function or stance of blaming a person (or animal) is what leads to the indexing of the social meanings *rude* and *angry*.

As mentioned previously, the social meanings contained in Figure 7.9 are really socio-pragmatic meanings, rather than social meanings tied to particular social types or personae; that is, they relate to social meanings that are relevant to that specific conversation, rather than necessarily relating to traits that endure once that conversation has ended. One could imagine, however, that variation in animacy, and the socio-pragmatic meanings that this variation can index, could contribute to indexing a social type. It may be that particular types of people are more likely to swear in a way that suggests blame to a person than to express a general heightened emotion. While this is highly speculative and I have no particular social group(s) in mind, previous work in sociolinguistics has suggested that structural variation can be used in different ways by different groups of speakers (e.g., Snell, 2010).

In their production study on tag questions (e.g., (4), with the tag questions in **bold**), Moore and Podesva (2009) track different uses of the construction in conversations between female students at an English high school. Dividing their speakers into four communities of practice (Eckert & McConnell-Ginet, 1992), Moore and Podesva (2009) identify the different stances that members of these groups take using tag questions. For example, the *Populars* group more frequently used tag questions to talk about girls in other social groups (4-a), compared to the other three social groups, with the authors suggesting that the *Populars* use tag questions to construct themselves as *cool* and *distinct*. This contrasts with the use of tag questions by the *Townies* community of practice. The *Townies*' use of tag questions focused more on their own group dynamics, with tag questions specifically used to "encourage involvement in the telling of the narrative" (Moore & Podesva, 2009, p. 471), as in (4-b). The authors suggest that this is done to conduce a shared viewpoint on the group's rebellious activities.

(4) a. they've never got along really, **have they?**

b. [He was playing] all these songs over the thing, **weren[ø] ['e?**...

As well as the overall frequency of tag questions differing between the groups, therefore, the internal composition of the tag questions used by the different groups also differed, serving different conversational purposes for each (Moore & Podesva, 2009). One could imagine a situation in which this could also be the case for the animate and inanimate swearing constructions discussed in Chapter 6 of this thesis. For people wishing to use their swearwords to express heightened emotion towards specific people, in order to effectively convey anger and to be more conversationally transgressive, using a swearword before an animate noun is the best option. One might imagine that the sports coach in Howell and Giuliano (2011) might do this when motivating his team to beat another. Similarly, the workplace whinging discussed in Daly et al. (2004) might involve blaming someone, or it might involve whinging about a general state of affairs, with one of these more typical of a certain set of people in the factory. This is again highly speculative, however.

7.4.3 The indexical field of swearing and variable (ING)

Figure 7.10 provides a representation of the indexical fields for alveolar [ɪn], swearing and velar [ɪŋ]. For this indexical field, I have used the label SWEARING to represent the collection of 9 swearwords used in Experiment I of Chapter 4, which includes the subset of 3 swearwords used in Experiment II of the same chapter. I have chosen to use multiple, intersecting indexical fields to illustrate the findings of this chapter. Each variant of variable (ING) (excluding velar nasal plus), and the concept of swearing, are represented by their own fields. The indexical fields for the two (ING) variants do not intersect, as their use is mutually exclusive. The field for SWEARING is intersected by both (ING) fields to represent the two ways of pronouncing a swearword ending with '*-ing*' (e.g., *fucking* or *fuckin*). Those sections of the (ING) fields that do not intersect with the SWEARING field represent occurrences of the respective variant on a neutral word.

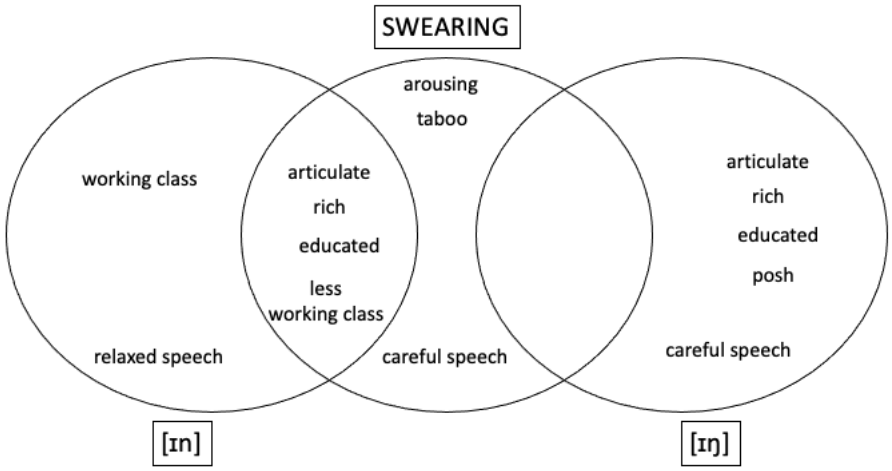

Figure 7.10: The intersecting indexical fields of alveolar [ɪn], swearing and velar [ɪŋ]

The [ɪŋ] field contains the meanings *articulate*, *rich*, *educated* and *posh*, reflecting the fact that, overall in the matched-guise experiment, the guises that were saturated with the velar variant received higher ratings on scales depicting these meanings. This is also what was found by Schleef et al. (2017). The same is true for the presence of the label *working class* in the [ɪn] field.

To reflect the fact that participants tended towards treating alveolar [ɪn] tokens in swearwords as if they were velar [ɪŋ] tokens - at least on these four scales - the labels *articulate*, *rich*, *educated* and *less working class* are included in the intersection between the [ɪn] field and the SWEARING field.

The reader will recall that I offered two possible explanations for the effect observed in Chapter 4. I have attempted to account for each of these explanations in the indexical field of swearing; as the meanings in an indexical field only represent *potential* meanings (Eckert, 2008), I can do this without committing to either explanation. The cognitive explanation involved the increased attentional resources required to process a swearword due to its increased levels of arousal and tabooness (MacKay et al., 2004; Eilola & Havelka, 2011). If the swearwords did take up more attentional resources, participants may have been distracted from the variable (ING) tokens and, as a result, assumed that the speaker had used the most common variant more generally, namely the velar [ɪŋ] variant (see also Vaughn & Kendall, 2018). To reflect this, I have included the labels *arousing* and *taboo* in the field of SWEARING.

The social explanation for the results observed in Chapter 4 concerned the tendency for swearwords to be used to create verbal emphasis (Stapleton, 2010) and, therefore, be more frequent in *careful speech*, compared to *relaxed speech*. The alveolar [ɪn] variant is more frequently used in casual speech, while the velar [ɪŋ] variant is more frequent in careful speech (Labov, 1972). As a result, the association of swearing with careful speech

may have led listeners to expect the velar variant, regardless of what they actually heard. To reflect this, the label *careful speech* has been included in both the [ɪŋ] and SWEARING fields, and *relaxed speech* has been included in the [ɪn] field.

This is obviously contrary to how I originally expected the indexical field(s) of swearing and variable (ING) to look. Originally it was expected that, because swearing and alveolar [ɪn] share similar socio-indexical associations (see Section 4.4.1.1.2 of Chapter 4), the combined presence of the two (e.g., in *fuckin*) would lead to a unique combination of emergent social meanings. I have illustrated this in Figure 7.11, which represents the type of indexical field I had predicted before conducting the experiments. Both swearing and alveolar [ɪn] are broadly associated with working class speech (S. E. Hughes, 1992; Stapleton, 2010; Schleef et al., 2017), informality and casualness (Stapleton, 2010; Schleef et al., 2017) and particular speaker qualities such as lower intelligence (DeFrank & Kahlbaugh, 2019; Schleef et al., 2017), hence the inclusion of these labels in the indexical fields of each in Figure 7.11, and the inclusion of the labels *very informal*, *very casual*, *very working class* ad *very unintelligent* in the intersection of the two fields.

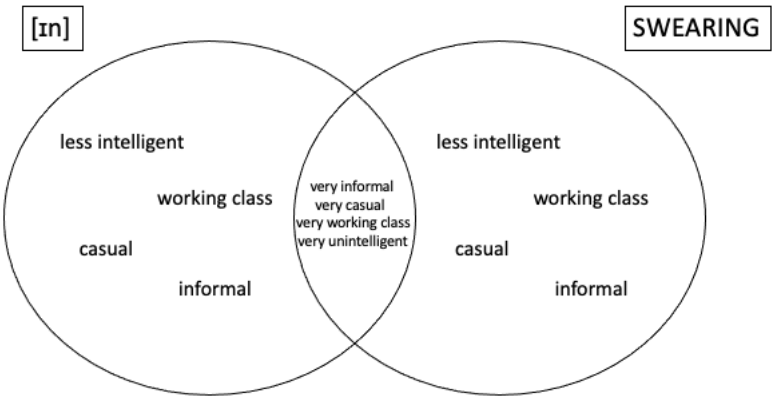

Figure 7.11: The *predicted* indexical fields of alveolar [ɪn] and swearing

7.5 Consequences for swearing research

To re-orient this chapter towards the bigger questions about swearing that this thesis aims to address, I am going to quote the final paragraph of Section 1.1, Chapter 1:

> *In sum, there is a large amount of linguistic variation in swearing. This variation occurs at multiple levels of structure, including phonetics, morpho-phonology and syntax and semantics. Given that perceptions of swearing are shown to vary as a function of language-external factors, including speaker- and context-specific factors, we might also expect variation as a function of language-internal factors.* ***Are all pronunciations of a swearword perceived the same? All word formations?***

> ***All sentence structures?*** *If the social meanings attributed to swearing depend on who and in what context the swearing occurs, it would be logical to expect those meanings to also depend on how they occur.*

The main questions are included in bold and relate to Chapters 4, 5 and 6. In answer to the first question - whether all pronunciations of a swearword are perceived the same - the findings in Chapter 4 suggest that, at least in relation to variable (ING), this might sometimes be the case. That is, the results of the experiments suggest that, when a swearword ending in '*-ing*' is ambiguous between velar and alveolar pronunciations, listeners will tend to hear it as a velar pronunciation. In the context of full sentences, even if the swearword is pronounced with an unambiguous alveolar [ɪn] ending, listeners may still tend towards hearing it as a velar [ɪŋ] ending, especially if other (ING) words in a person's speech have the velar ending. This is just a tendency, however; people still frequently perceive e.g., *fucking* and *fuckin*, *shitting* and *shittin*, *dicking* and *dickin* etc., as distinct from one another.

In answer to the second question - whether all word formations of a swearword are perceived the same - my findings in Chapter 5 suggest that, at least in relation to infixation, they are not. That is, the results of my experiments suggest that e.g., *fan-fucking-tastic* and *fucking fantastic*, activate different social meanings; this is further modulated by the well-formedness of the construction.

In answer to the third question - whether all sentence structures containing a swearword are perceived the same - my findings in Chapter 6 suggest that, at least in relation to the property of animacy and not the property of gradability, they are not. That is, the same sentence activates different social meanings dependent on whether it contains a determiner phrase like e.g., *the fucking kids* or *the fucking paint*.

These findings have significant consequences for swearing research. As I discussed at length in Section 2.4 of Chapter 2, swearing has been linked to a variety of social meanings. In Section 2.4.2 of Chapter 2 specifically, I reviewed numerous experimental perception studies that have suggested that swearing leads to speakers being socially evaluated differently than if they had not sworn. These various studies relied on a binary classification of swearing in their experiments, with the only two options being *swearing* or *not swearing*. In these studies, swearing in English was linked to perceptions of increased *offensiveness* (Beers Fägersten, 2012), decreased *intelligence*, *trustworthiness* and *likeability* (DeFrank & Kahlbaugh, 2019), increased *persuasiveness* (Scherer & Sagarin, 2006), increased *intensity* (Patton et al., 2017) and decreased *capableness* and *effectiveness* (Howell & Giuliano, 2011). In Dutch, swearing was linked to an increase in perceived *credibility* in the courtroom and in Italian, swearing was linked to a decrease in *persuasiveness* for a male politician (Cavazza & Guidetti, 2014).

My own findings call into question the generalizability of these and other similar

findings. It isn't really possible to say that swearing leads to a person being perceived as more or less anything when perceptions of the same swearword can depend on how that swearword is produced. Swearing will be perceived differently if it is infixed than if it is not infixed. It will also be perceived differently if it precedes either an animate or inanimate object. Finally, even though I have concluded that it is sometimes possible that the realisation of variable (ING) won't affect how a swearword is perceived, the presence of a swearword can affect how other socially meaningful forms are perceived. Perceptions of the rest of a person's speech style may therefore be affected by the presence of a swearword, meaning that other sociolinguistic variants in a person's speech may not affect how they are perceived in the same way.

Each finding has more specific consequences for swearing research too. My finding for swearing and the sociophonetic variable suggests that swearwords are not just socially salient to the point that they affect how a person is socially evaluated, but rather, they are so salient in a person's speech that they can influence the social perception of other forms around them. This might suggest that, if a person is swearing, the presence of other socially meaningful variation in their speech becomes less relevant.

My finding for infixation suggests that people have a perception of which swearing constructions are more or less genuine expressions of heightened emotion. If a person uses a less frequent swearing expression, especially one which is more structurally marked such as swearing infixation, they may be perceived as swearing not for the purpose of expressing emotion, but for expressing sarcasm. Depending on the hearer, this might be perceived to be funny or it might be perceived to be rude.

My finding for animacy suggests that people modulate their reactions to swearing depending on who they think the target of the swearing is. People reason over the possible intended meanings of a swearword and the other possible structures a speaker could have used to convey those meanings with minimal ambiguity. Having reached a conclusion about what the most likely target of the swearword is, people react accordingly. Much like the slur *favelado* (Beaton & Washington, 2015), the indexical field of swearing is not limited to social meanings and stances that pertain to the speaker, but also to the entities to which or whom they make reference.

This is all to say that myriad other factors are likely to be at play when a person is confronted with swearing, other than just the fact that a person chose to swear. Insofar as the principal aim of this thesis was to move swearing research forward and highlight as-yet unexplored factors that could affect how swearing is perceived, my research and results contribute significantly to this aim.

7.6 Limitations

One obvious limitation of this thesis is its concentration on English. I will discuss this in more detail in the following section on future directions for this research, so for now I will simply say that all of my findings are restricted to English. Furthermore, and as I will again touch on in the next section, my findings are limited to the four variables on which I have chosen to focus. Rather than providing definitive evidence that phonetic, morphological and semantic/pragmatic variation in swearing definitely affects social evaluation of a speaker, I have instead provided a set of examples that suggest that this might be the case, with the added hope that this might open the door to further use of this approach to the study of swearing.

There are also limitations with respect to the conclusions I have been able to make about the specific sociolinguistic variables that I have tested. In Chapter 2, in which I comprehensively reviewed the previous literature on swearing, I highlighted the lack of work in swearing research on language-internal factors; this was the gap I aimed to fill with this thesis. Part of the motivation for doing this was that previous literature had already established that language-external factors affect how swearing is perceived (Martin, 1997; T. Jay & Janschewitz, 2008; Howell & Giuliano, 2011; Jacobi, 2014; Cavazza & Guidetti, 2014), so it stood to reason that language-internal factors may also be playing a role.

The effects of language-internal factors on perceptions of swearing do not occur independent of language-external factors, however. That is, whatever the effect might be predicted to be of using a particular pronunciation, word formation or sentence structure over another, this effect could still be modulated by the identity of the speaker and/or the setting in which they are in. As much as I have shown that linguistic variation can affect how a person is socially evaluated, the effects will not be the same in every context.

As an example, while swearing infixation was shown in Chapter 5 to make a person seem more sarcastic, funnier and/or ruder, the emergence of these social meanings is very likely to depend on context. Perception experiments like matched-guise tasks are something of a black box. On the one hand, they allow experimenters to control for other potentially interfering factors, meaning that any changes to how participants evaluate a speaker can be exclusively explained by the relevant experimental manipulations. That is, in Experiment II of Chapter 5, changes in participants' responses on the Likert scales could be attributed to the presence or absence of infixation.

On the other hand, however, such tasks involve participants creating abstract representations of speakers and situations in their minds and drawing an inference based on that abstract representation. In my task, they likely thought of the most likely speaker that would utter a sentence containing an infixed swearing construction and responded to the

Likert scales based on how they would socially evaluate that person. They did the same for sentences containing non-infixed constructions and, using this data, I was able to draw some conclusions about how infixed and non-infixed constructions differed with respect to social meaning in the abstract. This is accounted for by using indexical fields to model the social meanings of swearing infixation, as each and every social meaning included in an indexical field is a potential social meaning, not a something that is fixed.

What my experiment did not test is whether people have different reactions to swearing infixation depending on the person doing it and the situation in which they do it. T. Jay and Janschewitz (2008) showed that swearing is perceived as less appropriate in certain situations. Would swearing infixation activate the same social meanings on a building site as it would in the classroom? What about when spoken by a man vs spoken by a woman? While swearing infixation might be considered, on the aggregate, ruder than a non-infixed construction, this effect could be modulated by numerous language-external factors.

Further research combining language-external and language-internal factors in the social perception of swearing is therefore required in order to gain a fuller picture. Until that happens, the results in this thesis are only indicative of how different swearing constructions is socially evaluated in the abstract.

7.7 Future directions

There are at least two directions in which this research should be taken to further improve our understanding of swearing and the role that internal linguistic variation plays in how it is perceived. The first is to try to extend my findings to other linguistic variables in English. My conclusions about phonetic, morphological and semantic/pragmatic variation are limited to variable (ING), infixation and animacy and gradability. But there are numerous other forms of linguistic variation from these domains that may also affect social evaluation of the speaker. I will discuss one example from each here.

With respect to phonetics, we might expect variation in vowel production to affect how the use of a swearword is perceived. We know, based on work by Gold and McIntyre (2016), that the duration of the /ʌ/ vowel in *fuck* varies depending on whether the person is expressing surprise or insulting someone. If people are aware of the relationship between vowel duration and pragmatic meaning, either implicitly or explicitly, then this may affect the social meanings that they associate with the use of *fuck*. The same could be true of the /ɪ/ vowel in *shit*, the /ʌ/ vowel in *cunt* etc. Conversely, it may be that, as may have been the case with variable (ING), such variation wouldn't affect social perception because the semantic meaning of the swearword is so cognitively salient that its pronunciation receives less attention compared to the pronunciation of neutral words.

With respect to morphology, I previously discussed several other examples of expres-

sive morphology that could also be sources of social meaning in the same way that swearing infixation is. One of these examples is compounding. In English, swearing compound constructions such as *a fuckton* and *a fuckload* are well-attested; both have entries in the Urban Dictionary, for example. These constructions contrast with their non-compounded counterparts *a fucking ton* and *a fucking load*. The combining of a swearword and a measure term in this way is an example of expressive morphology and, much like infixation, could be perceived to be more playful due to being more structurally marked.

Finally, with respect to semantics and pragmatics, as I discussed in Section 2.3.3.2 of Chapter 2, while swearwords, as expressives, are in some ways semantically homogenous, they vary with respect to the meaning contributions they make to a sentence. Two further examples from English that might be sources of different social meanings are the literal/metaphorical distinction and expressive modifiers. Concerning the former, Beers Fägersten (2012) has already shown that literal uses of a swearword (e.g., *she fucked this guy*) are considered more *offensive* than non-literal uses of the same word (e.g., *she fucked up the exam*); further work could establish whether a wider range of social meaning differences emerge as a function of this distinction. Regarding expressive modifiers, further work could establish whether people perceive swearing differently if its sole function is to add further heightened emotion to an already emotionally laden sentence (e.g., *you're a fucking bastard*).

The second direction in which this research could be taken is to extend my findings to other languages. The majority of the sociolinguistic studies I reviewed in Chapter 2 pertained to English, with a few exceptions. All of my own findings also pertain exclusively to English. This is obviously a problem, because if we want to be able to make claims about swearing, as a cross-linguistic and cross-cultural phenomenon, we must verify whether these kinds of results can also be found for languages that differ significantly from English. I will briefly highlight one case in which one could expect different effects to those that I found.

In Chapter 6, I tested the difference between swearing modification of animate and inanimate nouns on sentence processing (Experiment I) and social evaluation of the speaker (Experiment II); this had also been done using a judgment task by Frazier et al. (2015). I did this using English sentences such as those in (5). Importantly, English is a language without grammatical gender.

(5) a. The damn bear broke the fence clean off
    b. The damn wind broke the fence clean off

As I highlighted in Chapter 1 (see Section 1.1), however, in languages like German, swearwords are marked for grammatical gender. The German translation of the inanimate-subject sentence in (5-b) would therefore be as in (6), with the word *verdammt* marked to

reflect the masculine gender of *Wind*.

(6) Der verdammt**e** Wind hat den Zaun komplett abgebrochen

I suggested in Chapter 6 that, when reading a sentence in which a swearing modifier is prejacent to an inanimate noun, people will be more likely to think that the modifier applies non-locally, signalling the speaker's heightened emotion towards the state of affairs, rather than signalling their heightened emotion toward that specific individual. This might not be the case in (6), however, because of the gender marking. As the marking depends on the gender of that specific noun, in this case masculine, people may reason that the speaker chose to do this because they wanted to convey the heightened emotion towards the wind specifically as the cause of the state of affairs. As the adjective must agree with the gender of the noun, people may be more likely to think that the expressive meaning also applies to the noun, rather than to the state of affairs described by the sentence. This, in turn, might mean that the effect I observed for animacy on social evaluations of swearing would not replicate for languages with grammatical gender. This is admittedly speculative, but it offers food for thought.

As well as conducting further research on other variables and in other languages, there is further analysis that could be done with that data that I have already collected. Across all of my experiments, I collected demographic information about my participants, including their gender, age, household income, level of education and self-identified social class. I took certain decisions about how to include these in statistical models; in general, these variables didn't appear to significantly affect my results in any particular way. In some cases, however, such as for household income and level of education, there may have been issues with collinearity between these fixed effects. Furthermore, taking away the threshold for statistical significance for a moment, there may be trends in the data that could lead to further research. It may be that particular portions of the participant samples are leading in the observed effects in interesting ways. An area for further research is therefore to conduct further analysis of inter-participant variability. Similarly, particularly for Experiment I of Chapter 4, further analysis of inter-item variability would be a worthwhile pursuit, as it is highly plausible that certain items may be leading the effect here.

7.8 The final word

In the introduction to Chapter 2 on *Studying Swearing*, I pitched swearing as an interesting phenomenon for academic study due to what it can teach us about human behaviour. I did so on the basis of three main contributions it can make to understanding human behaviour.

I argued that studying swearing can improve our understanding of language, of language processing, and of the relationship between language and the people that use it. Previous studies have typically focussed on just one of these areas. Some studies have focussed solely on the linguistic characteristics of swearwords that differentiate them from other words. Some studies have focussed solely on the way in which swearwords are processed by listeners or readers. Finally, some studies have focussed solely on the social function(s) of swearing. A minority of studies have looked, in-depth, at more than one of these areas (e.g., Frazier et al., 2015).

I have made contributions in each of these areas throughout my thesis. In each chapter, I analysed a linguistic characteristic of swearwords that differentiates them from other words in English. I did this using evidence in the literature, as well as my own original observations. Using this analysis, I made principled predictions about how different types of swearing would be *perceived* by listeners or readers. I successfully tested these predictions, showing that linguistic variation in swearing can affect language processing. Using these results, I made further predictions about how listeners or readers would *notice* the different types of swearing when socially evaluating a speaker. I successfully tested these predictions, showing that linguistic variation in swearing can also affect how people react to the person swearing.

As well as making the contributions to swearing research discussed in Section 7.5 of this chapter, this thesis also makes a significant methodological contribution. The three-step process I have used throughout this thesis - starting with linguistic analysis, testing processing then testing social evaluation - can serve as a blueprint, not just for future research on swearing, but for any research on sociolinguistic perception. Such an approach has already been advocated for by Squires (2016) (see Chapter 3 for more detail). With this thesis, I hope to have provided an example of how we can integrate our understandings of language, language processing and sociolinguistic perception to build holistic accounts of how linguistic forms come to affect the way people are perceived.

Appendix A

A.1 Infixing Experiment I: Trisyllabic adjectives

Word	Stress	Word	Stress	Word	Stress	Word	Stress	Word	Stress	Word	Stress
Feminine	initial	Fortunate	initial	Brilliant	initial	Generic	medial	Seductive	medial	Repellent	medial
Masculine	initial	Merciless	initial	Talented	initial	Destructive	medial	Progressive	medial	Repugnant	medial
Dangerous	initial	Heavenly	initial	Masterly	initial	Forgiving	medial	Untactful	medial	Obnoxious	medial
Natural	initial	Cowardly	initial	Excellent	initial	Protective	medial	Unpleasant	medial	Accomplished	medial
Innocent	initial	Trustworthy	initial	Erudite	initial	Unhappy	medial	Unplayful	medial	Outstanding	medial
Different	initial	Negligent	initial	Wonderful	initial	Dishonest	medial	Uncheerful	medial	Tremendous	medial
Positive	initial	Amoral	initial	Fabulous	initial	Elated	medial	Softhearted	medial		
Possible	initial	Courteous	initial	Important	medial	Precocious	medial	Clearheaded	medial		
Accurate	initial	Rational	initial	Imperfect	medial	Discerning	medial	Uneager	medial		
Serious	initial	Virtuous	initial	Exciting	medial	Unable	medial	Unthankful	medial		
Difficult	initial	Confident	initial	Iconic	medial	Compliant	medial	Regretful	medial		
Radical	initial	Relevant	initial	Attractive	medial	Assertive	medial	Deflated	medial		
Delicate	initial	Spirited	initial	Committed	medial	Indecent	medial	Effective	medial		
Generous	initial	Credulous	initial	Aggressive	medial	Uncertain	medial	Unfocused	medial		
Capable	initial	Sensitive	initial	Convincing	medial	Supportive	medial	Uncivil	medial		
Powerful	initial	Diligent	initial	Respectful	medial	Persistent	medial	Contented	medial		
Limited	initial	Genial	initial	Determined	medial	Unlucky	medial	Unshapely	medial		
Colorful	initial	Mystical	initial	Unstable	medial	Unhealthy	medial	Distrustful	medial		
Elegant	initial	Ignorant	initial	Consistent	medial	Malicious	medial	Unsocial	medial		
Powerless	initial	Moneyless	initial	Appealing	medial	Immoral	medial	Unhelpful	medial		

Violent	initial	Merciful	initial	Defiant	medial	Unwitting	medial	Unworldly	medial		
Glorious	initial	Gluttonous	initial	Sarcastic	medial	Displeasing	medial	Receptive	medial		
Negative	initial	Humorless	initial	Impressive	medial	Impartial	medial	Tyrannic	medial		
Juvenile	initial	Sensible	initial	Courageous	medial	Perceptive	medial	Unmanly	medial		
Adequate	initial	Infantile	initial	Neurotic	medial	Constructive	medial	Regressive	medial		
Dateable	initial	Sorrowful	initial	Possessive	medial	Hardhearted	medial	Shambolic	medial		
Logical	initial	Principled	initial	Abnormal	medial	Informal	medial	Fantastic	medial		
Typical	initial	Cretinous	initial	Inspiring	medial	Unselfish	medial	Appaling	medial		
Credible	initial	Reverent	initial	Offensive	medial	Forgetful	medial	Horrific	medial		
Competent	initial	Humorous	initial	Outgoing	medial	Pragmatic	medial	Horrendous	medial		
Tolerant	initial	Prejudiced	initial	Submissive	medial	Repentant	medial	Atrocious	medial		
Dominant	initial	Civilised	initial	Dramatic	medial	Unloving	medial	Abhorrent	medial		
Likeable	initial	Laughable	initial	Dynamic	medial	Resentful	medial	Distressing	medial		
Practical	initial	Terrible	initial	Alarming	medial	Ungracious	medial	Substandard	medial		
Fallible	initial	Horrible	initial	Impolite	initial	Unworthy	medial	Pathetic	medial		
Critical	initial	Hideous	initial	Eccentric	medial	Untruthful	medial	Abysmal	medial		
Liberal	initial	Sickening	initial	Surprising	medial	Judgmental	medial	Disgusting	medial		
Scrupulous	initial	Pitiful	initial	Magnetic	medial	Coldhearted	medial	Revolting	medial		
Loveable	initial	Odious	initial	Decisive	medial	Warmhearted	medial	Repulsive	medial		

A.2 Infixing Experiment II - Sentences and adjectives

Context	Adjective
I've just been promoted to become manager	Fantastic
My brother completely ignored me at Christmas	Coldhearted
I've been sat watching TV all week	Unhealthy
My dad has been hit by lightening twice in his life	Unlucky
My housemate has left a huge mess in our house	Unpleasant
John is refusing to come to the party on Friday	Unsocial
Our neighbour has been screaming and shouting all night long	Unstable
The new table I bought broke within two days	Substandard
My neighbours have been playing loud music all night	Uncivil
My brother always walks around on all fours	Abnormal
My dad has been smashing up the patio	Destructive
Maria pretended to have a dentist appointment to leave work	Dishonest
For the party I'm wearing a dress covered in sequins	Fabulous
The man exposed himself to everyone in the park	Indecent
My friends stopped talking to me since I came out	Judgmental
This guy keeps coming back to the shop to ask me out	Persistent

A.3 Infixing Experiment II - Example stimuli list

Person A	Person B	Infixed	Formedness	Type
I've just been promoted to become manager	That's fan-fucking-tastic	yes	well-formed	test
My brother completely ignored me at Christmas	That's coldhear-fucking-ted	yes	ill-formed	test
I've been sat watching TV all week	That's fucking unhealthy	no	well-formed	test
My dad has been hit by lightening twice in his life	That's fucked unlucky	no	ill-formed	test
My housemate has left a huge mess in our house	That's un-fucking-pleasant	yes	well-formed	test
John is refusing to come to the party on Friday	That's unso-fucking-cial	yes	ill-formed	test
Our neighbour has been screaming and shouting all night long	That's fucking unstable	no	well-formed	test
The new table I bought broke within two days	That's fucked substandard	no	ill-formed	test
My neighbours have been playing loud music all night	That's un-fucking-civil	yes	well-formed	test
My brother always walks around on all fours	That's abnor-fucking-mal	yes	ill-formed	test
My dad has been smashing up the patio	That's fucking destructive	no	well-formed	test
Maria pretended to have a dentist appointment to leave work	That's fucked dishonest	no	ill-formed	test
For the party I'm wearing a dress covered in sequins	That's fabu-fucking-lous	yes	well-formed	test
The man exposed himself to everyone in the park	That's inde-fucking-cent	yes	ill-formed	test
My friends stopped talking to me since I came out	That's fucking judgmental	no	well-formed	test
This guy keeps coming back to the shop to ask me out	That's fucked persistent	no	ill-formed	test
My son was just in a car crash on the motorway	That's really shit	N/A	N/A	dis
My cat died over the weekend from cancer	That's shit	N/A	N/A	dis
I just won a million pounds on the lottery	That's bloody great	N/A	N/A	dis
On my way home I saw my neighbour fall into a hedge	That's damn funny	N/A	N/A	dis

The new car I just bought has already broken down	That's really crap	N/A	N/A	dis
My boyfriend forgot to get me a present for my birthday	That's crap	N/A	N/A	dis
Peter lied about having to go to the dentist	That's bollocks	N/A	N/A	dis
My teacher is making me stay after class to clean the whiteboard	That's bullshit	N/A	N/A	dis
Jessie's sister ruined Christmas for eberybody	That's crappy	N/A	N/A	dis
Paul made us all leave and walk out in the cold	That's shitty	N/A	N/A	dis
I've just been told I'll get a raise next month	That's great	N/A	N/A	dis
Simon threw a brick through my car window	That's awful	N/A	N/A	dis
Liverpool have just won the league by a record margin	That's amazing	N/A	N/A	dis
My bike tyre got a puncture while it was raining	That's terrible	N/A	N/A	dis
Our cousin scored 100% on his piano exam	That's very surprising	N/A	N/A	dis
Michelle accidentally sat on her birthday cake	That's really funny	N/A	N/A	dis
I thought I'd won the race but I got disqualified	That's too bad	N/A	N/A	dis
Dan tripped over his laces and fell in a pond	That's so hilarious	N/A	N/A	dis
Martha won the race ahead of her sister	That's quite unsurprising	N/A	N/A	dis
My daughter broke the schools 100m record on Wednesday	That's amazing	N/A	N/A	dis

Appendix B

B.1 Animacy/Gradability Experiment I - Animacy: Test Nouns

Noun	Log frequency	Dom POS	Length	Syllables	Index
dogs	4.862609677	noun	4	1	1
bear	4.877436431	noun	4	1	2
chef	4.69521426	noun	4	1	3
dog	5.168214255	noun	3	1	4
kids	5.188761359	noun	4	1	5
son	5.165646064	noun	3	1	6
daughter	4.971346491	noun	8	2	7
dad	5.414337051	noun	3	1	8
neighbour	4.23644051	noun	9	2	9
maid	3.771647488	noun	4	1	10
friend	5.285428471	noun	6	1	11
cats	4.415067013	noun	4	1	12
teacher	4.67542065	noun	7	2	13
nurse	4.497824984	noun	5	1	14
twins	4.082947527	noun	5	1	15
shoes	4.659523554	noun	5	1	1
wind	4.997525073	noun	4	1	2
heat	4.764959683	noun	4	1	3
car	5.441200895	noun	3	1	4
room	5.598429239	noun	4	1	5
van	4.688068285	noun	3	1	6
carpet	4.249133236	noun	6	2	7
key	5.073824605	adjective	3	1	8
snowstorm	2.737128411	noun	9	2	9
rain	5.083178946	noun	4	1	10
truck	4.216219259	noun	5	1	11
toys	4.335421387	noun	4	1	12
curtain	3.851981917	noun	7	2	13
storm	4.463040044	noun	5	1	14
paint	4.704846532	noun	5	1	15

B.2 Animacy/Gradability Experiment I - Animacy stimuli

Animate sentence	Inanimate sentence
I just saw that the MOD dogs were on the sofa again	I just saw that the MOD shoes were on the sofa again
I just heard that the MOD bear broke the fence clean off	I just heard that the MOD wind broke the fence clean off
Later I saw that the MOD chef burnt my steak all over	Later I saw that the MOD heat burnt my steak all over
Earlier I heard that the MOD dog was making lots of noise	Earlier I heard that the MOD car was making lots of noise
Then I remembered that the MOD kids ruined the holiday for us	Then I remembered that the MOD room ruined the holiday for us
Then I realised that our MOD son was covered in ice cream	Then I realised that our MOD van was covered in ice cream
Then I realised that our MOD daughter smelt of something really bad	Then I realised that our MOD carpet smelt of something really bad
I walked in and my MOD dad was lying on our bed	I walked in and my MOD key was lying on our bed
I woke up and the MOD neighbour had torn down the tree	I woke up and the MOD snowstorm had torn down the tree
I went outside and the MOD maid had filled the swimming pool	I went outside and the MOD rain had filled the swimming pool
I'm so angry that my MOD friend is caught in the snow	I'm so angry that my MOD truck is caught in the snow
I'm so annoyed that my MOD cats are still in their box	I'm so annoyed that my MOD toys are still in their box
I'm quite angry that my MOD teacher is now in my kitchen	I'm quite angry that my MOD curtain is now in my kitchen
I couldn't believe that the MOD nurse made her fall over again	I couldn't believe that the MOD storm made her fall over again
I'm so shocked that the MOD twins stained the carpet on Friday	I'm so shocked that the MOD paint stained the carpet on Friday

B.3 Animacy/Gradability Experiment I - Gradability: Test adjectives

Adjective	Noun	Dom POS	LOG freq	Valency	Arousal	A-N	f-A	p-A	LD	Index
inconsistent	map	adjective	3.23481182	3.32	3.2	0	0	0	0.97	1
geographical	map	adjective	3.39730371	6.21	3.18	0	0	0	0.94	1
memorable	table	adjective	3.81300702	6.54	4.83	0	1	0	1	2
hexagonal	table	adjective	2.65952355	NA	NA	0	0	0	0.73	2
smooth	frame	adjective	4.35896536	6.42	2.76	0	0	1	0.97	3
purple	frame	adjective	4.28512675	5.6	4.05	0	0	2	1	3
creaky	bookcase	adjective	2.69137092	NA	NA	0	0	0	0.79	4
wooden	bookcase	adjective	4.29496359	5.27	4	0	0	1	0.91	4
innovative	camera	adjective	3.77091758	6.38	4.17	0	0	0	0.97	5
electronic	camera	adjective	3.83968684	6.52	3.64	2	0	0	1	5
enormous	window	adjective	4.62346529	5.68	5.05	4	2	1	0.97	6
circular	window	adjective	3.67208571	5.47	3.4	4	1	0	1	6
improbable	solution	adjective	2.89163538	4.05	3.95	1	0	0	0.81	7
impossible	solution	adjective	4.60928468	3.5	4.52	0	1	10	0.94	7
bitter	beer	adjective	4.18681714	3.63	4.64	1	0	0	0.94	8
German	beer	adjective	4.7628438	NA	NA	0	2	0	0.97	8
cheap	cake	adjective	4.64925381	5.24	4.47	1	0	6	0.97	9
Swiss	cake	adjective	3.96055355	NA	NA	0	0	5	0.97	9

Continued on next page

Table B.1 – continued from previous page

Adjective	Noun	Dom POS	LOG freq	Valency	Arousal	A-N	f-A	p-A	LD	Index
fascinating	fossil	adjective	4.47243691	7.37	5.14	1	0	0	1	10
prehistoric	fossil	adjective	3.48531644	5.68	4.55	1	0	0	0.91	10
faded	rug	verb	3.58447669	NA	NA	0	0	0	0.97	11
woven	rug	verb	3.32412466	NA	NA	2	0	0	0.94	11
ancient	vase	adjective	4.62972889	7.24	3.48	0	0	1	0.97	12
British	vase	adjective	5.48663284	NA	NA	1	1	2	0.97	12
formal	jacket	adjective	4.16878453	5.37	4.48	2	0	0	0.97	13
hooded	jacket	adjective	2.90525074	NA	NA	2	0	0	#N/A	13
casual	event	adjective	3.69443088	6.05	3	0	0	1	0.97	14
weekly	event	adjective	4.01650495	5.43	2.86	4	0	0	0.94	14
unstable	metal	adjective	3.61218967	3.43	5.05	0	1	0	0.97	15
magnetic	metal	adjective	3.70262943	6.17	3.55	0	0	0	1	15
sensational	table	adjective	3.92130344	7.21	5.7	0	0	0	0.94	16
rectangular	table	adjective	3.17285698	5.25	2.7	2	0	0	1	16
starchy	stew	adjective	2.52181053	NA	NA	0	0	0	0.85	17
Russian	stew	adjective	4.45176294	NA	NA	2	2	1	1	17
flexible	knife	adjective	3.98129304	6.74	4.45	9	0	0	1	18
serrated	knife	adjective	2.71692503	NA	NA	11	0	0	0.5	18

B.4 Animacy/Gradability Experiment I - Gradability stimuli

Preceding context	Modifier	Gradable	Non-gradable	Following context	Index
I realised that the map was	fucking/probably	inconsistent	geographical	and that it had mistakes	1
I learnt that the table was	fucking/probably	memorable	hexagonal	and that it had stripes	2
I learnt that the frame was	fucking/probably	smooth	purple	and that it was pretty	3
I realised that the bookcase was	fucking/probably	creaky	wooden	and that it had grooves	4
I realised that the camera was	fucking/probably	innovative	electronic	and that it was tiny	5
I learnt that the window was	fucking/probably	enormous	circular	and that it was cheap	6
I realised that the solution was	fucking/probably	improbable	impossible	and that it was stupid	7
I realised that the the beer was	fucking/probably	bitter	German	and that it was strong	8
I realised that the cake was	fucking/probably	cheap	Swiss	and that it had icing	9
I learnt that the fossil was	fucking/probably	fascinating	prehistoric	and that it was priceless	10
I learnt that the rug was	fucking/probably	faded	woven	and that it had fringes	11
I learnt that the vase was	fucking/probably	ancient	Chinese	and that it was patterned	12
I realised that the jacket was	fucking/probably	formal	hooded	and that it was torn	13
I realised that the metal was	fucking/probably	unstable	magnetic	and that it was thick	14
I learnt that the table was	fucking/probably	sensational	rectangular	and that it was long	15
I realised that the book was	fucking/probably	ornate	Russian	and that it was old	16
I realised that the knife was	fucking/probably	flexible	serrated	and that it was sharp	17
I learnt that the event was	fucking/probably	casual	weekly	and that it was boring	18

B.5 Animacy/Gradability Experiment II - Animacy Stimuli: Animate sentences

Animate sentences
I got home at about 6 o'clock. Then I saw that the MOD dogs were on the sofa again.
I just got off the phone with John. Apparently the MOD bear broke the fence clean off.
I was really looking forward to dinner. Then the MOD chef completely burnt my steak all over.
I've working from home this week. The MOD dog was making lots of noise.
I got a great night's sleep. I woke up and the MOD neighbour had torn down the tree.
I had just finished eating lunch. I went outside and the MOD maid had filled the swimming pool.
We were supposed to meet at midday. Now my MOD friend is stuck in the snow.
I hoped to have everything unpacked by now. For some reason my MOD cats are still in their box.
We were looking at photo albums earlier. Then I remembered that the MOD kids ruined the holiday for us.
We walked back to the car park at 4. Then I saw that our MOD son was covered in ice cream.
I was cleaning the whole house yesterday. Then I realised that our MOD daughter smelt of something really bad.
My husband opened the door for me. I walked in and my MOD dad was lying on our bed.
I had a tough day today. Then I saw that my MOD teacher was in my kitchen.
I went to visit my grandma in hospital earlier. I couldn't believe that the MOD nurse made her fall over again.
I'd just got home from work. I didn't expect to see that the MOD twins had stained the carpet.
I've been stuck at home all day. The MOD plummer made things really uncomfortable.

B.6 Animacy/Gradability Experiment II - Animacy Stimuli: Inanimate sentences

Inanimate sentences
I got home at about 6 o'clock. I saw that the MOD shoes were on the sofa again.
I just got off the phone with John. Apparently the MOD wind broke the fence clean off.
I was really looking forward to dinner. Then the MOD heat completely burnt my steak all over.
I've working from home this week. The MOD car was making lots of noise.
I got a great night's sleep. I woke up and the MOD snowstorm had torn down the tree.
I had just finished eating lunch. I went outside and the MOD rain had filled the swimming pool.
We were supposed to meet at midday. Now my MOD truck is stuck in the snow.
I hoped to have everything unpacked by now. I'm so annoyed that my MOD toys are still in their box.
We were looking at photo albums earlier. Then I remembered that the MOD room ruined the holiday for us.
We walked back to the car park at 4. Then I realised that our MOD van was covered in ice cream.
I was cleaning the whole house yesterday. Then I realised that our MOD carpet smelt of something really bad.
My husband opened the door for me. I walked in and my MOD key was lying on our bed.
I had a tough day today. Then I saw that my MOD curtain was in my kitchen.
I went to visit my grandma in hospital earlier. I couldn't believe that the MOD storm made her fall over again.
I'd just got home from work. I didn't expect to see that the MOD paint had stained the carpet.
I've been stuck at home all day. The MOD cold made things really uncomfortable.

References

Aaronson, D., & Scarborough, H. S. (1976). Performance theories for sentence coding: Some quantitative evidence. *Journal of Experimental Psychology: Human perception and performance*, *2*(1), 56.

Abramowicz, Ł. (2007). Sociolinguistics meets exemplar theory: Frequency and recency effects in (ing). *University of Pennsylvania Working Papers in Linguistics*, *13*(2), 3.

Acton, E. (2019). Pragmatics and the social life of the English definite article. *Language*, *95*(1), 37–65.

Acton, E. (2020). *Pragmatics and the third wave.* Cambridge: Cambridge University Press.

Acton, E., & Potts, C. (2014). That straight talk: Sarah Palin and the sociolinguistics of demonstratives. *Journal of Sociolinguistics*, *18*(1), 3–31.

Agha, A. (2003). The social life of cultural value. *Language & communication*, *23*(3-4), 231–273.

Albers, C., & Lakens, D. (2018). When power analyses based on pilot data are biased: Inaccurate effect size estimators and follow-up bias. *Journal of experimental social psychology*, *74*, 187–195.

Algom, D., Chajut, E., & Lev, S. (2004). A rational look at the emotional stroop phenomenon: a generic slowdown, not a stroop effect. *Journal of experimental psychology: General*, *133*(3), 323.

Al-Hindawe, J. (1996). Considerations when constructing a semantic differential scale. *La Trobe papers in linguistics*, *9*(7), 1–9.

Amaral, P., Roberts, C., & Smith, E. A. (2007). Review of the logic of conventional implicatures by Chris Potts. *Linguistics and Philosophy*, *30*(6), 707–749.

Anderson, A. K. (2005). Affective influences on the attentional dynamics supporting awareness. *Journal of experimental psychology: General*, *134*(2), 258.

Anderson, E. (1997). Violence and the inner-city street code. *Violence and childhood in the inner city*, 1–30.

Andersson, L., & Trudgill, P. (2007). Swearing. In L. Monaghan & J. Goodman (Eds.), *A cultural approach to interpersonal communication.* Oxford, UK: Blackwell.

Andics, A., McQueen, J. M., Petersson, K. M., Gál, V., Rudas, G., & Vidnyánszky, Z. (2010). Neural mechanisms for voice recognition. *Neuroimage*, *52*(4), 1528–1540.

Anwyl-Irvine, A., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. (2018). *Gorilla in our midst: An online behavioral experiment builder.* Behavior Research Methods.

Aquino, J. M., & Arnell, K. M. (2007). Attention and the processing of emotional words: Dissociating effects of arousal. *Psychonomic bulletin & review*, *14*(3), 430–435.

Aronoff, M. (1976). Word formation in generative grammar. *Linguistic Inquiry Monographs Cambridge, Mass.*(1), 1–134.

Aryani, A., Conrad, M., Schmidtke, D., & Jacobs, A. (2018). Why 'piss' is ruder than 'pee'? the role of sound in affective meaning making. *PloS one*, *13*(6), e0198430.

Aston, C. (2021, Feb). *Are you more f$#king, f*kin' or f%!ken?* BuzzFeed. Retrieved from `https://www.buzzfeed.com/clareaston/are-you-more-fucking-fuckin-or-fucken`

Aston, G., & Burnard, L. (1998). *The BNC handbook: exploring the British National Corpus with SARA*. Capstone.

Attardo, S., Eisterhold, J., Hay, J., & Poggi, I. (2003). *Multimodal markers of irony and sarcasm*. Walter de Gruyter GmbH & Co. KG Berlin, Germany.

Auckle, T. (2017). Youth discourse in multilingual Mauritius: The pragmatic significance of swearing in multiple languages. *Stellenbosch Papers in Linguistics Plus*, *51*(1), 61–86.

Bach, K. (1999). The myth of conventional implicature. *Linguistics and philosophy*, 327–366.

Bailey, G. (2015). *Orthographic reflections of (ing): A Twitter-based corpus study.* (Manchester Forum in Linguistics, University of Manchester, 6th November)

Bailey, G. (2019). Emerging from below the social radar: Incipient evaluation in the North West of England. *Journal of Sociolinguistics*, *23*(1), 3–28.

Balota, D. A., Yap, M. J., Hutchison, K. A., Cortese, M. J., Kessler, B., Loftis, B., ... Treiman, R. (2007). The English lexicon project. *Behavior research methods*, *39*(3), 445–459.

Barrett, L. F. (1998). Discrete emotions or dimensions? the role of valence focus and arousal focus. *Cognition & Emotion*, *12*(4), 579–599.

Baruch, Y., & Jenkins, S. (2007). Swearing at work and permissive leadership culture: When anti-social becomes social and incivility is acceptable. *Leadership & Organization Development Journal*, *28*(6), 492–507.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48. doi: 10.18637/jss.v067.i01

Bayard, D., & Krishnayya, S. (2001). Gender, expletive use, and context: Male and female expletive use in structured and unstructured conversation among New Zealand university students. *Women and Language*, *24*(1), 1.

Bayer, J. (1991). German particles in a modular grammar: Neurolinguistic evidence. *Discourse Particles: Descriptive and Theoretical Investigations on the Logical, Syntactic and Pragmatic Properties of Discourse Particles in German., ed. by Werner Abraham*, 253–302.

Bayer, M., Sommer, W., & Schacht, A. (2010). Reading emotional words within sentences: the impact of arousal and valence on event-related potentials. *International Journal of Psychophysiology*, *78*(3), 299–307.

Beaton, M. E., & Washington, H. B. (2015). Slurs and the indexical field: The pejoration and reclaiming of favelado 'slum-dweller'. *Language Sciences*, *52*, 12–21.

Beers Fägersten, K. (2012). *Who's swearing now?: the social aspects of conversational swearing.* Cambridge Scholars Publishing.

Bell, A., & Holmes, J. (1992). H-droppin': Two sociolinguistic variables in New Zealand English. *Australian Journal of Linguistics*, *12*(2), 223–248.

Bellert, I. (1977). On semantic and distributional properties of sentential adverbs. *Linguistic Inquiry*, *8*(2), 337–351.

Beltrama, A. (2018). Totally between subjectivity and discourse. Exploring the pragmatic side of intensification. *Journal of Semantics*, *35*(2), 219–261.

Beltrama, A. (2020). Social meaning in semantics and pragmatics. *Language and Linguistics Compass*, *14*(9), e12398.

Beltrama, A., & Bochnak, M. R. (2015). Intensification without degrees cross-linguistically. *Natural Language & Linguistic Theory*, *33*(3), 843–879.

Beltrama, A., & Staum Casasanto, L. (2017). Totally tall sounds totally younger: Intensification at the socio-semantics interface. *Journal of sociolinguistics*, *21*(2), 154–182.

Beltrama, A., & Trotzke, A. (2019). Conveying emphasis for intensity: Lexical and syntactic strategies. *Language and Linguistics Compass*, *13*(7), e12343.

Bender, E. M. (2005). On the boundaries of linguistic competence: Matched-guise experiments as evidence of knowledge of grammar. *Lingua*, *115*(11), 1579–1598.

Benor, S. B. (2001). The learned/t: phonological variation in Orthodox Jewish English. *University of Pennsylvania Working Papers in Linguistics*, *7*(3), 2.

Bergen, B. K. (2016). *What the f: What swearing reveals about our language, our brains, and ourselves*. Hachette UK.

Bertels, J., Kolinsky, R., Pietrons, E., & Morais, J. (2011). Long-lasting attentional influence of negative and taboo words in an auditory variant of the emotional stroop task. *Emotion*, *11*(1), 29.

Bestelmeyer, P. E., Rouger, J., DeBruine, L. M., & Belin, P. (2010). Auditory adaptation in vocal affect perception. *Cognition*, *117*(2), 217–223.

Blakemore, D. (2002). *Relevance and linguistic meaning: The semantics and pragmatics of discourse markers* (Vol. 99). Cambridge University Press.

Blakemore, D. (2011). On the descriptive ineffability of expressive meaning. *Journal of Pragmatics*, *43*(14), 3537–3550.

Boersma, P. (2019). Praat: doing phonetics by computer. *http://www. praat. org/*.

Breadmore, H. L., Krott, A., & Olson, A. C. (2014). Agreeing to disagree: Deaf and hearing children's awareness of subject–verb number agreement. *Quarterly Journal of Experimental Psychology*, *67*(3), 474–498.

Brown, P., & Levinson, S. C. (1987). *Politeness: Some universals in language usage* (Vol. 4). Cambridge university press.

Brysbaert, M., & Stevens, M. (2018). Power analysis and effect size in mixed effects models: A tutorial. *Journal of cognition*, *1*(1).

Buchanan, T. W., Etzel, J. A., Adolphs, R., & Tranel, D. (2006). The influence of autonomic arousal and semantic relatedness on memory for emotional words. *International journal of psychophysiology*, *61*(1), 26–33.

Bucholtz, M. (2001). The whiteness of nerds: Superstandard English and racial markedness. *Journal of linguistic anthropology*, *11*(1), 84–100.

Buchstaller, I. (2006). Social stereotypes, personality traits and regional perception displaced: Attitudes towards the 'new' quotatives in the UK. *Journal of Sociolinguistics*, *10*(3), 362–381.

Bukmaier, V., Harrington, J., & Kleber, F. (2014). An analysis of post-vocalic /s-ʃ/ neutralization in Augsburg German: evidence for a gradient sound change. *Frontiers in psychology*, *5*, 828.

Burnett, H. (2014). A delineation solution to the puzzles of absolute adjectives. *Linguistics and philosophy*, *37*(1), 1–39.

Burnett, H. (2017a). *Gradability in natural language: Logical and grammatical foundations* (Vol. 7). Oxford University Press.

Burnett, H. (2017b). Sociolinguistic interaction and identity construction: The view from game-theoretic pragmatics. *Journal of Sociolinguistics*, *21*(2), 238–271.

Bybee, J. (2002). Phonological evidence for exemplar storage of multiword sequences. *Studies in second language acquisition*, *24*(2), 215–221.

Cai, Z. G., Gilbert, R. A., Davis, M. H., Gaskell, M. G., Farrar, L., Adler, S., & Rodd, J. M. (2017). Accent modulates access to word meaning: Evidence for a speaker-model account of spoken word recognition. *Cognitive Psychology*, *98*, 73–101.

Callier, P. R. (2013). *Linguistic context and the social meaning of voice quality variation* (Unpublished doctoral dissertation). Georgetown University.

Campbell-Kibler, K. (2005). *Listener perceptions of sociolinguistic variables: The case of (ING)* (Unpublished doctoral dissertation). Citeseer.

Campbell-Kibler, K. (2007). Accent,(ING), and the social logic of listener perceptions. *American speech*, *82*(1), 32–64.

Campbell-Kibler, K. (2009). The nature of sociolinguistic perception. *Language Variation and Change*, *21*(1), 135–156.

Campbell-Kibler, K. (2010a). The effect of speaker information on attitudes toward (ING). *Journal of Language and Social Psychology*, *29*(2), 214–223.

Campbell-Kibler, K. (2010b). Sociolinguistics and perception. *Language and linguistics compass*, *4*(6), 377–389.

Campbell-Kibler, K. (2010c). The sociolinguistic variant as a carrier of social meaning. *Language Variation and Change*, *22*(3), 423–441.

Campbell-Kibler, K. (2011). Intersecting variables and perceived sexual orientation in men. *American Speech*, *86*(1), 52–68.

Campbell-Kibler, K. (2016). Towards a cognitively realistic model of meaningful sociolinguistic variation. *Awareness and control in sociolinguistic research*, 123–151.

Casasanto, L. S. (2008). Does social information influence sentence processing? In *Proceedings of the annual meeting of the cognitive science society* (Vol. 30).

Cavazza, N., & Guidetti, M. (2014). Swearing in political discourse: why vulgarity works. *Journal of Language and Social Psychology*, *33*(5), 537–547.

Chambers, J. K., & Trudgill, P. (1998). *Dialectology*. Cambridge University Press.

Cheshire, J. (1982). Variation in an English dialect: A sociolinguistic study. *Cambridge Studies in Linguistics London*, *37*.

Cheshire, J. (1987). *Syntactic variation, the linguistic variable, and sociolinguistic theory.* Walter de Gruyter, Berlin/New York.

Cheshire, J. (2005). Syntactic variation and beyond: Gender and social class variation in the use of discourse-new markers. *Journal of Sociolinguistics*, *9*(4), 479–508.

Chomsky, N. (1965). *Aspects of the theory of syntax* (Vol. 11). MIT press.

Christie, C. (2013). The relevance of taboo language: An analysis of the indexical values of swearwords. *Journal of Pragmatics*, *58*, 152–169.

Citron, F. M., Weekes, B. S., & Ferstl, E. C. (2013). Effects of valence and arousal on written word recognition: Time course and erp correlates. *Neuroscience letters*, *533*, 90–95.

Clark, L., & Schleef, E. (2010). The acquisition of sociolinguistic evaluations among Polish-born adolescents learning English: Evidence from perception. *Language Awareness*, *19*(4), 299–322.

Coates, J. (1986). *Women, men and language: A sociolinguistic account of gender differences in language*. Routledge.

Cohen, J. (1992). Statistical power analysis. *Current directions in psychological science*, *1*(3), 98–101.

Corver, N. (2008). Uniformity and diversity in the syntax of evaluative vocatives. *The Journal of Comparative Germanic Linguistics*, *11*(1), 43–93.

Coughlan, S. (2021, Jun). More swearing but parents want children protected. *BBC News*. Retrieved from `https://www.bbc.co.uk/news/education-57419263`

Cruse, D. A. (1986). *Lexical semantics*. Cambridge university press.

Culpeper, J. (1996). Towards an anatomy of impoliteness. *Journal of pragmatics*, *25*(3), 349–367.

Culpeper, J. (2011). *Impoliteness: Using language to cause offence* (Vol. 28). Cambridge University Press.

Culpeper, J. (2018). Taboo language and impoliteness. *The Oxford Handbook of Taboo Words and Language*, *10*.

Dahl, O., & Fraurud, K. (1996). Animacy in grammar and discourse. *Pragmatics and Beyond New Series*, 47–64.

Daly, N., Holmes, J., Newton, J., & Stubbe, M. (2004). Expletives as solidarity signals in FTAs on the factory floor. *Journal of Pragmatics*, *36*(5), 945–964.

Davies, C. (2021, Aug). A bloody shame: Britons find a new favourite swearword. *The Guardian*. Retrieved from `https://www.theguardian.com/uk-news/2021/aug/20/bloody-loses-its-place-atop-swearing-lexicon-as-british-use-of-expletives-falls`

DeFrank, M., & Kahlbaugh, P. (2019). Language choice matters: When profanity affects how people are judged. *Journal of Language and Social Psychology*, *38*(1), 126–141.

De Klerk, V. (1997). The role of expletives in the construction of masculinity. *Language and masculinity*, 144–158.

de Vries, R., Hilbig, B., Zettler, I., Dunlop, P., Holtrop, D., Kibeom, L., & Ashton, M. (2018). Still Misinterpreting Lie Scales: Reply to Feldman's Rejoinder. *Unpublished manuscript. Retrieved from https://drive. google. com/open*.

de Vries, R. E., Hilbig, B. E., Zettler, I., Dunlop, P. D., Holtrop, D., Lee, K., & Ashton, M. C. (2018). Honest people tend to use less—not more—profanity: Comment on

Feldman et al.'s (2017) Study 1. *Social psychological and personality science*, *9*(5), 516–520.

Dewaele, J.-M. (2004). The emotional force of swearwords and taboo words in the speech of multilinguals. *Journal of multilingual and multicultural development*, *25*(2-3), 204–222.

Dewaele, J.-M. (2010). Christ fucking shit merde! *Language preferences for swearing among maximally proficient multilinguals. Sociolinguistic Studies*, *4*(3), 595–614.

Dewaele, J.-M. (2015). British 'bollocks' versus American 'jerk': Do native British English speakers swear more–or differently–compared to American English speakers? *Applied Linguistics Review*, *6*(3), 309–339.

Dewaele, J.-M. (2016). Thirty shades of offensiveness: L1 and LX English users' understanding, perception and self-reported use of negative emotion-laden words. *Journal of Pragmatics*, *94*, 112–127.

Dewaele, J.-M. (2017). Self-reported frequency of swearing in English: Do situational, psychological and sociobiographical variables have similar effects on first and foreign language users? *Journal of Multilingual and Multicultural Development*, *38*(4), 330–345.

Dhooge, E., & Hartsuiker, R. J. (2011). How do speakers resist distraction? evidence from a taboo picture-word interference task. *Psychological Science*, *22*(7), 855–859.

Dines, E. R. (1980). Variation in discourse: "and stuff like that". *Language in society*, 13–31.

Dingemanse, M., & Thompson, B. (2020). Playful iconicity: Structural markedness underlies the relation between funniness and iconicity. *Language and Cognition*, *12*(1), 203–224.

Donahoo, S. A. (2019). *Does swearing slow us down or speed us up? a comparison of descriptive and expressive adjectives in sentence processing.* Conference Talk at Swearing in Scandanavia 2019, Södertörn University.

Donahoo, S. A., & Lai, V. T. (2020). The mental representation and social aspect of expressives. *Cognition and Emotion*, 1–16.

D'Onofrio, A. (2015). Persona-based information shapes linguistic perception: Valley Girls and California vowels. *Journal of Sociolinguistics*, *19*(2), 241–256.

D'Onofrio, A. (2018). Personae and phonetic detail in sociolinguistic signs. *Language in Society*, *47*(4), 513–539.

D'Onofrio, A. (2019). Complicating categories: Personae mediate racialized expectations of non-native speech. *Journal of Sociolinguistics*.

Dörre, L., Czypionka, A., Trotzke, A., & Bayer, J. (2018). The processing of German modal particles and their counterparts. *Linguistische Berichte*(255), 58–91.

Dowty, D. (1991). Thematic proto-roles and argument selection. *language*, *67*(3), 547–619.

Drummond, R. (2020). Teenage swearing in the UK. *English World-Wide*, *41*(1), 59–88.

Dynel, M. (2012). Swearing methodologically: the (im)politeness of expletives in anonymous commentaries on Youtube. *Journal of English studies*, *10*, 25–50.

Eckert, P. (1989). *Jocks and burnouts: Social categories and identity in the high school*. Teachers college press.

Eckert, P. (2000). *Language variation as social practice: The linguistic construction of identity in Belten High*. Wiley-Blackwell.

Eckert, P. (2008). Variation and the indexical field. *Journal of sociolinguistics*, *12*(4), 453–476.

Eckert, P. (2019). The limits of meaning: Social indexicality, variation, and the cline of interiority. *Language*, *95*(4), 751–776.

Eckert, P., & McConnell-Ginet, S. (1992). Think practically and look locally: Language and gender as community-based practice. *Annual review of anthropology*, *21*(1), 461–488.

Eckert, P., & McConnell-Ginet, S. (2013). *Language and gender*. Cambridge University Press.

Eilola, T. M., & Havelka, J. (2011). Behavioural and physiological responses to the emotional and taboo Stroop tasks in native and non-native speakers of English. *International Journal of Bilingualism*, *15*(3), 353–369.

Elfner, E., & Kimper, W. (2008). Reduplication without red: Evidence from diddly-infixation. In *Proceedings of the 27th west coast conference on formal linguistics* (Vol. 150, p. 158).

Fairchild, S., Mathis, A., & Papafragou, A. (2020). Pragmatics and social meaning: Understanding under-informativeness in native and non-native speakers. *Cognition*, *200*, 104171.

Fanfuckingtastic. (2021). Retrieved from `https://www.urbandictionary.com/define.php?term=Fanfuckingtastic`

Feldman, G., Lian, H., Kosinski, M., & Stillwell, D. (2017). Frankly, we do give a damn: The relationship between profanity and honesty. *Social psychological and personality science*, *8*(7), 816–826.

Fischer, J. L. (1958). Social influences on the choice of a linguistic variant. *Word*, *14*(1), 47–56.

Forrest, J. (2017). The dynamic interaction between lexical and contextual frequency: A case study of (ING). *Language Variation and Change*, *29*(2), 129–156.

Fortin, A. (2011). *The morphology and semantics of expressive affixes* (Unpublished doctoral dissertation). Oxford University, UK.

Foss, D. J. (1969). Decision processes during sentence comprehension: Effects of lexical item difficulty and position upon decision times. *Journal of Verbal Learning and Verbal Behavior*, *8*(4), 457–462.

Foulkes, P. (2010). Exploring social-indexical knowledge: A long past but a short history. *Laboratory Phonology*, *1*(1), 5–39.

Foulkes, P., & Docherty, G. (2006). The social life of phonetics and phonology. *Journal of phonetics*, *34*(4), 409–438.

Frauenfelder, U. H., & Segui, J. (1989). Phoneme monitoring and lexical processing: Evidence for associative context effects. *Memory & Cognition*, *17*(2), 134–140.

Frazier, L., Dillon, B., & Clifton, C. (2015). A note on interpreting damn expressives: Transferring the blame. *Language and Cognition*, *7*(2), 291–304.

Gauthier, M., & Guille, A. (2017). Gender and age differences in swearing. *Advances in Swearing Research: New languages and new contexts*, *282*, 137.

Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: an integrative review of implicit and explicit attitude change. *Psychological bulletin*, *132*(5), 692.

Gibson, E. (2000). The dependency locality theory: A distance-based theory of linguistic complexity. *Image, language, brain*, *2000*, 95–126.

Giles, H. (1970). Evaluative reactions to accents. *Educational review*, *22*(3), 211–227.

Glick, P., & Fiske, S. T. (1996). The ambivalent sexism inventory: Differentiating hostile and benevolent sexism. *Journal of personality and social psychology*, *70*(3), 491.

Gold, E., & McIntyre, D. (2016). What the f***: An acoustic-pragmatic analysis of meaning in The Wire. *The Journal of the Acoustical Society of America*, *140*(4), 3228–3228.

Gordon, E. (1997). Sex, speech, and stereotypes: Why women use prestige speech forms more than men. *Language in society*, 47–63.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: attitudes, self-esteem, and stereotypes. *Psychological review*, *102*(1), 4.

Grice, H. P. (1975). Logic and conversation. In *Speech acts* (pp. 41–58). Brill.

Grosser, G. S., & Walsh, A. A. (1966). Sex differences in the differential recall of taboo and neutral words. *The Journal of psychology*, *63*(2), 219–227.

Guillet, R., & Arndt, J. (2009). Taboo words: The effect of emotion on memory for peripheral information. *Memory & Cognition*, *37*(6), 866–879.

Gutzmann, D. (2011). Expressive modifiers and mixed expressives. *Empirical issues in syntax and semantics*, *8*, 123–141.

Gutzmann, D. (2015). *Use-conditional meaning: Studies in multidimensional semantics* (Vol. 6). OUP Oxford.

Gutzmann, D. (2019). *The grammar of expressivity* (Vol. 72). Oxford University Press.

Gutzmann, D., & Turgay, K. (2014). Expressive intensifiers and external degree modification. *The Journal of Comparative Germanic Linguistics*, *17*(3), 185–228.

Gutzmann, D., & Turgay, K. (2019). Secondary content: an introduction. In *Secondary content* (pp. 1–25). Brill.

Hadley, C. B., & MacKay, D. G. (2006). Does emotion help or hinder immediate memory? arousal versus priority-binding mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(1), 79.

Haiman, J. (1998). *Talk is cheap: Sarcasm, alienation, and the evolution of language*. Oxford University Press on Demand.

Hair, M., & Ozcan, T. (2018). How reviewers' use of profanity affects perceived usefulness of online reviews. *Marketing Letters*, *29*(2), 151–163.

Harris, C. L., Ayçíçeğí, A., & Gleason, J. B. (2003). Taboo words and reprimands elicit greater autonomic reactivity in a first language than in a second language. *Applied Psycholinguistics*, *24*(4), 561–579.

Harris, J. A., & Potts, C. (2009). Perspective-shifting with appositives and expressives. *Linguistics and Philosophy*, *32*(6), 523–552.

Hay, J. (2018). Sociophonetics: The role of words, the role of context, and the role of words in context. *Topics in cognitive science*, *10*(4), 696–706.

Hay, J., Jannedy, S., & Mendoza-Denton, N. (1999). Oprah and /ay/: Lexical frequency, referee design and style. In *Proceedings of the 14th international congress of phonetic sciences* (pp. 1389–1392).

Hay, J., Warren, P., & Drager, K. (2006). Factors influencing speech perception in the context of a merger-in-progress. *Journal of phonetics*, *34*(4), 458–484.

Hazen, K. (2008). (ING): A vernacular baseline for English in Appalachia. *American Speech*, *83*(2), 116–140.

Hess, L. (2018). Perspectival expressives. *Journal of Pragmatics*, *129*, 13–33.

Hobbs, J. R. (1979). Coherence and coreference. *Cognitive science*, *3*(1), 67–90.

Hoeksema, J. (2019). Taboo terms and their grammar. In *The Oxford handbook of taboo words and language*.

Hoey, E. M., Hömke, P., Löfgren, E., Neumann, T., Schuerman, W. L., & Kendrick, K. H. (2020). Using expletive insertion to pursue and sanction in interaction. *Journal of Sociolinguistics*.

Holmes, J. (2018). Sociolinguistics vs pragmatics. *Pragmatics and its Interfaces*, *294*, 11.

Houston, A. C. (1985). *Continuity and change in English morphology: The variable (ING)*.

Howell, J. L., & Giuliano, T. A. (2011). The effect of expletive use and team gender perceptions of coaching effectiveness. *Journal of Sport Behavior*, *34*(1).

Hughes, G. (1998). *Swearing: A social history of foul language, oaths and profanity in English*. Penguin UK.

Hughes, S. E. (1992). Expletives of lower working-class women. *Language in Society*, *21*(2), 291–303.

Hunt, M., & Acton, E. K. (2022). "how's the wife?": Pragmatic reasoning in spousal reference. *Journal of Pragmatics*, *188*, 152-170. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0378216621003830` doi: https://doi.org/10.1016/j.pragma.2021.11.005

Hymes, D. (1974). *Foundations in sociolinguistics*. Philadelphia: University of Pennsylvania Press.

Irwin, P. (2014). So [totally] speaker-oriented: An analysis of "Drama SO". *Microsyntactic Variation in North American English*, 29–70.

Irwin, P. (2015). Expressive meaning in an AAE attributive construction. *Language Sciences*, *50*, 12–29.

Jackendoff, R. (1978). *Grammar as evidence for conceptual structure*. MIT Press, Cambridge, MA.

Jackendoff, R. (2007). Linguistics in cognitive science: The state of the art. *The Linguistic Review*, *24*(4), 347–401.

Jacobi, L. L. (2014). Perceptions of profanity: How race, gender, and expletive choice affect perceived offensiveness. *North American Journal of Psychology*, *16*(2).

Janschewitz, K. (2008). Taboo, emotionally valenced, and emotionally neutral word norms. *Behavior Research Methods*, *40*(4), 1065–1074.

Jay, K., & Jay, T. (2015). Taboo word fluency and knowledge of slurs and general pejoratives: Deconstructing the poverty-of-vocabulary myth. *Language Sciences*, *52*, 251–259.

Jay, T. (2009a). Do offensive words harm people? *Psychology, public policy, and law*, *15*(2), 81.

Jay, T. (2009b). The utility and ubiquity of taboo words. *Perspectives on Psychological Science*, *4*(2), 153–161.

Jay, T., Caldwell-Harris, C., & King, K. (2008). Recalling taboo and nontaboo words. *The American journal of psychology*, 83–103.

Jay, T., & Janschewitz, K. (2008). The pragmatics of swearing. *Journal of Politeness Research. Language, Behaviour, Culture*, *4*(2), 267–288.

Jegerski, J. (2014). Self-paced reading. *Research methods in second language psycholinguistics*, 20–49.

Jennison, C., & Turnbull, B. W. (1999). *Group sequential methods with applications to clinical trials*. CRC Press.

Johnson, D. I. (2012). Swearing by peers in the work setting: Expectancy violation valence, perceptions of message, and perceptions of speaker. *Communication Studies*, *63*(2), 136–151.

Johnson, D. I., & Lewis, N. (2010). Perceptions of swearing in the work setting: An expectancy violations theory perspective. *Communication Reports*, *23*(2), 106–118.

Johnstone, B., Andrus, J., & Danielson, A. E. (2006). Mobility, indexicality, and the enregisterment of "Pittsburghese". *Journal of English linguistics*, *34*(2), 77–104.

Kachel, S., Radtke, A., Skuk, V. G., Zäske, R., Simpson, A. P., & Steffens, M. C. (2018). Investigating the common set of acoustic parameters in sexual orientation groups: A voice averaging approach. *PloS one*, *13*(12), e0208686.

Kaplan, D. (1999). What is meaning? explorations in the theory of meaning as use. *Brief version–draft*, *1*.

Karttunen, L. (1973). Presuppositions of compound sentences. *Linguistic inquiry*, *4*(2), 169–193.

Karttunen, L., & Peters, S. (1979). Conventional lmplicature. In *Presupposition* (pp. 1–56). Brill.

Kawahara, H., Morise, M., Banno, H., & Skuk, V. G. (2013). Temporally variable multi-aspect N-way morphing based on interference-free speech representations. In *2013 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference* (pp. 1–10).

Kawahara, H., Morise, M., Takahashi, T., Nisimura, R., Irino, T., & Banno, H. (2008). TANDEM-STRAIGHT: A temporally stable power spectral representation for periodic signals and applications to interference-free spectrum, F0, and aperiodicity estimation. In *2008 IEEE International Conference on Acoustics, Speech and Signal Processing* (pp. 3933–3936).

Kawahara, H., Nisimura, R., Irino, T., Morise, M., Takahashi, T., & Banno, H. (2009). Temporally variable multi-aspect auditory morphing enabling extrapolation without objective and perceptual breakdown. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing* (pp. 3905–3908).

Kehler, A. (2002). *Coherence, reference, and the theory of grammar*. CSLI publications Stanford, CA.

Kendall, T. (2013). *Speech rate, pause and sociolinguistic variation: Studies in corpus sociophonetics*. Springer.

Kennedy, C. (2013). *Projecting the adjective: The syntax and semantics of gradability and comparison*. Routledge.

Kennedy, C., & McNally, L. (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 345–381.

Kiesling, S. F. (1998). Men's identities and sociolinguistic variation: The case of fraternity men. *Journal of Sociolinguistics*, *2*(1), 69–99.

Kiesling, S. F. (2009). Style as stance. *Stance: sociolinguistic perspectives*, 171–194.

Kim, J. (2016). Perceptual associations between words and speaker age. *Laboratory Phonology: Journal of the Association for Laboratory Phonology*, *7*(1).

Kleiman, E. (2017). EMAtools: Data Management Tools for Real-Time Monitoring/Ecological Momentary Assessment Data [Computer software manual]. Retrieved from `https://CRAN.R-project.org/package=EMAtools` (R package version 0.1.3)

Knickerbocker, H., Johnson, R. L., & Altarriba, J. (2015). Emotion effects during reading: Influence of an emotion target word on eye movements and processing. *Cognition and Emotion*, *29*(5), 784–806.

Koops, C., Gentry, E., & Pantos, A. (2008). The effect of perceived speaker age on the perception of PIN and PEN vowels in Houston, Texas. *University of Pennsylvania Working Papers in Linguistics*, *14*(2), 12.

Kousta, S.-T., Vinson, D. P., & Vigliocco, G. (2009). Emotion words, regardless of polarity, have a processing advantage over neutral words. *Cognition*, *112*(3), 473–481.

Kratzer, A. (1999). Beyond ouch and oops: How descriptive and expressive meaning interact. In *Cornell conference on theories of context dependency* (Vol. 26).

Kuperman, V., Estes, Z., Brysbaert, M., & Warriner, A. B. (2014). Emotion and language: Valence and arousal affect word recognition. *Journal of Experimental Psychology: General*, *143*(3), 1065.

Labov, W. (1963). The social motivation of a sound change. *Word*, *19*(3), 273–309.

Labov, W. (1966). *The social stratification of English in New York city*. ERIC.

Labov, W. (1972). *Sociolinguistic patterns* (No. 4). University of Pennsylvania Press.

Labov, W. (1994). *Principles of linguistic change. vol. 1: Internal factors*. Oxford: Blackwell.

Labov, W. (2001). Principles of linguistic change volume 2: Social factors. *Language in Society*, *29*.

Labov, W. (2006). *The social stratification of English in New York city*. Cambridge University Press.

Labov, W., Ash, S., Baranowski, M., Nagy, N., Ravindranath, M., & Weldon, T. (2006). Listeners' sensitivity to the frequency of sociolinguistic variables. *University of Pennsylvania Working Papers in Linguistics*, *12*(2), 10.

Labov, W., Ash, S., Ravindranath, M., Weldon, T., Baranowski, M., & Nagy, N. (2011). Properties of the sociolinguistic monitor. *Journal of Sociolinguistics*, *15*(4), 431–463.

Lakens, D. (2014). Performing high-powered studies efficiently with sequential analyses. *European Journal of Social Psychology*, *44*(7), 701–710.

Lambert, W. E., Hodgson, R. C., Gardner, R. C., & Fillenbaum, S. (1960). Evaluational reactions to spoken languages. *The Journal of Abnormal and Social Psychology*, *60*(1), 44.

LaMotte, S. (2021, Apr). Why swearing is a sign of intelligence, helps manage pain and more. *The Observer*. Retrieved from `https://edition.cnn.com/2021/01/26/health/swearing-benefits-wellness/index.html`

Lavandera, B. R. (1978). Where does the sociolinguistic variable stop? *Language in society*, *7*(2), 171–182.

Lawson, R., & Milani, T. (2015). Fight narratives, covert prestige, and performances of "tough" masculinity: Some insights from an urban center. *Language and Masculinities*, 53–76.

Leon, A. C., Davis, L. L., & Kraemer, H. C. (2011). The role and interpretation of pilot studies in clinical research. *Journal of psychiatric research*, *45*(5), 626–629.

Levon, E. (2014). Categories, stereotypes, and the linguistic perception of sexuality. *Language in Society*, *43*(5), 539–566.

Levon, E., & Buchstaller, I. (2015). Perception, cognition, and linguistic structure: The effect of linguistic modularity and cognitive style on sociolinguistic processing. *Language Variation and Change*, *27*(3), 319–348.

Levon, E., & Fox, S. (2014). Social salience and the sociolinguistic monitor: A case study of ING and TH-fronting in Britain. *Journal of English Linguistics*, *42*(3), 185–217.

Levon, E., & Ye, Y. (2019). Language, indexicality and gender ideologies: contextual effects on the perceived credibility of women. *Gender and Language*, *14*(2), 123–151.

Lewis, R. L., & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive science*, *29*(3), 375–419.

Limbrick, P. (1991). A study of male and female expletive use in single and mixed-sex situations. *Te Reo*(34), 71–89.

Ljung, M. (2010). *Swearing: A cross-cultural linguistic study*. Springer.

Love, R. (2017). Bad language revisited: swearing in the Spoken BNC2014. In *Corpus Linguistics 2017 Conference. University of Birminham.*

Love, R. (2021). Swearing in informal spoken English: 1990s–2010s. *Text & Talk*, *41*(5-6), 739–762.

Lutzky, U., & Kehoe, A. (2016). Your blog is (the) shit: A corpus linguistic approach to the identification of swearing in computer mediated communication. *International Journal of Corpus Linguistics*, *21*(2), 165–191.

Lüdecke, D. (2021). sjplot: Data visualization for statistics in social science [Computer software manual]. Retrieved from `https://CRAN.R-project.org/package=sjPlot` (R package version 2.8.7)

MacKay, D. G., Shafto, M., Taylor, J. K., Marian, D. E., Abrams, L., & Dyer, J. R. (2004). Relations between emotion, memory, and attention: Evidence from taboo stroop, lexical decision, and immediate memory tasks. *Memory & Cognition*, *32*(3), 474–488.

Madan, C. R., Shafer, A. T., Chan, M., & Singhal, A. (2017). Shock and awe: Distinct effects of taboo words on lexical decision and free recall. *The Quarterly Journal of Experimental Psychology*, *70*(4), 793–810.

Martin, R. (1997). "girls don't talk about garages!": Perceptions of conversation in same-and cross-sex friendships. *Personal Relationships*, *4*(2), 115–130.

Mathewson, K. J., Arnell, K. M., & Mansfield, C. A. (2008). Capturing and holding attention: The impact of emotional words in rapid serial visual presentation. *Memory & Cognition*, *36*(1), 182–200.

McCarthy, J. J. (1982). Prosodic structure and expletive infixation. *Language*, 574–590.

McCawley, J. D. (1978). Where you can shove infixes. *Syllables and segments*, 213–221.

McCready, E. (2010). Varieties of conventional implicature. *Semantics and Pragmatics*, *3*, 8–1.

McCready, E., & Schwager, M. (2009). Intensifiers. In *Workshop on expressives and other kinds of non-truth-conditional meaning, university of osnabrück.*

McEnery, T. (2004). *Swearing in English: Bad language, purity and power from 1586 to the present*. Routledge.

McEnery, T., & Xiao, Z. (2004). Swearing in modern British English: the case of fuck in the BNC. *Language and Literature*, *13*(3), 235–268.

McGowan, K. B., & Babel, A. M. (2020). Perceiving isn't believing: Divergence in levels of sociolinguistic awareness. *Language in Society*, *49*(2), 231–256.

McLeod, L. (2011). Swearing in the 'tradie' environment as a tool for solidarity. *Griffith Working Papers in Pragmatics and Intercultural Communication*, *4*, 1–10.

Meibauer, J. (2013). Expressive compounds in German. *Word Structure*, *6*(1), 21–42.

Mendoza-Denton, N. (2008). *Homegirls: Language and cultural practice among Latina youth gangs*. John Wiley & Sons.

Mercury, R.-E. (1995). Swearing: A "bad" part of language; a good part of language learning. *TESL Canada journal*, 28–36.

Methven, E. (2020). 'A Woman's Tongue': Representations of Gender and Swearing in Australian Legal and Media Discourse. *Australian Feminist Law Journal*, 1–25.

Mitchell, D. C., & Green, D. W. (1978). The effects of context and content on immediate processing in reading. *Quarterly Journal of Experimental Psychology*, *30*(4), 609–636.

Montagu, A. (1967). *The anatomy of swearing*. University of Pennsylvania press.

Moore, E., & Podesva, R. (2009). Style, indexicality, and the social meaning of tag questions. *Language in Society*, *38*(4), 447–485.

Morzycki, M. (2016). *Modification*. Cambridge University Press.

Myers, S. A., & Knox, R. L. (1999). Verbal aggression in the college classroom: Perceived instructor use and student affective learning. *Communication Quarterly*, *47*(1), 33–45.

Napoli, D. J., & Hoeksema, J. (2009). The grammatical versatility of taboo terms. *Studies in Language. International Journal sponsored by the Foundation "Foundations of Language"*, *33*(3), 612–643.

Neal-Holder, P. (2020, Jan). Should schools be allowed to ban slang words like 'peng'? *BBC News*. Retrieved from `https://www.bbc.co.uk/news/education-51064279`

Needle, J. M., & Pierrehumbert, J. B. (2018). Gendered associations of English morphology. *Laboratory Phonology*, *9*(1).

Nevins, A., & Vaux, B. (2003). Metalinguistic, shmetalinguistic: The phonology of shmreduplication. In *Proceedings from the annual meeting of the chicago linguistic society* (Vol. 39, pp. 702–721).

Newman, M. L., Pennebaker, J. W., Berry, D. S., & Richards, J. M. (2003). Lying words: Predicting deception from linguistic styles. *Personality and social psychology bulletin*, *29*(5), 665–675.

Niedzielski, N. (1999). The effect of social information on the perception of sociolinguistic variables. *Journal of language and social psychology*, *18*(1), 62–85.

Norris, D., & McQueen, J. M. (2008). Shortlist B: a bayesian model of continuous speech recognition. *Psychological review*, *115*(2), 357.

Nygaard, L. C., & Lunders, E. R. (2002). Resolution of lexical ambiguity by emotional tone of voice. *Memory & cognition*, *30*(4), 583–593.

Ofcom (Office of Communications). (2016). *Attitudes to potentially offensive language and gestures on TV and radio: Quick reference guide.* London: Ipsos MORI Social Research Institute. Retrieved from `https://www.ofcom.org.uk/research-and-data/tv-radio-and-on-demand/tv-research/offensive-language-2016`

ONS (Office for National Statistics). (2010). *National Statistics Socio-economic Classification (NS-SEC) Coding Tool* (Tech. Rep.). Retrieved from `http://www.neighbourhood.statistics.gov.uk/HTMLDocs/dev3/ONS_NSSEC_discovery`

Oxford English Dictionary. (2020). Oxford University Press.

Pahl, R. (2018). GroupSeq: A GUI-Based Program to Compute Probabilities Regarding Group Sequential Designs [Computer software manual]. Retrieved from `https://CRAN.R-project.org/package=GroupSeq` (R package version 1.3.5)

Pantos, A. J., & Perkins, A. W. (2013). Measuring implicit and explicit attitudes toward foreign accented speech. *Journal of Language and Social Psychology*, *32*(1), 3–20.

Paradise, L. V., Cohl, B., & Zweig, J. (1980). Effects of profane language and physical attractiveness on perceptions of counselor behavior. *Journal of counseling psychology*, *27*(6), 620.

Patton, C. L., Asken, M., Fremouw, W. J., & Bemis, R. (2017). The influence of police profanity on public perception of excessive force. *Journal of Police and Criminal Psychology*, *32*(4), 340–357.

Pharao, N., Maegaard, M., Møller, J. S., & Kristiansen, T. (2014). Indexical meanings of [s+] among Copenhagen youth: Social perception of a phonetic variant in different prosodic contexts. *Language in Society*, 1–31.

Pierrehumbert, J. B. (2001). Exemplar dynamics: Word frequency, lenition and contrast. *Typological studies in language*, *45*, 137–158.

Pierrehumbert, J. B. (2002). Word-specific phonetics. *Laboratory phonology*, *7*.

Pietilä, P., Tainio, L., Lappalainen, S., & Lahelma, E. (2020). Swearing as a method of antipedagogy in workshops of rap lyrics for 'failing boys' in vocational education. *Gender and Education*, 1–15.

Podesva, R. J. (2008). Three sources of stylistic meaning. In *Texas linguistic forum* (Vol. 51, pp. 1–10).

Podesva, R. J., Roberts, S. J., & Campbell-Kibler, K. (2002). Sharing resources and indexing meanings in the production of gay styles. *Language and sexuality: Contesting meaning in theory and practice*, 175–189.

Potts, C. (2005). *The logic of conventional implicatures* (No. 7). Oxford University Press on Demand.

Potts, C. (2007). The expressive dimension. *Theoretical linguistics*, *33*(2), 165–198.

Potts, C., Asudeh, A., Cable, S., Hara, Y., McCready, E., Alonso-Ovalle, L., ... Roeper, T. (2009). Expressives and identity conditions. *Linguistic Inquiry*, *40*(2), 356–366.

Potts, C., & Kawahara, S. (2004). Japanese honorifics as emotive definite descriptions. In *Semantics and linguistic theory* (Vol. 14, pp. 253–270).

Preston, D. R. (1999). *Handbook of perceptual dialectology* (Vol. 1). John Benjamins Publishing.

Prolific. (2022). *Prolific academic.* Retrieved from `https://www.prolific.co`

Qualtrics, I. (2013). Qualtrics. *Provo, UT, USA*.

R Core Team. (2018). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from `https://www.R-project.org/`

Rassin, E., & Heijden, S. V. D. (2005). Appearing credible? swearing helps! *Psychology, Crime & Law*, *11*(2), 177–182.

Rastall, P. (2004). Playful English: kinds of reduplication. *English Today*, *20*(4), 38–41.

Rastle, K., Harrington, J., & Coltheart, M. (2002). 358,534 nonwords: The ARC nonword database. *The Quarterly Journal of Experimental Psychology Section A*, *55*(4), 1339–1362.

Reilly, J., Kelly, A., Zuckerman, B. M., Twigg, P. P., Wells, M., Jobson, K. R., & Flurie, M. (2020). Building the perfect curse word: A psycholinguistic investigation of the form and meaning of taboo words. *Psychonomic Bulletin & Review*, *27*(1), 139–148.

Revelle, W. (2021). psych: Procedures for psychological, psychometric, and personality research [Computer software manual]. Evanston, Illinois. Retrieved from `https://CRAN.R-project.org/package=psych` (R package version 2.1.3)

Roberts, C. (1996). Information structure in discourse: Toward an integrated for-al theory of pragmatics. In J. H. Yoon & A. Kathol (Eds.), *Osuwpl, vol. 49, 91–136.*

Roberts, C., Davies, E., & Jupp, T. (2014). *Language and discrimination*. Routledge.

Romaine, S. (1980). The relative clause marker in Scots English: Diffusion, complexity, and style as dimensions of syntactic change. *Language in society*, 221–247.

Romaine, S. (1999). *Communicating gender*. Psychology Press.

Sabri, O., & Obermiller, C. (2012). Consumer perception of taboo in ads. *Journal of business research*, *65*(6), 869–873.

Säily, T. (2011). Variation in morphological productivity in the BNC: Sociolinguistic and methodological considerations. *Corpus Linguistics and Linguistic Theory*, *7*, 119–141.

Saka, P. (2007). *How to think about meaning* (Vol. 109). Springer Science & Business Media.

Sanchez, T. (2008). Accountability in morphological borrowing: Analyzing a linguistic subsystem as a sociolinguistic variable. *Language Variation and Change*, *20*(2), 225–253.

Sankoff, G. (1973). Above and beyond phonology in variable rules. In *New ways of analysing variation english* (pp. 42–62). Washington, D.C., Georgetown Uni Press.

Scherer, C. R., & Sagarin, B. J. (2006). Indecent influence: The positive effects of obscenity on persuasion. *Social Influence*, *1*(2), 138–146.

Schleef, E., & Flynn, N. (2015). Ageing meanings of (ing): age and indexicality in Manchester, England. *English World-Wide*, *36*(1), 48–90.

Schleef, E., Flynn, N., & Barras, W. (2017). Regional diversity in social perceptions of (ING). *Language Variation and Change*, *29*(1), 29–56.

Schleef, E., Meyerhoff, M., & Clark, L. (2011). Teenagers' acquisition of variation: A comparison of locally-born and migrant teens' realisation of english (ing) in edinburgh and london. *English world-wide*, *32*(2), 206–236.

Schlenker, P. (2007). Expressive presuppositions. *Theoretical Linguistics*, *33*(2), 237–245.

Schlenker, P. (2010). Supplements within a unidimensional semantics ii: Epistemic status and projection. In *Proceedings of northeastern linguistic society 2009, glsa* (p. 167—182).

Schmidt, R. W. (1990). The role of consciousness in second language learning. *Applied linguistics*, *11*(2), 129–158.

Schütze, C. T. (2009). Web searches should supplement judgements, not supplant them. *Zeitschrift für Sprachwissenschaft*, *28*(1), 151–156.

Schütze, C. T., & Sprouse, J. (2014). Judgment data. *Research methods in linguistics*, *27*, 50.

Schwarz, F. (2007). Processing presupposed content. *Journal of Semantics*, *24*(4), 373–416.

Schwarz, F. (2015). Presuppositions vs. asserted content in online processing. In *Experimental perspectives on presuppositions* (pp. 89–108). Springer.

Schweinberger, M. (2018). Swearing in Irish English–a corpus-based quantitative analysis of the sociolinguistics of swearing. *Lingua*, *209*, 1–20.

Scott, G. G., O'Donnell, P. J., & Sereno, S. C. (2012). Emotion words affect eye fixations during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*(3), 783.

Searle, J. R. (1969). *Speech acts: An essay in the philosophy of language* (Vol. 626). Cambridge university press.

Seizer, S. (2011). On the uses of obscenity in live stand-up comedy. *Anthropological Quarterly*, *84*(1), 209–234.

Siegel, D. C. (1974). *Topics in English morphology.* (Unpublished doctoral dissertation). Massachusetts Institute of Technology.

Siegrist, M. (1995). Effects of taboo words on color-naming performance on a stroop test. *Perceptual and Motor Skills*, *81*(3_suppl), 1119–1122.

Silverstein, M. (1976). Hierarchy of features and ergativity. In R. Nixon (Ed.), *Grammatical categories in australian languages.*

Silverstein, M. (2003). Indexical order and the dialectics of sociolinguistic life. *Language & communication*, *23*(3-4), 193–229.

Simons, M., Tonhauser, J., Beaver, D., & Roberts, C. (2010). What projects and why. In *Semantics and linguistic theory* (Vol. 20, pp. 309–327).

Skuk, V. G., & Schweinberger, S. R. (2014). Influences of fundamental frequency, formant frequencies, aperiodicity, and spectrum level on the perception of voice gender. *Journal of Speech, Language, and Hearing Research*.

Snell, J. (2010). From sociolinguistic variation to socially strategic stylisation. *Journal of Sociolinguistics*, *14*(5), 630–656.

Sobieraj, S. (2018). Bitch, slut, skank, cunt: Patterned resistance to women's visibility in digital publics. *Information, Communication & Society*, *21*(11), 1700–1714.

Spencer-Oatey, H. (2008). *Culturally speaking second edition: Culture, communication and politeness theory*. Bloomsbury Publishing.

Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition* (Vol. 142). Citeseer.

Sprouse, J. (2018). Acceptability judgments and grammaticality, prospects and challenges. In *Syntactic structures after 60 years* (pp. 195–224). De Gruyter Mouton.

Sprouse, J., & Almeida, D. (2012). Assessing the reliability of textbook data in syntax: Adger's core syntax. *Journal of Linguistics*, *48*(3), 609–652.

Sprouse, J., Schütze, C. T., & Almeida, D. (2013). A comparison of informal and formal acceptability judgments using a random sample from linguistic inquiry 2001–2010. *Lingua*, *134*, 219–248.

Squires, L. (2016). Processing grammatical differences: Perceiving versus noticing. *Awareness and control in sociolinguistic research*, 80–103.

Stapleton, K. (2003). Gender and swearing: A community practice. *Women and Language*, *26*(2), 22.

Stapleton, K. (2010). Swearing. *Interpersonal pragmatics*, *6*, 289.

Stapleton, K. (2020). Swearing and perceptions of the speaker: A discursive approach. *Journal of Pragmatics*, *170*, 381–395.

Stenström, A.-B. (2017). Swearing in English and Spanish teenage talk. *Advances in swearing research. Amsterdam: John Benjamins*, 157–182.

Stephens, R., Atkins, J., & Kingston, A. (2009). Swearing as a response to pain. *Neuroreport*, *20*(12), 1056–1060.

Steriopolo, O. (2009). Form and function of expressive morphology: a case study of Russian. *Russian Language Journal*, *59*, 149–194.

Strand, E. A. (1999). Uncovering the role of gender stereotypes in speech perception. *Journal of language and social psychology*, *18*(1), 86–100.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of experimental psychology*, *18*(6), 643.

Sulpizio, S., Toti, M., Del Maschio, N., Costa, A., Fedeli, D., Job, R., & Abutalebi, J. (2019). Are you really cursing? neural processing of taboo words in native and foreign language. *Brain and language*, *194*, 84–92.

Szewczyk, J. M., & Schriefers, H. (2011). Is animacy special?: ERP correlates of semantic violations and animacy violations in sentence processing. *Brain research*, *1368*, 208–221.

Tagliamonte, S. (2004). Someth [in]'s go [ing] on!: Variable ing at ground zero. *Language Variation in Europe: Papers from ICLAVE*, *2*, 390–403.

Tamminga, M. (2017). Matched guise effects can be robust to speech style. *The Journal of the Acoustical Society of America*, *142*(1), EL18–EL23.

Terkourafi, M. (2008). *Toward a unified theory of politeness, impoliteness, and rudeness*. De Gruyter Mouton.

Tessier, A.-M., & Becker, M. (2018). Vowel but not consonant identity and the very informal English lexicon. In *Proceedings of the annual meetings on phonology* (Vol. 5).

Thelwall, M. (2008). Fk yea i swear: cursing and gender in MySpace. *Corpora*, *3*(1), 83–107.

Thomas, E. (2010). *Sociophonetics: An introduction*. Macmillan International Higher Education.

Thomas, L., & LaBar, K. (2005). Emotional arousal enhances word repetition priming. *Cognition & Emotion*, *19*(7), 1027–1047.

Tiemann, S., Schmid, M., Bade, N., Rolke, B., Hertrich, I., Ackermann, H., ... Beck, S. (2011). Psycholinguistic evidence for presuppositions: On-line and off-line data. In *Proceedings of sinn und bedeutung* (Vol. 15, pp. 581–596).

Time for united's moral minority to move on. (2011, Apr). *The Daily Mail*. Retrieved from `https://www.pressreader.com/uk/daily-mail/20110427/288548789144083`

Trudgill, P. (1974). *The social differentiation of English in Norwich* (Vol. 13). CUP Archive.

Tuft, S. E., McLennan, C. T., & Krestar, M. L. (2018). Hearing taboo words can result in early talker effects in word recognition for female listeners. *Quarterly Journal of Experimental Psychology*, *71*(2), 435–448.

Tyler, J. C. (2015). Expanding and mapping the indexical field: Rising pitch, the uptalk stereotype, and perceptual variation. *Journal of English Linguistics*, *43*(4), 284–310.

Van Berkum, J. J., Van den Brink, D., Tesink, C. M., Kos, M., & Hagoort, P. (2008). The neural integration of speaker and message. *Journal of cognitive neuroscience*, *20*(4), 580–591.

Van Heuven, W. J., Mandera, P., Keuleers, E., & Brysbaert, M. (2014). SUBTLEX-UK: A new and improved word frequency database for British English. *The Quarterly Journal of Experimental Psychology*, *67*(6), 1176–1190.

Vaughn, C. (2021). *Internal constrains and the sociolinguistic monitor.* Retrieved from `https://nors.ku.dk/english/calendar/2021/5th-variation-and-language-processing-conference/` (5th Variation and Language Processing Conference)

Vaughn, C., & Kendall, T. (2018). Listener sensitivity to probabilistic conditioning of sociolinguistic variables: The case of (ING). *Journal of Memory and Language*, *103*, 58–73.

Vinson, D., Ponari, M., & Vigliocco, G. (2014). How does emotional content affect lexical processing? *Cognition & emotion*, *28*(4), 737–746.

Walker, A., & Hay, J. (2011). Congruence between 'word age'and 'voice age' facilitates lexical access. *Laboratory Phonology*, *2*(1), 219–237.

Wang, W., Chen, L., Thirunarayan, K., & Sheth, A. P. (2014). Cursing in English on Twitter. In *Proceedings of the 17th acm conference on computer supported cooperative work & social computing* (pp. 415–425).

Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior research methods*, *45*(4), 1191–1207.

Watts, E. (2006). *Mobility-induced dialect contact: A sociolinguistic investigation of speech variation in Wilmslow, Cheshire* (Unpublished doctoral dissertation). University of Essex.

Weiner, E. J., & Labov, W. (1983). Constraints on the agentless passive. *Journal of linguistics*, 29–58.

Wells, J. C. (1982). *Accents of English: Volume 1* (Vol. 1). Cambridge University Press.

Westbury, C., & Hollis, G. (2019). Wriggly, squiffy, lummox, and boobs: What makes some words funny? *Journal of Experimental Psychology: General*, *148*(1), 97.

Wickham, H. (2016). *ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York. Retrieved from `https://ggplot2.tidyverse.org`

Williams, A., Kerswill, P., Foulkes, P., & Docherty, G. (1999). Dialect levelling: change and continuity in Milton Keynes, Reading and Hull. *Urban voices: Accent studies in the British Isles*, 141–162.

Woods, K. J., Siegel, M. H., Traer, J., & McDermott, J. H. (2017). Headphone screening to facilitate web-based auditory experiments. *Attention, Perception, & Psychophysics*, *79*(7), 2064–2072.

Wyczesany, M., & Ligeza, T. S. (2015). Towards a constructionist approach to emotions: verification of the three-dimensional model of affect with EEG-independent component analysis. *Experimental brain research*, *233*(3), 723–733.

Yamamoto, M. (1999). *Animacy and reference*. Amsterdam Philadelphia.

Yardy, B. J. (2010). *Sound symbolism, sonority, and swearing: An affect induction perspective* (Unpublished doctoral dissertation). Lethbridge, Alta.: University of Lethbridge, Dept. of Psychology, 2010.

Yougowords.com. (2019). *Yougowords.com.* Retrieved from `http://www.yougowords.com/speech/adjective/3-syllables`

Zaenen, A., Carletta, J., Garretson, G., Bresnan, J., Koontz-Garboden, A., Nikitina, T., ... Wasow, T. (2004). Animacy encoding in English: Why and how. In *Proceedings of the workshop on discourse annotation* (pp. 118–125).

Zhang, Q. (2008). Rhotacization and the 'Beijing Smooth Operator': The social meaning of a linguistic variable 1. *Journal of Sociolinguistics*, *12*(2), 201–222.

Zheng, Y., & Samuel, A. G. (2017). Does seeing an Asian face make speech sound more accented? *Attention, Perception, & Psychophysics*, *79*(6), 1841–1859.

Ziegler, J., & Pylkkänen, L. (2016). Scalar adjectives and the temporal unfolding of semantic composition: An MEG investigation. *Neuropsychologia*, *89*, 161–171.

Zwicky, A. M., & Pullum, G. K. (1987). Plain morphology and expressive morphology. In *Annual Meeting of the Berkeley Linguistics Society* (Vol. 13, pp. 330–340).