Investigating the phonetic and linguistic features used by speakers to communicate an intent to harm

Sarah Kelly

Doctor of Philosophy

University of York

Language and Linguistic Science

September 2018

Abstract

This research aims to examine the phonetic and linguistic features which can be associated with a threatening intent. At present, there is a range of threat assessment resources and descriptions in legal cases which provide insight surrounding the content and production of threatening language. However, the veracity of these descriptions has not been thoroughly explored in empirical research. Through the examination of authentic and simulated threatening language data, this research provides a broad overview of the usage of phonetic and linguistic features to convey a threatening intent to harm. A set of 10 authentic speech recordings where a direct (or explicitly-worded) threat was present were analysed in relation to a sample of non-threatening speech. In addition, simulated threatening and non-threatening speech and texts were collected from 41 participants under experimental conditions. These threatening and non-threatening data were compared with respect to mean fundamental frequency, intensity, articulation rate and changes to vocal tract features and vocal settings. The simulated data were also examined for the use of lexical features which have previously been associated with the actualisation of harm.

The results of these analyses suggest that there is no compelling evidence to support the assertion of a 'threatening tone of voice'. There were, however, tendencies for these speakers to raise their mean fundamental frequency, intensity and articulation rate during threatening speech production relative to their non-threatening speech. There was also evidence to suggest that a number of lexical features used by these participants also corresponded to previous examinations of authentic threatening texts. It is suggested that on the basis of these findings, the production of threatening language is a considerably more complex and varied behaviour than might be expected. These findings have notable implications for the development of threat assessment tools, and for the description of a 'threatening manner' in legal contexts.

Contents

List of Tables

List of Figures

Acknowledgements

First and foremost, I would like to give special thanks to my supervisor Dr. Dominic Watt. When we first met, I was so full of insecurities and self-doubt. I have no doubt that these qualities wormed their way into our meetings and frustrated him to no end. With his guidance and seemingly never-ending patience over the years, Dom has helped me to complete this thesis to the best of my ability, and to become a more self-assured person in the process. I'll always be grateful for his support.

I'd also like to thank my fellow threat speech researcher James Tompkinson. Maybe without realising it, he has helped me to persist with this research more than just about anybody. Watching his achievements over the years has pushed me to be the best researcher that I can be.

Thanks too to all the staff based in the Language and Linguistic Science department at The University of York. I've been fortunate enough to have spent the last eight years studying there, and I wouldn't have changed it for the world.

My thanks also go out to Dr Harriet Smith, Dr Natalie Braber and Dr David Wright based at Nottingham Trent University, as well as to Jeremy Robson over at DeMontfort University. Our brief time working together made me feel as though there was a place for me in academia for the very first time.

Finally, the completion of this thesis would not have been possible without the support of my friends and family. I really hope that I have done you all proud.

Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as references.

I would like to dedicate this thesis to my late mum and grandma.

1 Introduction

The current study presents a series of investigations into threatening language. Firstly, examining whether there are phonetic features which can be associated with spoken threats which differ from non-threatening speech. Secondly, by examining whether authors of written threats opt for specific lexical and/or grammatical cues (jointly referred to as 'linguistic features' in this research) when communicating a threatening intent. The advantages and disadvantages to the use of authentic and simulated data are also considered, as is the use of data provided by actors and non-actors. Finally, the use of typographical features in written threatening texts is discussed in relation to the texts subsequent phonetic production. It is hoped that by incorporating both forensic phonetic and forensic linguistic approaches to this little-understood speech act, this study can further our understanding of threatening language, and encourage future collaborations for the study of other language-based offences.

The following subsection outlines the structure of this thesis.

1.1 Overview of the thesis

Chapter 2 provides an overview of the existing literature on threatening behaviour. In particular, the methods used by practitioners of various disciplines to determine when a threat has been made. The chapter begins with data which suggests that threatening behaviour is a common experience. As such, there are a large number of resources available to help victims assess the risk of a threat, and/or how to report the threat to an authority figure. These resources will be outlined, along with a discussion about their usability for a lay-person audience. The chapter then provides examples of threat assessment procedures which are used in workplace settings. These procedures are designed to report threatening behaviour targeted toward a workplace. These threat assessment procedures provide a rationale for examining threatening language and its relation to the expression of anger or aggression.

As this research is motivated by threats which occur in a forensic setting, case law taken from criminal cases in the United Kingdom and the United States is discussed. These cases mostly centre on establishing whether or not a specific utterance or text should be considered illegal. Following this, the chapter draws on the existing linguistic literature into threatening language. From this literature, it becomes apparent that there is currently no research which has examined the phonetic features of threatening speech. However, there are predictions made about the phonetic features which might characterise threats. These predictions are derived from an evolutionary perspective and provide a basis for the phonetic features selected for analysis in this thesis. There is, however, substantially more research about the lexical and grammatical features of written threatening language in the forensic linguistic domain. For the purposes of this research, the term 'linguistic features' refers to the lexical and grammatical features of texts or speech, in order to differentiate those from 'phonetic features' which describe properties which are only relevant to speech production. In particular, this research will consider previous attempts to build threat detection systems and the linguistic features associated with terrorist communications. This chapter also provides a rationale for examining threats produced by both males and females. Literature which explores the idea that threatening behaviour is learned or is inherent is also discussed in this chapter.

Based on this literature, a number of gaps in the domain of threat production are identified. Most notably, the lack of research which examines whether threats differ from non-threatening language in terms of the lexical and grammatical features (linguistic features) used, as well as whether threatening and non-threatening speech differ in terms of their phonetic correlates. With these gaps in the literature in mind, the following research questions are posed in Chapter 2.4:

1. What phonetic features can be associated with a threatening tone of voice?
   - (a) Are the phonetic features under investigation in this study comparable to previous descriptions of emotional speech? In particular, how do the phonetic features that are associated with threatening speech in this research relate to the phonetic features associated with angry speech in previous literature?
   - (b) Are there phonetic features which differ between the production of direct and indirect threats? Is there a compensatory effect present during the production of indirect threats?
   - (c) Do these phonetic features differ between authentic and simulated threats?
   - (d) Do these phonetic features differ between speakers with acting experience, and speakers with no formal acting experience?
   - (e) Do these phonetic features differ between male and female speakers?

2. Are there linguistic features present in the simulated threats collected for this study which relate to the linguistic features said to predict the likelihood that violence will be actualised?

    (a) Do the written direct and indirect threats conform to the definitions of 'direct' and 'indirect' threats coined in this study? If not, how could the current definitions of direct and indirect threats be changed to better reflect this difference?

3. Is there a relationship between the typographical features of a scripted threat and its subsequent phonetic production?

These questions intend to address the gaps of knowledge highlighted in Chapter 2.

Chapter 3 begins with an overview of the how the current study defines various types of threatening language. Details are then provided about the data that were collected for this study, how these data were analysed both phonetically and linguistically, and justifies these approaches. From the existing literature on threatening speech presented in Chapter 2, it is apparent that there is currently no agreed upon method for the collection of threatening language data. As such, this thesis will use approaches that have been established in research about the phonetic features of emotional speech. The methodological constraints posed by collecting emotional and threatening speech data are also discussed.

Chapters 4 and 5 details the results of the data analysis. This analysis initially examines the lexical and grammatical properties of the threatening texts collected for this research under experimental conditions. Later this thesis explores whether there is empirical evidence to support the notion of a 'threatening tone of voice.' For each phonetic and linguistic feature analysed for this study, the results will be immediately followed by a discussion relating to that feature. This chapter is broadly divided into three sections: the linguistic analysis of simulated written threats, the phonetic features of authentic and simulated spoken threats, and examining the relationship between the typographical features in the simulated handwritten threatening texts and the author's subsequent phonetic production.

This thesis concludes in Chapter 6. This chapter briefly summarises the main findings of this research, and in addition, recommendations for future research in this area. The results of this research are then more thoroughly discussed in relation to the research questions posed in Chapter 2. The original contributions of this research are emphasised, and the results are placed in the broader context of other relevant literature. Finally, suggestions are offered for improvements to the data collection and data analysis discussed in this thesis.

At present, there is no single, widely-accepted definition of a threat. Nor is there a single, agreed method for determining whether a threat has been made. Chapter 2 describes various threat assessment, legal and linguistic approaches to conceptualising threats. In particular, this chapter focuses on how practitioners or experts within these groups differentiate threatening language from non-threatening language. The difficulties and limitations of these approaches are also considered.

2 Literature Review

This chapter will begin with an overview of the extent to which threats have been reported or documented across a variety of workplaces and institutions based both in the United Kingdom as well as in other countries. This is followed by a brief consideration of the issues faced when trying to ascertain the scale of threatening behaviour based on police or government statistics.

This chapter then discusses the use of threatening behaviour across a range of contexts as a way of predicting future harmful acts (§2.0.2). In particular, the notion that making threats has been linked with an increased risk of harm being actualised. This leads into an overview of how lay-people responsible for identifying and/or reporting threatening behaviour have been instructed or trained to perform this task - or whether they received any support or guidance with this role (§2.0.3).

From this point, the chapter outlines a variety of legal cases in the United Kingdom and the United States which have centred on the legality of alleged threatening speech or actions (§2.1). The purpose of describing these cases is to further consider how threats have been defined - this time in a legal context, as opposed to a workplace or school setting. Section 2.2 proceeds to describe the current contributions of linguistic research into defining or identifying threatening language. This ranges from pragmatic approaches which aimed to conceptualise actions performed through speech, to more recent attempts to detect threatening language usage on the internet. The linguistic and phonetic features linked with threatening language in previous relevant literature will also be presented in this chapter (§2.2-2.3). This provides a stronger rationale for the linguistic and phonetic features chosen for analysis in this research.

2.0.1 Prevalence of threatening language and behaviour

Threats can be loosely defined as a communication of an intent to commit harm (Fraser, 1998). While threatening language is often thought of as an inherently criminal act, threats can also serve an important interactive and communicative function. It is not uncommon to hear stories of parents who have threatened to ground their children, or about employees who have made threats to leave their job. What distinguishes these type of relatively harmless acts from criminal acts is not altogether clear. Yet around the world, legal systems are tasked with making such distinctions. Assessing whether or not a 'threatening' statement, utterance, or other expression is a criminal act, is a process which begins at the moment an alleged offence is committed and is continued throughout the legal process. At each stage of the legal process, individuals draw upon their own understanding of what it means to make an illegal threat. Their perception might be based on their own life experience or an innate sense of what threats are.

While there are legal guidance and case law available to assist with these decisions, often this information is vague and/or heavily subject to interpretation. Though it is not fully understood how people conceptualise or produce threats, there is compelling evidence to suggest that threatening language is widely experienced across different areas of society. There are reports from a variety of sources which document the high frequency of threats made towards staff working across a range of sectors. This includes staff working in hospitals (Winstanley and Whittington, 2004; Pich et al., 2011), the emergency services (Archer, 1999), the education system (Wilson, Douglas and Lyon, 2011), the hospitality industry (Karatepe, Yorganci and Haktanir, 2009), and the prison system (Vartia and Hyyti, 2002).

Threats are also directed towards high-profile targets, such as politicians (Mullen at al., 2008) and celebrities (Dietz et al., 1991). Threatening language is also a commonly-reported occurrence on online social media platforms such as Facebook (Kwan and Skoric, 2013) and Twitter (Ronson, 2015). In addition, threats are a commonly reported behaviour in surveys about intimate partner violence (Burczycka and Conroy, 2018). Interestingly, while threats are commonly reported by victims in person-based crime surveys (van Dijk, 2012), these offences are also among the crimes least reported to the police (Van Kesteren et al., 2000).

In the United Kingdom, there are no statistics available for the number of threats reported to the police. Nor is there data on the number of people convicted for making threats. What is publicly-accessible are data for the number of reports and criminal charges of offences which *include* threats among a possible number of offences committed. For crimes which centre on, for example, domestic abuse or public disorder, threats are among a number of offences which are specified in relevant legislation. Therefore, the data discussed in this subsection should not be interpreted as evidence of a rise or fall in threatening behaviour. Instead, these data should be taken as evidence towards an increase or decrease in the reporting and subsequent recording of a wide range of criminal acts (including, but not limited to, threats).

The Office of National Statistics provides data for each calendar year for the crimes recorded by police in England and Wales (sourced from the Office for National Statistics, 2018). A number of types of crime described in these data include threats. For example, for crimes described as 'violence without injury,' there were 586,538 offences recorded between April 2017 and March 2018. For the same crime type, there were around 474,000 offences recorded in the previous calendar year. This crime type consists of threats or conspiracy to commit murder, certain offences against children and assaults without injury. There have also been increases in the reporting of crimes involving knives or sharp instruments. There was a small increase from 2,744 these offences recorded from April 2016 to March 2017, to 2,972 offences recorded in the following year. This crime type includes threats to kill, as well as other crimes such as attempted murder, assault with injury, and rape (all involving a knife or sharp instrument). Threats are also featured in data regarding public order offences. There were 385,864 public order offences recorded by police between April 2017 and March 2018. This represented a 26% increase from the previous calendar year.

The prosecution rates for threat-related crimes have also increased slightly between 2016 and 2017 (Criminal Justice Statistics quarterly, England and Wales 2017). This includes a 2% rise in the number of prosecutions for public order offences (from 17,500 to 17,800). In addition, there is a relatively high rate of guilty pleas for these offences of around 66%.

There is a large body of research devoted to the prevention of violence in which threats are often cited as a predictor of harm. The following subsection discusses the findings of research which examines the link between threats and the actualisation of violence.

2.0.2 Threatening language as a predictor of harmful acts

The relationship between threatening language and harmful acts is not entirely clear. In literature from across a range of disciplines, there is a focus on identifying characteristics that either increase (Breakwell, 1997) or decrease (Mullen, Pathé and Stuart, 1999) the likelihood that a threatener will act upon their threat. In other words, threats are treated as a symptom of a broader problem: the actualisation of violence. This is in contrast with linguistic approaches which often consider threats to be harmful acts in and of themselves. These linguistic approaches mostly focus on defining threats based on their pragmatic function or linguistic features (see Chapter 2.2.5).

As described earlier in this chapter, threats occur in a vast array of (often forensically-relevant) contexts. This has led to research which examines the 'usefulness' of threats as a predictor of violence. For example, there is research which examines how staff working for the police or security staff infer risk. Threats are often flagged up as a behaviour associated with higher risk of harm. This includes threats contributing to police impressions of domestic abuse (Robinson, Pinchevsky and Guthrie, 2016) and campus attacks (Deisinger, et al., 2008).

There is also research which has looked at the behaviours said to predict the likelihood that harm will be realised in other settings. In research about identifying potentially violent patients in a hospital setting, there are a number of phonetic and linguistic features which are associated with the increased risk of violence of patients (Jackson, Wilkes and Luck, 2014; Kim, Ideker and Todicheeney-Mannes, 2011). These features include patients producing speech with an increase in volume, yelling, 'making threats', adopting a 'demeaning' or 'sarcastic' tone of voice, using abusive language (such as swearing) and making 'aggressive' statements. Making threats has also been said to predict the likelihood of violence actualised by psychiatric patients (Linaker and Busch-Iversen, 1995).

It is not clear, however, as to what constitutes a threat in many of these studies. It appears to be common in such research to not define or operationalise threats. As such, it seems that it is left to the individual to identify threats based on their own life experiences, or possibly in reference to their own training. This implies that threats are regarded as an easily identifiable behaviour by both researchers who examine threats as a predictor of violence, and by individuals who participate in such research.

It is worth adding that while threats have been linked with the actualisation of violence, it is possible that this is not the motivation of the threatener. Their motivation might be, for example, to instil fear or to cause disruption. For instance, a threatener might deliberately communicate a bomb threat to inconvenience emergency service staff in the knowledge that there is no such device.

The evidence presented so far in this research has indicated that making and receiving threats (whether criminal or otherwise) appears to be a widespread experience. In addition, it appears that there is no obvious or consistent demographic target to threatening behaviour. People from a wide range of backgrounds are threatened as part of their day-to-day life. In turn, making threats appears to be treated by researchers and threat assessment practitioners as a precursor to violence. This raises the question of whether these victims are using the same, or similar criteria to assess if they consider an action to be threatening. Their assessment might be informed (in part) by publicly-available threat assessment resources. These resources provide information on the characteristics of threatening behaviour that is written in such a way so as to be accessible to a lay-audience.

The following subsection will discuss a range of threat assessment resources that are targeted toward a lay-person audience. Many of these resources are designed to help lay-people to identify threatening behaviour or to record threatening behaviour.

2.0.3 Threat assessment resources for lay- audiences

There are a wide number of threat assessment resources available online for a lay-audience. For example, there are resources that are designed to assist people in abusive relationships, or for people who encounter abuse in the workplace. This section will present and discuss several of these resources. These resources were selected because of their advice regarding the linguistic and phonetic production of threatening language. It should be noted that as these resources were designed for a lay-audience (i.e. not for academic research), there are issues with identifying the specific author(s) of the materials and their level of expertise.

The following extract is taken from from a website which offers advice to nurses based in the United States about predicting violence (Centers for Disease Control and Prevention, 2013):

- Verbal cues: speaking loudly or yelling, swearing, threatening tone of voice
- Non-verbal or behavioural cues: Physical appearance (clothing and hygiene neglected), arms held tight across chest, clenched fists, heavy breathing, pacing or agitation, a terrified look signifying fear and high anxiety, a fixed stare, aggressive

or threatening posture, thrown objects, sudden changes in behavior, indications of drunkenness or substance abuse

From this extract, there are a number of linguistic and phonetic features which are linked to the prediction of violence. These predictors include, the adoption of a so-called 'threatening tone of voice', speaking loudly (or yelling) and using profane words. Currently, there is no widely-accepted definition across a range of disciplines as to what constitutes a threatening tone of voice. It is also not clear whether changes to vocal effort such as shouting and yelling can predict violence in clinical (or other) settings. As will be described later in this chapter, many studies have aimed to link linguistic features to the prediction of violence.

Another application of threat assessment resources is concerned with documenting alleged threats of violence or harm towards a specified location or person(s). For example, a business might receive a threatening phone call. This event might be documented by a lay-person for features which could provide clues to the caller's identity or which indicate that the threat was 'genuine'.

Figure 1 shows an example of a publicly-available threat assessment protocol.

Figure 1: Bomb threat checklist (GOV.UK, 2016)

Protective Marking: Restricted when Completed

Form 5474

ACTIONS TO BE TAKEN ON RECEIPT OF A BOMB THREAT

1 Remain calm and talk to the caller

2 Note the caller's number if displayed on your phone

3 If the threat has been sent via email or social media see appropriate section below

4 If you are able to, record the call

5 Write down the exact wording of the threat:

When Where What How Who Why Time

ASK THESE QUESTIONS & RECORD ANSWERS AS ACCURATELY AS POSSIBLE:

1. Where exactly is the bomb right now?	
2. When is it going to explode?	
3. What does it look like?	
4. What does the bomb contain?	
5. How will it be detonated?	
6. Did you place the bomb? If not you, who did?	
7. What is your name?	
8. What is your address?	
9. What is your telephone number?	
10. Do you represent a group or are you acting alone?	
11. Why have you placed the bomb?	
Record time call completed:	

INFORM BUILDING SECURITY/ COORDINATING MANAGER

Name and telephone number of person informed:	

DIAL 999 AND INFORM POLICE

Time informed:	

This part should be completed once the caller has hung up and police/ building security/ coordinating manager have all been informed

Date and time of call:	
Duration of call:	
The telephone number that received the call:	

ABOUT THE CALLER:		**Male**	**Female**	**Nationality?**		**Age?**
THREAT LANGUAGE:		**Well-spoken**	**Irrational**	**Taped**	**Foul**	**Incoherent**
CALLER'S VOICE:		**Calm**	**Crying**	**Clearing throat**	**Angry**	**Nasal**
Slurred	**Excited**	**Stutter**	**Disguised**	**Slow**	**Lisp**	***Accent**
Rapid	**Deep**	**Familiar**	**Laughter**	**Hoarse**	**Other** *(please specify)*	

***What accent?**	
If the voice sounded familiar, who did it sound like?	

BACKGROUND SOUNDS:		**Street noises**	**House noises**	**Animal noises**	**Crockery**	**Motor**
Clear	**Voice**	**Static**	**PA system**	**Booth**	**Music**	
Factory machinery		**Office machinery**		**Other** *(please specify)*		

REMARKS:

ADDITIONAL NOTES:

Signature: .... Print Name: .... Date:

ACTIONS TO BE TAKEN ON RECEIPT OF A BOMB THREAT SENT VIA EMAIL OR SOCIAL MEDIA

1 DO NOT reply to, forward or delete the message

2 If sent via email note the address

3 If sent via social media what application has been used and what is the username/ID?

4 Dial 999 and follow police guidance

5 Preserve all web log files for your organisations to help the police investigation (as a guide, 7 days prior to the threat message and 48 hours after)

Signature: .... Print Name: .... Date: ....

SAVE AND PRINT – HAND COPY TO POLICE AND SECURITY/ COORDINATING MANAGER

Retention Period: 7 years
MP 925/10

Like many of the publicly-available resources used by institutions or companies for spoken threat assessment, this resource appears to centre on assessing threats made by telephone communications. Generally, these telephone threat assessments are carried out by lay-people who have no formal linguistic training. Typically, their role in the company involves receiving telephone calls from the wider public or other organisations. Such protocols are usually employed as and when a member of staff is handling a suspected threatening communication. With the aid of a protocol, the member of staff typically asks the caller a series of questions which aim to collect information about the nature of the threat (e.g. *'What kind of weapon is it?'*, *'What will make the bomb explode?'*). Questions concerning the identity of the caller (*'What is your name?' 'Who are you affiliated with?'*), or the purpose of the threat (*'Why did you place the bomb?'*) might also be asked. This information would be passed on to police to assist in their subsequent investigations.

Crucially, these resources do not contain information about how to identify a potential threat. It is unclear whether there is accompanying training available which can help users identify threats. If so, it is not known whether threats are associated with particular linguistic and phonetic features. It is also possible that these protocols are used based on subjective decisions about what is threatening to the call receiver. As such, there is a possibility that some threats might not be identified, while other communications are misinterpreted as threatening. It is also not known whether all of the questions or listed voice descriptions have been tested on lay-people. This would help to ensure that these options are likely to be consistently understood in the same way across lay-listeners.

These protocols are also problematic from a phonetic perspective. As will be discussed in Chapter 2.3.4, there is uncertainty about how a speaker's emotional state or personality is encoded within the speech signal. In this research domain, the perception of these qualities from a speech signal is also not a straightforward process. As will be discussed in §2.3.4, there is linguistic research which suggests that inferring emotions based on a voice can be error-prone. Despite this, these protocols appear to operate on several assumptions about speech perception:

- that speakers either communicate their actual emotional state through their speech signal or can convey that they are feeling a specific emotion (even if they are, in fact, experiencing different emotions)
- that listeners are able to accurately infer the speaker's emotional state using some

combination of vocal tract features and word usage

- that information on qualities about a caller's voice and/or word usage is valuable when documenting a threat

These assumptions about language can also be found in other resources for lay-people experiencing threatening behaviour. These resources might offer either: a description of how threats are actualised, or (more commonly) no information whatsoever about what constitutes a threat. In the former approach, threats are reduced to a set of behaviours which emerge when someone is producing a threat. The latter approach appears to be based on a number of possible assumptions about threatening behaviour. Firstly, the assumption that threats are so varied that it is not possible to provide a generalised description. Alternatively, there could be an assumption that lay-people are somehow able to perceive threats accurately. As such, it would not be viewed as necessary to provide resources which explain what threatening behaviour looks or sounds like. This thesis hopes to shed light on whether any of these threat assessment approaches can be substantiated by linguistic analysis.

So far, this research has presented threat assessment resources that are available to lay-people or threat assessment practitioners. It appears that threats are mostly treated as predictors of violence, rather than harmful actions in and of themselves. In addition, advice varies on what constitutes a threat or 'threatening manner'. For some resources, people are advised to trust their instincts when it comes to assessing threats (de Becker, 1997). As shown earlier in this section, in other resources designed for a lay-audience, there are more specific details about threatening behaviour. Relating to this research, there are also descriptions of the vocal production of threats, or features which could reasonably affect speech.

However, there have currently been no publicly-available empirical or peer-reviewed studies which have examined the phonetic features of threatening speech. As such, it is possible that the advice provided in these resources is non-specific or even inaccurate. It is hoped that this research can eventually contribute toward the design of resources that might help people to assess threatening behaviour.

For many countries around the world, decisions about whether or not something could be considered an illegal threat are made by juries consisting of lay-people. It is possible that the assumptions about threatening language presented earlier might feed in to the decision making of jurors. The next section provides an overview of the issues relating to how threatening language evidence might be presented in courts. In addition, legal cases will be discussed which relate to the linguistic and phonetic properties of threats.

2.1 Threatening language in legal settings

In many jurisdictions across the world, threats are an example of a language crime. The term 'language crimes' used in this thesis, refers to utterances or written texts which are considered to be illegal in a given jurisdiction. This deviates slightly from the definition provided by Solan and Tiersma (2005), who use the phrases 'illegal speech acts' and 'language crimes' interchangeably. This is because strictly speaking, the term 'speech acts' only refer to spoken utterances. This thesis will examine data taken from both spoken and written threats. In addition, several of the legal disputes discussed here centre on written language.

In legal disputes centred on voice evidence, there are innumerable references made to a so-called 'tone of voice' by earwitnesses. Some of these descriptions are grounded in the perception of acoustic properties such as 'loud tone of voice' (David Antonio Mendes v DPP [2015]) or 'high pitched voice' (Elizabeth A. Vacura and Raymond Vacura v. Carol L. Plott, 666 F.2d 1200 (8th Cir. 1981)). Other descriptions are considerably more vague and open to interpretation, these could include terms such as *'high tone of voice'*[1]. There are also descriptions which are currently not grounded in linguistic research. Instead, these descriptions seem to derive from folk-linguistic knowledge. For example, inferring the speaker's personality *('sarcastic sounding voice;'* Lewis-Webb v. Qualico Steel Co., Inc., 929 F. Supp. 385 (M.D. Ala. 1996)), or their emotional state *('calm voice;'* United States of America, v. Antonio Clemmons, 461 F.3d 1057 (8th Cir. 2006)) based on vocal tract features alone. The term 'tone of voice' is also used to illustrate the performance of speech acts *('...pleading, persuading tone of voice'*; People v. Miller (Crim. No. 3921. Second Dist., Div. Three. Jan. 18, 1946.)). The presence (or lack of presence) of a 'threatening tone of voice' is also asserted in cases centred on threatening speech (State v. McDowell 620 A.2d 94 (1993); Packer v. Skid Roe, Inc., 938 F. Supp. 193 (S.D.N.Y. 1996); State v. Milner 571 N.W.2d 7 (1997)).

[1] 'High' could feasibly refer to acoustic properties such as volume or pitch (or both). Alternatively, this descriptor could be referring to the effect of substance abuse on the speaker's voice.

In more recent years, there have been legal cases which have acknowledged the limitations of presenting voice descriptions in the form of a transcript as evidence (R v JDP [2010] EWCA Crim 3224). While transcriptions of voice testimony serve as a useful record of an alleged offence, they cannot provide a complete picture of the voice to a court and/or jury panel. Without accompanying audio or video recording, conveying concepts such as 'tone of voice' is challenging [2]. This is the case even when the transcript attempts to describe the manner in which the utterances were spoken. It is possible that, for example, 'high' or 'threatening' voices are perceived differently across a given population. This has significant implications for how a reported language crime is treated in the legal system. Differences in the perception of voice evidence might affect the decisions concerning whether to bring a case to trial, or in the decisions of a jury when assessing guilt.

Despite the gaps in knowledge about speech perception, courts rarely question the validity of voice descriptions that are provided as testimony. A witness's opinion about *how* an utterance sounded, seem to carry enough weight in legal settings, without a need for further scrutiny. This thesis aims to examine whether descriptions such as 'threatening tone of voice,' have any phonetic basis. If there is no evidence to support these descriptions, arguably people working within the legal system should be made aware of the potential for variability among 'threatening' voices. This would allow the evidence to be considered in relation to relevant empirical research.

As previously mentioned in this research, not all threats are necessarily criminal. Only under specific circumstances can a threat be considered to be illegal. Across legal systems in the United Kingdom, there are numerous offences which explicitly reference unlawful, threatening language or behaviour. These range from threats to kill, to the use of threats as a means of procuring unwanted sexual activity from a vulnerable person. In addition, threats overlap considerably with other offences such as abusive language, or those related to terrorism. Due to the scope of these offences, determining whether a threat is illegal will differ depending on the circumstances surrounding an alleged threat. There is also some degree of subjectivity involved in these processes.

This thesis will not be able to provide a truly comprehensive examination of threatening

[2]This limitation also applies to the documentation of witness testimony about visual information, such as facial expressions or body language.

language cases. However, the following subsections will show that there are numerous approaches used by courts across the United Kingdom and the United States to determine whether an illegal threat was made. To this end, numerous relevant legal cases from the United Kingdom and the United States will be described in the following sections.

2.1.1 Relevant legal cases in the United Kingdom

In 2012, there was an English court case which centred on a custody officer's perception of a supposed 'threatening tone of voice'. In this case, the custody officer heard the following utterance being produced by a suspect as he was being removed from his cell:

'When I get out of here [a police cell] *I'm going to do something about this'.*

This utterance was also recorded via a CCTV camera in the cell. The custody officer perceived this utterance to be a reiteration of a threat, and as such, he argued that this made the recorded utterance a threat in its own right. It was also argued by the custody officer that this utterance 'sounded like a threat'.

In contrast, the defendant argued that this comment was not a threat. Moreover, this utterance did not suggest that any harm would be directed toward the custody officer. Because there were no obvious references to harm in this alleged threat, in effect this case became a stalemate between two perspectives: the perception of a threat based on 'tone of voice', and the argument that no threat had been made because the words which were used did not clearly state any harmful action. These opposing arguments can be likened to the locutionary and illocutionary acts described as part of Speech Act Theory (Austin, 1969). An overview of this theory will be described in Chapter 2.2.1.

The case R v Colin Albert Jennison [2000] 2 Cr. App. R. S. 213, centred on a series of allegedly threatening utterances. These 'threats' were made to the relatives of a woman the appellant had been convicted of killing years previously. The appellant was alleged to have said 'I mean what I say, you will die.' Afterwards, he was described by witnesses as raising his voice while saying 'I have killed your sister and I'm going to kill you, that's a promise.' When interviewed by police, Jennison acknowledged that he had spoken to the relatives in question, but said that he did not make any threats to kill. During his trial, the Attorney General listed the following aggravating factors of the threats to kill:

- the threats were directed towards the relatives of a person the appellant had killed
- there were multiple threats made over a period of time
- the threats had a 'real' effect on the recipients, and that this effect was long-lasting
- the appellant expressed no remorse
- the view that the appellant would continue his behaviour

In the subsequent court of appeal trial, it was noted that the most important factor to be considered in this case was the recipient's opinion about the likelihood that the threat would be carried out. In addition, the appellant's violent background, and the grudge held against the recipients of his threats was information that was also considered to be relevant to the case. Notably, the appellant's assertion that he did not make a threat (and as such, was not intending to threaten) was not a successful defence.

There are a number of aspects to this case which are of particular linguistic interest. There were numerous references found in the witness testimony evidence which described the appellant's voice as 'raised' or shouted while he made the alleged threats. There is also the use of the word 'promise' during one of his alleged threats. While threats and promises can serve different pragmatic functions, in this example the court took the view that the utterance overall was an example of a threat. It is possible that the inclusion of a 'promise' might have been viewed by the court as a statement conveying a commitment to carry out the threat. The pragmatic functions of threats and promises will be discussed in Chapter 2.2.4. There are also differences in the use of active voice between the two examples of threats provided: 'I'm going to kill you' (active), and '...you will die' (intransitive). For the case of Jennison, his utterances were viewed by the recipients as evidence that the appellant was likely to carry these harmful acts.

In the case R v Cakmak; R v Cavcav; R v Talay; R v Can; R v Karaaslan; R v Durukanoglu [2002] EWCA Crim 500, six appellants had been convicted of making threats to damage property. At the time of the alleged offence, the men were inside one of the 'pods' or capsules at the London Eye tourist attraction in London. While inside this pod, the appellants had allegedly threatened to set themselves on fire as a form of protest toward perceived injustices by the Turkish political government of the time. This threat had been communicated to the operator of the London Eye using an intercom.

The operator believed that the men posed a genuine risk to the London Eye itself, as well as to the other tourists who were onboard the attraction. The appellants said that they had not intended to harm these other tourists. The appellants were convicted under section 2(a) of the Criminal Justice Act (1971). This act refers to making threats to destroy or damage property. However, there is an argument to be made that these appellants were not threatening the London Eye per se. Instead, it can be suggested that they were in fact threatening to endanger themselves, even if other people and property would also be likely to be endangered. Therefore, the target(s) of threats could be inferred in legal cases as opposed to being explicitly stated.

The case R v Russell Williams [2010] EWCA Crim. 778, centred on a series of alleged death threats. These 'threats' were delivered by Williams in a series of phone calls to his former partner. These included utterances such as *'I am coming home tonight. I will show you what for. You know when I cut your head off tonight. Swear on baby's life. I am coming tonight.'* and *'I am going to cut your head off. I will cut you from your toes to your fucking neck.'* The former partner had only heard these utterances over the phone, as opposed to face-to-face.

These utterances were described in the court of appeal as 'direct'. It is not known made clear what exactly is mean by 'direct'. This could refer to the phone call being delivered *directly* to the alleged target. It is also possible that 'direct' refers to the clear descriptions of violent acts in these communications. In addition, the appellant had described how he had communicated these messages because he was 'emotional and angry'.

This case relates to the current research in a number of ways. Firstly, in the meaning of a 'direct threat'. This term could reasonably refer to a variety of different properties of threatening language. This study intends to collect both indirect and direct threats in order to examine the linguistic features pertaining to each. Secondly, this case relates anger to the expression of threat. In this instance, making a threat is said to be motivated by feelings of anger. Therefore, it is possible that the linguistic and phonetic production of threatening language could be expected to overlap with those of angry speech. This issue will also be examined in more detail in this research.

In the case DPP v Richie Smith [2017] EWHC 3193, the appellant was approached by two police officers at a train station. The officers had been informed of Smith's intoxicated state, and unruly behaviour while on board his train. The appellant was described as being 'quite loud,' speaking in a 'raised voice.' During their interaction, Smith said the following phrase: *'go deal with some niggers and Pakis'.* The officers present found these words offensive and thought others within earshot would have also felt offended. The appellant argued that because he was intoxicated, he did not intend to use those terms and that he regretted saying them. However, Smith was charged with one count contravening section 5 of the Public Order Act 1986, and one count of contravening section 31 of the Crime and Disorder Act 1998 (for racially-aggravated abuse). This case was subsequently overturned by an appeal court for three main reasons.

Firstly, as outlined in section 6 (5) of the Public Order Act 1986, intoxication is not a defence for being unable to express an unlawful intention. Therefore, the consumption of drugs or alcohol is not a mitigating factor (except for unforeseen reactions to prescribed medications or the appellant unknowingly consuming a substance). The original trial also set about proving his intent. Under section 5 of the Public Order Act 1986, it is not necessary to prove an intention to cause harassment or distress. If this offence had been prosecuted under s4A of the same Act, this would have been an appropriate line of enquiry. Finally, the appeal court considered the offensiveness of the terms used by the appellant. It was surmised that while language can be offensive (as these terms were), this does not automatically mean that alarm or distress was caused. This finding has implications for the possible link between the perception of offensive words and threatening language. In particular, whether the use of offensive language can be considered to be threatening.

In Scotland, it is an offence to behave in a threatening or abusive manner which would likely cause a reasonable person to suffer fear or alarm (Criminal Justice and Licensing (Scotland) Act 2010). In addition, it must be shown that the accused either intended their behaviour to cause fear or alarm or 'is reckless as to whether the behaviour would cause fear or alarm.' In 2013, there were two cases where a conviction was secured under this Act. However, the interpretation of the 'reasonable person' part of the legislation differed between these cases.

In Jolly v HMA [2013] SCCR 511, the appellant allegedly made a series of threats during a series of interviews before his release from a young offenders institution. These alleged threats were not directed towards any person present at these interviews, but rather to various people outside of the young offenders institution. In this case, it was surmised that threats did not necessarily have to be subjected to a wholly objective reasonable person test, The effect the threat(s) had on those present would be sufficient evidence to secure a conviction. In contrast, another case, Rooney v Brown [2013] SCCR 334, did apply a reasonable person test. This case centred on a series of abusive comments shouted to police officers during the appellant's arrest. While there was no evidence to suggest that the police officers present felt fear, it was determined that a reasonable person would be likely to react as such. A subsequent appeal on behalf of a number of similar legal cases which followed overruled the decisions made in Jolly v. HM Advocate (Ewan Paterson v PF Airdrie [2014] HCJAC 87). It was recommended that the objective approach taken in Rooney v. Brown was in keeping with the wording of Section 38 of the Criminal Justice and Licensing (Scotland) Act 2010.

Therefore, one approach taken in the Scottish courts is to consider what a reasonable person would interpret from an utterance(s). As such, the witness's account of how they interpreted an alleged threat is not necessarily considered. What exactly would cause a reasonable person to fear or alarm is not apparent. As seen in Chapter 2.0.3 (threat assessment protocols), there seems to be an expectation that people will interpret alleged threats uniformly. Alleged threats can also be understood by witnesses in a way which might be regarded by the courts as 'unreasonable.' Also, a threatening intention must be shown. Alternatively, it must be shown that the appellant lacked insight that his/her behaviour could be interpreted as a threat. As with English and Welsh cases, deciding whether someone held a threatening intent (or should have foreseen this interpretation) seems to be a somewhat arbitrary process.

2.1.2 Relevant legal cases in the United States

In the United States, the First Amendment serves (in part) to protect freedom of speech. 'Speech' here, can include spoken, written or non-linguistic forms of expression (such as body gestures, or images). However, not all speech is protected by the constitution. Examples where speech has been judged to:

- incite imminent unlawful action
- be obscene

- use ‘fighting words’
- facilitate a bribe
- violate trademarks and copyrights
- use false or misleading advertising
- constitute a ‘true threat’

are not protected. Such exceptions to freedom of speech have emerged over time, due to cases being brought to the Supreme Court of the United States. Legal cases have, and continue to focus on determining whether speech has ‘overstepped’ a legal boundary. This can involve the creation, and implementation of a ‘test’ (or set of guidelines) which aim to distinguish protected speech from non-protected speech. The following sections briefly summarise some of these tests; beginning with the test for a ‘true threat.’

The term ‘true threat’ refers to a threat which is considered to be unconstitutional in nature. This creates a distinction between threats which are legal, and those which are not. Distinguishing unprotected ‘true threats’ from legal threats has been the subject of numerous cases brought to the Supreme Court. Most notably, Watts v. United States [1966], which attempted to differentiate ‘true threats’ from political hyperbole. In this case, the following words were stated at by an attendee at a public rally:

‘If they ever make me carry a rifle the first man I want to get in my sights is L. B. J’[3]

These words were interpreted by the Court of Appeals for the District of Colombia Circuit, as an unlawful death threat directed toward the President of the United States. The reaction of the audience to this statement (applause and laughter), was likened to the reaction of speeches given by Adolf Hitler and ancient Roman Emperors. As such, listener reaction was a highly important consideration in this initial legal decision. This suggests that a different audience reaction, could have led to a different outcome in the initial trial(s). However, it can be assumed that more overtly negative audience reactions, such as jeering or distress, would have also provided grounds for the statement to be non-protected.

Watts, the accused, had appealed against his original conviction. He argued that as his ‘threat’ was phrased conditionally, the stated outcome (shooting the President of the United States) was purely hypothetical. This outcome would only have been realised had

[3]‘L.B.J’ here refers to Lyndon Baines Johnson, the then President of the United States.

Watts actually been forced to enlist in the Army (hence *'If they ever make me carry a rifle...'*). According to Watts, a conviction on these grounds would have violated the First Amendment. The Court of Appeals rejected this reasoning; citing that the 'context' of the utterance would enable it to be inferred as a threat. 'Context' here referred to both the heated nature of the rally, and its location in a public space. The Court also considered that any threats directed towards the President should be non-protected, in the interest of national security. Therefore, Watts' speech was ruled to violate the First Amendment. This decision was later overturned by the United States Supreme Court on appeal.

The Supreme Court in their summation of the appeal case, applied the following reasoning. It was reiterated that a law which prohibits threats made towards the President of the United States was constitutional. However, the 'context' of the alleged threat gave the overall impression of hyperbolic speech, rather than a legitimate threat towards the President of the United States. Here, 'context' referred both to the conditional nature of the statement, and how it was produced at a public rally during heated discourse. It was deemed unlikely that a 'true threat' would be produced under these circumstances. Note, that while the Supreme Court and Court of Appeals considered the same contextual information to reach their verdicts, confusingly, their interpretation of these events were vastly different. What is clearer perhaps, is the weight that these courts gave to audience reaction, and the perceived likelihood of harm occurring (rather than other factors).

The notion of a 'true threat' did not reemerge in the Supreme Court until the case National Association for the Advancement of Colored People v. Claiborne Hardware Co. [1982]. In this case, Charles Evers, a secretary of the NAACP made a series of speeches calling for the boycott of certain businesses which had white owners. Evers was alleged to have said the following:

'If we catch any of you going in any of them racist stores, we're going to break your damn neck.'

This statement was initially judged to be an example of both 'fighting words', and a true threat directed towards members of the audience. Fighting words differ from verbal threats, in that the former is designed to incite or encourage hatred, and/or violence. True threats, on the other hand, are an expression of an *intention* to commit unlawful violence towards a person or group. While these terms are similar, and may overlap with one another, they can be independent language crimes.

However, the Supreme Court later found that Evers' statement was neither an example of 'fighting words,' or a true threat. Unfortunately, it was not made clear why this statement failed to meet the criteria for either language crime. As a result, this case did not provide a clear framework for the classification of true threats that could be applied in future disputes.

This lack of criteria for a true threat was also at the crux of United States v. Dinwiddie [1995], a case involving the harassment of staff working at a family planning clinic. The court considered the following factors when determining whether illegal threats were made, in what are now referred to as the 'Dinwiddie Factors:'

- whether the statement(s) are communicated to the target directly (as opposed to a third party)
- the situational 'context' of the statement
- the manner of the statement
- the number of statements
- whether the accused had a history of criminal behaviour
- the reaction of the target

In her defence, Dinwiddie had cited Watts v. United States as a case which was similar to her own; inasmuch as both cases were examples of speech which were ultimately judged to be protected. However, upon closer examination the two cases appear to be quite different. Watts had produced a single conditional statement which was responded to by laughter. Dinwiddie had made numerous statements which were not always conditionally-phrased. Examples included, *'Patty* [victim's name], *you have not seen violence yet until you see what we do to you,'* and *'Robert* [victim's name], *remember Dr. Gunn* [a physician who was killed a year earlier by an opponent of abortion]. *This could happen to you. He is not in the world anymore. Whoever sheds man's blood, by man his blood shall be shed.'*[4] Rather than a humoured response, the target of her alleged threats had begun to wear a bulletproof vest when venturing into public spaces.

[4]Interestingly, the latter example was referred to as a 'warning' by the United States Court of Appeals.

Although the cases of Watts and Dinwiddie both involved the examination of similar factors (namely audience response, and conditionality) these cases did not reach the same legal decision. In the case of Dinwiddie, it was not deemed necessary for an arbitrary number of the above 'Dinwiddie' factors to be 'proven.' The 'Dinwiddie Factors' functioned more as guidelines to enable the court to consider a wider range of relevant factors than in previous true threat cases.

More recent legal cases in the United States have also examined alleged threatening behaviour which was non-verbal. Virginia v. Black [2003], focused on whether the statute in Virginia banning cross burnings with 'an intent to intimidate a person or group of persons' violated the First Amendment. In the United States, the burning of a cross (or a crucifix) is strongly associated with white supremacist movements, in particular the Ku Klux Klan. Historically, crosses were erected and then set alight in areas visible to members of the black community, people in interracial relationships, or political opponents. Crosses are also burned in secluded areas for rituals attended exclusively by members of the Ku Klux Klan.

The Supreme Court ruled that while the state of Virginia was right to outlaw cross burnings with 'an intent to intimidate a person or group of persons,' it was unconstitutional to outlaw *all* cross burnings. This is in spite of the historical connotations of cross burnings, and its association with expressions of racial hatred. The Supreme Court recognised that cross burnings may serve to communicate alternative messages (for example, a belief in racial segregation or white supremacy). As distasteful as these beliefs may be to most, the right to express such views is still protected by the constitution. This decision acknowledges that while burning a cross *could* be an intentional attempt to threaten a person or group, this is not always the case. There must be evidence to suggest that the action of cross burning was performed with an unlawful intention to communicate a threat to harm. Cross burnings which are found to be motivated for other reasons, would be protected. Exactly *how* this intent could be demonstrated was not made clear.

'True threats' were also the subject of a more recent case involving the first examination in the U.S of non-protected speech communicated via social media (Elonis v. United States [2015]). Anthony Elonis was accused of making numerous threats to harm specific people (including his ex-wife and coworkers). In addition, he also wrote about more general targets ('the police', and an unspecified 'kindergarten class'). The defendant stood accused of writing threatening messages on the social media platform Facebook using a pseudonym. These messages were described by the defendant as being 'rap lyrics,' (performed using a stage name). The 'lyrics' in question frequently referenced violent acts. [5] Elonis argued that the same could be said for lyrics performed by popular musicians, spanning a range of musical genres. Therefore, in addition to not being a true threat, the messages also held some artistic or literary value.

Initially, Elonis was found to be guilty of sending threatening communications. This decision was upheld on appeal. Testimony from his ex-wife suggested that she had feared for her own safety as a result of the social media messages. It was this subjective reasoning (*'Did the target(s) of the message feel in fear for their own saftey?'*), rather than a more objective test (*'Would a reasonable/objective person have felt in fear for their own safety upon receiving the message?'*) that led to the defendent's conviction. Elonis' attempts to reverse his conviction were successful at the Supreme Court. A majority verdict was reached, stating that mere 'negligence' on behalf of the message writer was insufficient mens rea for offences of this nature (albeit proof of mens rea is still required). 'Mens rea' is the name given to the requirement that a suspect had the intention to commit an offence, or being aware that their actions constituted an offence. Simply put, because previous courts had not established *why* the defendant had posted the messages, or whether he understood that his messages could have been interpreted as threatening, his original conviction was unsound. Frustratingly, this case did not clarify whether related traits (such as 'recklessness') were sufficient mens rea. Nor did this case propose an alternative 'test' to establish whether a true threat had been made.

These cases highlight the benefit of a legal 'test' which could consistently identify unlawful threats that are communicated across different media. Difficulty arises however, when establishing whether a defendant held an unlawful intent at the time the alleged offence had taken place. The outcomes of the cases mentioned here have overall, contributed to a more objective, systematic treatment of proving an unlawful intent. In other words, courts have tended to focus on establishing how the targets of the threat(s) have reacted. If the target was left in fear for their safety due to the alleged threat(s), then this communication would be considered unlawful. By extension, the producer of the alleged threat would be

[5] For example, 'If I only knew then what I know now...I would have smothered your ass with a pillow...' and 'I'm not going to rest until your body is a mess, soaked in blood and dying from all the little cuts.'

assumed to hold an unlawful intent.

With the conclusion of Elonis, it is unclear how potentially unlawful threats in future Supreme Court cases should be examined in U.S courts. In short, it is unclear whether establishing the defendant's intention is necessary in such cases.

To summarise this subsection, there have been numerous methods used in courts across the United Kingdom and the United States to determine whether an illegal threat has been made. In general, there appears to be a tendency to assess whether an objective or reasonable person would have felt alarmed or distressed as a result of the alleged threat. This is balanced with gauging the intent of the appellant, and/or assessing whether he/she should have foreseen the consequences of their actions. How someone's intention is established appears to be a mostly implicit process. In other words, 'intention' seems to be judged on the actions (or lack thereof) of the appellant, and the effect of these actions on the recipient.

There are also other considerations that the legal system must make when identifying the exact nature of the offence. For example, specifying the exact target of the threat, or whether factors such as alcohol consumption or mental health are mitigating factors. There are also countless references made in witness testimony to threats being produced in a 'threatening manner' or angrily or aggressively. This thesis will focus on whether there are methods used by speakers/authors to communicate a threatening intent.

The following section provides an overview of the previous linguistic research into threatening language. This past research addresses some of the issues highlighted in these criminal cases. In particular, the problems which can arise when trying to infer an intention to harm using linguistic cues.

2.2 Linguistic research on threatening language

Threatening language involves an interaction between how the threatener communicates his/her intent to cause harm, and whether his/her target correctly interprets this intent. For legal disputes, it is also imperative that this harmful intent can be successfully conveyed to people throughout the justice system. Real life threats (and 'threats' collected under experimental conditions), often do not seem superficially threatening. Nevertheless, these messages are able to cause a listener to feel fearful. This reaction can occur irrespective of whether the sender did not foresee, or intend this outcome. Unlike other language-based offences, threats cannot be communicated through using a specific verb (e.g. 'to threaten'). This is in contrast with other types of utterances (such as warnings or promises), where an act can be explicitly stated using its verb. Also, as explored in this chapter, there is no grammatical structure that is exclusively used to make a threat in English. As such, threats are inherently difficult to identify through words alone. This creates substantial difficulty in identifying threats suitable for research purposes, as well as developing threat detection systems which examine linguistic features.

To further discuss these issues, the following subsections will detail previous descriptions of speech acts. Speech acts are, broadly speaking, utterances which have some performative function, or real-world effect. Verbal threats could be considered as a type of speech act, which are used to express an intent to harm another person or entity. In other words, expressing an intent to bring about harm would be the performative function (or real-world effect) of a verbal threat. Also of interest here, is how threats compare to other, similar kinds of speech act. This is because there are numerous examples of alleged threats which have been phrased using verbs which express other speech acts.

The next subsection will outline Speech Act Theory, as presented by Austin (1969), along with its later developments. This theory attempts to provide a pragmatic framework to establish the intended meaning of speech, as well as its subsequent effect on its audience. Notably, this theory does not necessarily rely on analysing the exact words used (although this remains an area of contention among speech act theorists). As mentioned previously, Speech Act Theory does not focus on threats specifically. However, threats can be considered to be a subcategory of performative speech acts.

2.2.1 Speech Act Theory

Austin (1969), in his theory of speech acts, introduced the theoretical distinction between 'performative' language, and 'constative' language. All utterances can be allocated to either category based on whether they:

- are a statement about the state of the speaker, or the world (constative) or

- perform an action through producing the utterance (performative)

Thanking, apologising and promising are given as examples of performative language. By explicitly vocalising gratitude (*'I thank you for helping me'*), an apology (*'I apologise for not helping you'*), or a promise (*'I promise that I will help you'*), the speaker has, at least in theory, performed the action that was stated upon completing their utterance.

Normally (according to Austin), English-language performatives are typically marked according to the formula below:

Person	**Noun**	**Tense**	**Mood**	**Voice**
1st	Singular	Present	Indicative	Active

The examples mentioned earlier of thanking, apologising and promising are all performatives which fit this model. Particularly in more formal settings, this formula can be changed while still enabling the utterance to function as a performative:

Person	**Noun**	**Tense**	**Mood**	**Voice**
2nd/3rd	Plural	Present	Indicative	Passive

Austin provides the utterances *'Passengers are warned to cross the track by the bridge only,'* and *'You are hereby authorised to pay...'* as exemplars of this alternative performative formula (Austin 1969, p. 57). As such, there is no single grammatical 'formula' for performatives. This creates some difficulty when trying to distinguish performatives from constatives based on words or grammar alone. Superficially, these speech acts can look identical to one another. For example, saying *'I conclude that visiting him is a bad idea'*, could be seen as a 'performance' of sorts or as a statement.

Furthermore, Austin describes constative language as statements that are either true or false. Performative language, on the other hand, can only be felicitous, or infelicitous. This is despite performative language sometimes appearing to resemble statements of fact. For example, in an utterance such as *'I'm warning you not to do that again,'* it is clear that the speaker is warning their audience.

In describing performative language, Austin identified three components:

- Locutionary acts: the production of a meaningful utterance

- Illocutionary acts: the action which is intended by the speaker when they produce an utterance. For example, when a speaker produces a joke or an accusation, his/her intent is to appear humorous or accusatory, respectively
- Perlocutionary acts: the effect that the utterance has on the speaker's audience (whether intentional or otherwise)

It should be emphasised that for Austin, when a speaker produces any element of a performative speech act (i.e. a locutionary, illocutionary, or perlocutionary act) the speaker - by default - realises all of these components simultaneously.

Of these components, Austin concentrates on establishing the illocutionary act (or the illocutionary force) of a performative utterance. If a speaker says *'I thank you for giving me another chance',* this utterance can reasonably be considered as an expression of gratitude. This assumption is reached through the explicit reference to the illocutionary act, 'I thank.' Note that this gratitude only comes from the speaker. It would not be possible to say 'I thank...' to express someone else's gratitude. Conversely, saying *'*I threat...(you with this knife)* assumes that the *recipient* was feeling threatened before the completion of the speech act [6]. Unlike 'I thank...', '*I threat...' does not necessarily describe the speaker's actions (i.e. the act of threatening someone). '*I threat..' also does not describe the speaker feeling threatened.

Alternatively (according to Austin), a speaker can express an illocutionary act without making an explicit reference to it. For example, saying *'I promise you that I will return your money'* would be an explicit promise. Rephrasing this as a statement of fact *'I'll return your money'*, still implies that a promise has been made. In saying this, Austin blurs the distinction between constative and performative utterances. Making statements can also be thought of as 'doing something,' not just describing something.

For these reasons, the true illocutionary act of an utterance cannot be fully established by only examining the locutionary act. The social context (that is to say the people, place and cultural conventions which surround the speech act), can allow for the illocutionary and perlocutionary acts to be better understood. While anyone can utter the words *'I now pronounce you husband and wife',* these words only have a performative effect if spoken in the right social context. The speaker must have the authority to create a legal and/or

[6] The use of an asterisk (*) indicates an ungrammatical, or an unnatural-sounding sentence or phrase.

cultural union between a couple. For some ceremonies, this union must occur in a specific time and place. Conditions such as these must be satisfied if the words uttered are to have any performative effect.

In a similar vein, threats may require certain conditions alongside the speech act itself. At present, there appears to be no research which has addressed what contextual information lay-people may require in order to infer whether a threat has been produced.

2.2.2 Responses to Speech Act Theory

Searle (1989), addresses this constative-performative dichotomy in his later work. Firstly, he argues that performatives can be either true or false, as with constatives. Secondly, he proposes that only utterances which explicitly state a speech act ('explicit performatives' using Austin's terminology), can be categorised as a performative speech act. Unlike Austin, for Searle, indirect speech acts (e.g. *'Would you mind opening the window?')* as opposed to *('I order you to open the window,'*) are not performative utterances.

In this approach, performatives are a type of declaration. A declaration (or declarative) is an illocutionary act that has a real-world effect; based on the locutionary act. This would mean, counter-intuitively, that an utterance such as 'I state that...' would be a performative, and not a statement.

For Searle, this is because performative utterances contain performative verbs which, when uttered, perform the action stated. These verbs are described as having two properties:

- they are defined, in part, as containing the notion of intention
- are able to be used self-referentially (e.g. *'I promise that I'll see you again'*)

When the verb is uttered within a performative sentence, this 'intention' is manifested in the real world. It is not necessary that the speaker wants to (or is able to) fulfil their promise, warning or other action. It is only necessary that the speaker conveys an intention to make a promise, an apology or another speech act. In other words, the speaker intends to communicate a promise or an apology to a recipient but does not necessarily intend to *carry out* that promise or apology. The speaker's intention (expressed by the performative verb) is what enables performative utterances to be either true or false. Using this framework, performative utterances should be able to be interpreted literally.

Utterances which are phrased indirectly are not considered to be illocutionary speech acts.

This framework leads to difficulties in defining spoken threats as a type of illocutionary speech act. Threats frequently feature verbs which do not explicitly reference the act of threatening. Using Searle's framework, an utterance such as *'I'm warning you not to do that again'* would be considered a warning. *'Don't do that again'* or *'If you do that again, it'll go badly for you'* are not examples of indirect warnings or indirect threats. These examples also do not qualify as illocutionary speech acts. This is due to the lack of a self-referential main verb which explicitly states the intended action. However, these examples could feasibly have the *perlocutionary* effect of causing the recipient to fear harm (intentionally or otherwise).

It might be more accurate then to treat the verb 'to threaten' as a type of perlocutionary verb using Searle's terminology. These verbs state a perlocutionary act, such as 'to persuade'. *'I persuade you to join our team'* is an example of a performative utterance featuring a perlocutionary main verb. When the speaker's audience hears this main verb, they, in theory, experience the action stated (in this case, feeling persuaded).

For the verb 'to threaten' to function as a perlocutionary verb within a grammatical performative utterance, the following criteria must be met:

- The verb must be self-referential (i.e. able to be preceded by the first person singular pronoun 'I')
- The verb must be marked for present tense. This is because speech acts are 'actions' carried out upon their completion, not reports on past events, or predictions of future events.

In addition, a perlocutionary verb *can* be followed by another verb within the same clause (in either an infinitive or gerund form). For example, *'I admit to taking the money'* where the main verb 'admit' is followed by the gerund 'taking.' *'I admit nothing'* is also an acceptable sentence, without an infinitive or gerund following the verb (meaning that it is a catenative verb). As with 'to admit,' 'to threaten' is also a catenative verb.

An utterance such as *'I'm threatening to kill you'* would satisfy these conditions, whereas *'I threatened you'* and *'*He threat(en) you'* do not. While *'I'm threatening to kill you'* is a grammatical sentence which adheres to the above criteria, it is still problematic. Such a sentence appears to be *describing* what the speaker has been doing recently, and not performing the action stated in the present moment. Therefore, it is not a performative utterance. Changing this sentence to *'I threaten you'* would avoid being descriptive, while still satisfying the other conditions for a performative sentence. However, making this change results in an issue seen in Chapter 2.2.1. By using 'threaten' as a type of 'performative', there is an assumption that the audience felt threatened upon hearing the utterance. Crucially, 'to threaten' is being treated here as a perlocutionary verb, within a performative utterance. It is paradoxical for a speaker to correctly infer the effect that their utterance has on their audience while producing the utterance itself.

For Searle, this is a limitation of certain performative verbs. This is because few verbs can ever truly bring about observable, real-world change. Arguably, this is also a limitation of using perlocutionary verbs both self-referentially, and in the present tense. Similar issues arise in comparable sentences, such as *'*I convince him',* or *'I amuse the audience'.* The main verb is said to convey an intention to perform an action. As a result, the utterance will always be a speech act which matches the main verb. Real-life alleged threats (criminal or not) seldom include the word 'threat'. Requiring that they do so, would heavily restrict the number of possible threatening speech acts. Conversely, an innumerable amount of real-life threats would be classified as 'inexplicit.' Therefore, such threats would not be classed as performatives using this framework.

2.2.3 Summary of Speech Act Theory

In summary, the approaches used by Austin and Searle to classify performative speech acts appear to be diametrically opposed. For Austin, every utterance (irrespective of its wording) has the potential to be a performative speech act. This is so long as the utterance conforms to societal conventions. Threats can be performative because an utterance can be produced with the intention of being threatening (an illocutionary act), and/or be perceived as threatening (a perlocutionary act). This means that, theoretically, any utterance has the potential to be a threatening speech act.

On the other hand, threats appear to be largely prohibited under Searle's performative speech act framework. This is because performatives must explicitly state the act in question. Real-life threats are rarely phrased using the word 'threat' (or even related words such as 'intimidate' or 'frighten'). The fact that alleged threats are disputed in both legal and non-legal settings suggests that their wording alone is not enough to determine whether a threat was made. However, it is feasible that a proportion of these alleged threats were intended to be, and/or were perceived as threatening.

2.2.4 Threats, warnings and promises

At present, there is no single, widely-accepted definition of a threat used across linguistic research. While the term 'threat' is sometimes used as a synonym for 'promise' (de Becker, 1997; Mullen, Pathé and Purcell, 2000) or 'warning' (Storey, 1995), there is a small body of linguistic literature which discusses the apparent differences between these speech acts. This section will provide an overview of these pragmatic descriptions, based on this previous literature (Fraser, 1998; Gales, 2012).

Notably, unlike threats, warnings and promises are said to be made for the ultimate benefit of the recipient of the message. This is despite some warnings and promises appearing to be detrimental to the recipient. For example, in the sentence *'if you don't do your homework, you'll be grounded',* at first glance may appear to be threatening. However, this sentence also serves to prevent the recipient from a future negative sanction (i.e. being grounded). This sentence (if directed to the intended target), can be seen to ultimately benefit the recipient, as they are made aware of this possible negative sanction prior to its implementation. In addition, this statement also acts as an incentive. The act of completing the homework would likely result in an improved state of the world for the recipient. Therefore, this phrase could be described as a warning.

Likewise, issuing a promise with a negative sanction, shifts the function of the speech act away from a prototypical 'promise', towards language which appears to be closer related to a threat. For example, *'I promise to find you after school and stab you'.* While this utterance is worded as a 'promise', it expresses an overtly harmful act towards the recipient. This utterance would likely be better described as a threat, despite the use of the performative verb 'promise'.

Related to this, when promises are not acted upon, this reflects negatively on the speaker. For example, failing to return a loaned book as promised before a predetermined date, would reflect badly on the speaker. Warnings arguably do not create the same response. This is because the ‘obligation’ to act on the warning is left to the recipient, and not the speaker. The negative sanction in a warning may also not be in the control of the speaker. For example, warning a friend not to go outside during dangerous weather conditions. In this scenario, the speaker is not able control, or influence the occurrence of the sanction (dangerous weather conditions).

Interestingly, none of these three speech acts are described in relevant linguistic research as in some way beneficial to the speaker. It could be argued that making (and successfully completing) a promise benefits the speaker by improving their social standing. Some threats may benefit the threatener because the threat acts as a form of coercion. In other words, the target is ‘forced’ into some action that is desirable to the threatener (at the target’s expense).

The differences and similarities between threats, warnings and promises are summarised in Table 1:

Table 1: Overview of threats, warnings and promises (adapted from Fraser, 1998)

Feature	**Threats**	**Warnings**	**Promises**
For speaker’s benefit	No	No	No
To addressee’s benefit	No	Yes	Yes
To the speaker’s detriment	No	No	No
To the addressee’s detriment	Yes	No	No
Speaker controls outcome	Yes	?	Yes
Addressee controls outcome	?	?	?
Speaker committed to act	No	No	Yes

2.2.5 Linguistic approaches to defining threatening language

The existing linguistic research on threatening language has largely focused on defining threats based on how they are inferred by others (Fraser, 1998; Gales, 2012; Gingiss, 1986; Napier and Mardigian, 2003; Shuy, 1993; Smith, 2008; Story, 1995; Yamanaka, 1995). In particular, how the words and the illocutionary force (i.e. speaker intention) used to express them are able to communicate a threat (Fraser, 1998). In addition, this linguistic literature on threats often distinguishes between different types of threat. These types of threat are often referred to as *direct*, *indirect* and *conditional* threats (Gales, 2010), although there are also alternative names given to these types of threat, such as 'genuine and explicit' (Shuy, 1993) and 'ambiguous' (Solan and Tiersma, 2005).

Across the relevant linguistic literature, there is a tendency to refer to 'direct threats' as threats which include an explicit reference to harm in their wording. As such, a reader (or listener) would be able recognise a clear reference to a violent or other unfavourable action. In contrast, 'indirect threats' are vaguely-worded, and the target is required to infer whether or not he/she is at risk of harm.

Understandably, indirect threats are problematic from a threat assessment perspective. This is because these threats are likely to be difficult to infer based on the words used alone. Uttering an indirect threat would thus provide a threatener with plausible deniability. In other words, a threatener can simply deny that he/she communicated an intention to harm, because there are a lack of words present in his/her 'threat' which could reasonably be inferred as such. Therefore, a target could sincerely feel that they are being threatened, but be unable to ever 'prove' that the 'threatener' meant to communicate any intention to harm.

The pragmatic approach seen in the previous literature alluded to in this section is also potentially problematic when applied to real world threat assessment. If threats are defined based on how they are inferred, this assumes the following:

- that a target is able to consistently and accurately infer a threatening intent.
- that different people could be exposed to the same language, in the same situational context, and would be able to infer the same threatening intent.
- that assessing the threatener's intention is not possible, or is a less useful tool for determining whether a threat has been made.

However, these assumptions are problematic from both a research and legal perspective. Currently, there is very little research into the perception of threatening language. As will be discussed later in this thesis (§2.2.7), the available literature on this topic has so far has only shown that listeners might use vocal cues to assess indirect threats and that listeners associate threatening speech with lowered pitch. There is still a large gap in our understanding regarding how lay-people (as well as threat assessment practitioners) infer a threatening intent based on linguistic features. As such, defining threats based on inference could be considered problematic when it is not yet understood *how* people infer threatening language.

Also, as shown previously in the current research, there are legal disputes which centre on whether threats which contain clear references to violence can even be considered threatening. While the bar set for establishing an illegal threat is high, it is also evident that proving that a witness had sufficiently 'felt threatened' is a difficult and inconsistent process. In addition, defining an illegal threat is heavily dependent on the location, time and nature of the threat. It is possible that legal decisions on threats can centre on whether the threatener intended to communicate an intent to harm, rather than assessing the perception of the 'threat'. As such, linguistic literature which focuses on conceptualising forensically-relevant threats should be mindful of the legal considerations relating to threats.

It is also possible that this body of literature has limited applications for real world threat assessment, which often examines the threatener's intent based on his/her actions. As noted by Fraser (1998), it might be unrealistic to aim for the accurate identification of threats based on linguistic features. That said, there has been an emerging body of research which explores the viability of linguistic threat detection. This area will be discussed in the next subsection.

2.2.6 Linguistic approaches to threat detection

This research is concerned with better understanding how lay-people define or conceptualise 'threatening speech' and, moreover, with how their ideas relate to the previous linguistic research about threatening language. In addition to the more pragmatic approaches outlined in the previous section, there have also been other linguistic approaches to identifying threats.

In recent years, there has been a range of computational linguistic approaches to the issue of automatic threat detection. Research of this kind is centred on developing tools which can flag up written text online which is deemed to be threatening. For many social media platforms, making threats is a prohibited behaviour (Help.twitter.com, n.d.; Facebook.com, 2018). However, owing to the sheer quantity of written communications on these sites, flagging up and responding to threatening language manually presents considerable difficulties for these platforms.

This has generated a small but growing interest in developing ways to 'accurately' identify online comments that are threatening. On the basis of publicly-available research in this area, this process seems to have focused on developing machine-learning algorithms that can automatically detect threatening comments. These algorithms are developed by collecting large quantities of written data and manually tagging text which is threatening. These data are also tagged for categories of interest to the researcher such as lexical features, parts-of-speech and lemmas. The software 'learns', through exposure, how to identify strings of texts which exhibit sufficient similarities with its database.

Research in this area has led to the creation of threatening language corpora such as the YouTube Threat Corpus (Hammer, 2014; Wester et al., 2016). This corpus consists of 9,845 comments taken from eight YouTube videos. These videos focus on topics such as immigration and religion, where it might be expected for the topics of discussion to be subjected to hateful or threatening language. Comments which were deemed to be either threatening or not threatening were tagged manually by the researcher. This tagging process was validated by calibration with other researchers. In this study, these data were tagged for bigrams of subject (e.g. 'I' and 'we') + 'aggressive' words (e.g. 'bomb', 'kill'). This is because it was expected that threats would be characterised by a subject that is performing (or is said to perform) a violent action (see also Cohen et al., 2014). Hammer (2016), linked the combination of subject + aggressive words, along with the use of the word 'Breivik' [7], with the classification of threatening texts.

There is also research on cyberbullying which has sought to identify linguistic features which can enable the automatic detection of this behaviours. Dinakar et al. (2011), created their own corpus of 50,000 YouTube comments taken from videos which had controversial content surrounding the categories 'sexuality', 'intelligence' and 'race and culture'. The comments taken from these videos were also grouped into these three categories. These comments were then manually tagged for comments which attacked members of

[7] Referring to Anders Behring Breivik who was convicted of the murders of 77 people in Norway in 2011.

these groups. For example, for the group concerned with cyberbullying based on 'sexuality', comments which were deemed to be attacking women or sexual minority groups were tagged. As with Hammer (2014), there was additional inter-judge calibration for this tagging process.

These data were tagged for unigrams, negative words, profane words, and frequently occurring parts of speech bigram tags. Negative words and profanity were deemed to be likely to be associated with cyberbullying. It was also noted that 'blatant' bullying involving profanity and abusive terms were both easier to model than vague language such as euphemism. This is likely due to the lack of wider contextual or linguistic information available which could contribute to the modelling of less overt cyberbullying.

These studies have both involved the use of human judges to assess whether a text is threatening or an example of cyberbullying. However, as seen in this research, threatening language is defined in a variety of ways across disciplines. As such, selecting appropriate or representative threat data from huge corpora such as online forum comments is an inherently difficult task.

This data collection method presents a number of possible issues. Firstly, there is potential for human judges to use different criteria when assessing whether a text is threatening. This could lead to disagreement or inconsistencies when determining whether an utterance is a threat. These inter-judge differences could be based on a number of factors, such as exposure to abuse or threatening situations (McCrory et al., 2011). There is still a great deal of uncertainty as to what external factors contribute towards threat perception. However, the use of inter-judge calibration discussed earlier could lessen the risk of misinterpreted or false-positive data being included in threat language corpora.In other words, there should be some agreed upon method to distinguish between threatening and non-threatening texts. Detailing this method would also enable the same process to be replicated in related future research using different data. This methodology might be somewhat straightforward when tagging direct threats. For these threats, judges might focus on looking for words which are associated with harmful actions. However, identifying indirect (or vaguely-worded) threats would be substantially more challenging. Without explicit references to violence or harmful actions, identifying indirect threats seems to be heavily subjective. The current study is motivated in part, by identifying ways in which authors might differentiate direct from indirect threats.

Another issue in using human judges to construct threatening language databases is that the data is selected based on a reader's perception. As discussed earlier in this research, threats are often thought of as expressions of intent by a speaker/author. As such, these data do not represent whether an author intended to communicate a threat - only that he/she appeared to do so. However, for online forums, whether an author intended to make a threat might be beside the point. Threat detection algorithms might be designed to identify comments which cause other users to feel fearful. Therefore, these threat detection systems might be expected to respond much like a 'typical' human reader, albeit one who is able to assess far larger quantities of language data automatically.

The issue of author intent also overlaps with the lack of available context for the alleged 'threat'. For example, a seemingly obvious threat might be a quotation which does not reflect the feelings of the author or is intended as a joke or within-group reference. Tagging threats without understanding the context in which it was made could result in collecting false-positive threat data that represents neither author intent or reader perception.

This research proposes that more studies are needed which more thoroughly examine the criteria used by lay-people and threat assessment experts to assess threatening language data. This information would better validate the current approaches used to collect threatening language data and the construction of threat detection algorithms. In order to contribute to this goal, the current research provides a quantitative and qualitative analysis of the linguistic content of threatening data collected under experimental conditions. By collecting experimental (or simulated) data, there can be some assurance that the authors intended to appear threatening. This allows for linguistic features to be associated with the expression of a threatening intent. These can later be incorporated into the design of future threat detection algorithms.

2.2.7 Linguistic features of threatening language

Currently, there is some relevant research available which could inform the design of threat detection tools. These studies have examined the linguistic features which occur in a wide variety of related texts. Of relevance to the current study, are studies which relate to the analysis of written texts authored by members of extremist groups. While such analyses do not focus on threatening language per se, they can provide a framework for how threatening language could be collected and analysed. These studies also associate specific linguistic features with strategies which are said to be associated with intimidating readers and the prediction of the actualisation of violence.

Pennebaker (2011) examined the linguistic features of 296 text files which were written by four groups associated with Islamic extremism. These groups included the al Qa'ida central group and Arabian Peninsula group, Hizb ut-Tahir, and the Movement for Islamic Reform in Arabia. The former two groups listed have previously committed acts of violence. The texts collected from these groups were translated from Arabic to English using unknown, Open-Source translators, and analysed using the text analysis software LIWC (Linguistic Inquiry and Word Count). This software was used to analyse a wide range of linguistic features. These included examining the quantities of function words (e.g. pronouns and modal auxiliary verbs) present in the texts written by authors affiliated with these extremist groups. Other outputs exclusive to LIWC software were also examined, such as the use of emotion words.

From these texts, it was found that the two violent groups appeared to write their texts differently to each other. In particular, the authors affiliated with these groups appeared to write texts which differed in language use in the months before a terrorist attack. As such, the author recommends that models which calculate the risk of an attack based on written texts should be group-specific. In addition, for both groups the LIWC summary variable 'Analytical' was associated with an impending attack. This variable represents the degree of 'complex thinking' expressed by the text(s). It was found that texts written before an impending attack had lowered analytical scores. This indicated that texts written in the months prior to an attack were associated with lower levels of complex thinking or thoughtfulness.

Other studies have used LIWC software to quantify the linguistic features of other forensically-relevant communications. In Chandler (2017), the focus of the study was to document the linguistic features used in 392 threatening communications authored by members of animal rights extremist groups. These threatening language data were procured online. Similarly to the study discussed above, LIWC software was also used to analyse these threatening texts. It was found that these texts had high scores for the LIWC summary variable 'Clout' and low scores for 'Tone'. These variables and their relation to the current study will be discussed later in the results of the linguistic analysis of this study (§4.2).

In addition, it was also noted that these threatening texts had low scores for negative words. This indicated that the threatening texts of animal rights extremists do not feature many words which can be associated with 'negative' emotions such as anger or sadness. There was also a tendency for these data to focus on events expressed in the present tense, rather than past or future events, and for these communications to feature very few instances of profane words. These findings suggest that authors of threatening texts might focus more on describing present events (perhaps the ongoing situation which has 'motivated' the threat) than describing past or future events.

As much of the data collected and analysed for the current research is based on a case involving threats delivered in response to (fictional) animal abuse, the findings of Chandler (2017) are of particular interest. Especially, because Chandler's findings allow for the comparison of simulated threatening texts concerning animal abuse against data sourced from real-life threatening communications.

Due to its usage in the analysis of forensically-relevant written texts (especially those relating to threats or terrorist communications), LIWC software will be used to measure the linguistic data collected for this study. This software and the variables it measures will be described in greater detail in the methodology chapter of this research.

Among the most detailed linguistic examinations of threatening language can be found in Smith (2006). This research focused on examining the features of threatening communications which can be associated with the actualisation of violence. These texts had been provided by the Federal Bureau of Investigations National Center for the Analysis of Violent Crime. The analysis was mostly performed by the computer program 'Psychiatric Content Analysis and Diagnosis' (Gottschalk & Bechtel, 2001).

This research found that threats which resulted in the actualisation of harm were associated with the following linguistic factors:

- using words indicating prejudices concerning religion

- repeatedly mentioning love, marriage, or romance
- using a polite tone in the threatening communication
- indicating the target/victim, either explicitly or implicitly, in the threatening communication
- specifying weapons threateners planned to use
- mentioning a reason or motive for making the threat
- indicating threateners were thinking about being with the target forever

However, no association was found for the following linguistic factors and the actualisation of harm:

- specifying the type of harm (direct, conditional, implied)
- indicating who will carry out the threat
- indicating when the threat will be carried out
- specifying a violent action
- using 'hypothetical structures' (i.e. modal auxiliaries such as 'should' and 'ought to')
- indicating his/her ability to carry out the action specified or implied in the threat
- indicating the threateners' plan to communicate with the target again in the future
- use of the passive voice

Some of these results contradict the previous findings of related research. In particular, there is research which has linked the specification of time and place in threats with the actualisation of violence towards celebrity targets (Dietz, Matthews, Van Duyne, et al., 1991).

It has also been reported that the Behavioral Analysis Unit (BAU) of the FBI has considered the use of linguistic data when assessing the veracity of written threatening communications (Simons and Tunkel, 2014). Of note to the current research are the following features:

- the use of pronouns
- the use of passive and active voice (where the passive voice could indicate a lack of commitment to an action)
- changes in language which indicates ownership (e.g. 'my gun', 'the gun', 'a gun')
- evidence that the text has been edited
- words such as 'maybe,' 'basically' (the use of which is also associated with a lack of commitment)

- quantity of information provided in text (threats which provide more information are typically treated as being more likely to be realised)

More recent research on threatening language has focused on looking at the context of a threat to establish its illocutionary and perlocutionary force (Kaplan, 2016), and examining the linguistic content of threatening texts (Gales, 2010, 2011; Nini, 2017). Gales (2010, 2011) and Nini (2017) both found that authentic written threats could be characterised by expressions of modality, particularly though the use of modal auxiliaries which are associated with certainty (e.g. 'will').

These features will be considered in relation to the written data collected for this research. As will be discussed in more detail later in this research (§3.3), these written data were collected under experimental conditions. Therefore, this research will not be able to assess the veracity (or likelihood of harmful actions) on the basis of these texts and recordings. However, what can be assessed is whether the authors of these texts consistently use specific linguistic features when writing threats, and in addition, whether these features are used differently by authors writing indirect and direct threats.

In summary, there are a growing number of studies which are documenting the linguistic features of threatening texts. It is also possible that there are other such algorithms in development outside of the public domain. In addition, there are threat detection algorithms currently under development. These studies have identified a number of linguistic features which could be used to predict the actualisation of violence. These include low levels of analytical thinking, high levels of 'clout', using words associating with religion or love, using a polite tone, indicating the target of the threat, specifying the time and place of the action, specifying the motive of the threat. These features will be considered when examining the data collected in this research.

In this research, there will be no comparison between threatening and non-threatening texts. Instead, this research examines whether there are linguistic differences between simulated indirect and direct threats, and consider whether these simulated data can be likened to what is currently known about real-life threats. This includes the linguistic features which have been linked with actual communications delivered from members of animal rights extremist groups. These include: low levels of 'emotional tone', focus on present tense and low frequencies of profane words.

The next section describes research which has described the phonetic features associated with threatening speech. Unlike the previously described linguistic studies, there does not appear to be any publicly-known research into the detection of threats based on vocal features.

2.3 Phonetic features of threatening speech

As noted in §2 and §2.1, the idea of a 'threatening tone of voice' has been expressed in legal cases and threat assessment resources. There has, however, been very little research which has explored the phonetic basis for a threatening tone of voice. This lack of attention is significant because at present, there is a possibility that the 'accurate' inference of threat based on vocal properties could be overestimated. This has implications for the advice given to victims of abuse which refers to the importance of 'tone' as a predictor or precursor to violence (Breakwell, 1997).

At present, there are only a few predictions for the vocal production of threatening utterances. The following sections will outline theories drawing from a range of academic disciplines, and describe their significance to the current study. Observations made about both human and non-human vocalisation will be presented. This is because there is an established and substantial body of research which examines the supposed 'threatening' vocal cues produced by different species. The findings of this area of research are highly relevant to theories of human communication, and especially to theories which derive from a biological, or evolutionary perspective. The possible connection between 'being threatening' and body size, in particular, has been noted about both human and non-human threat vocalisations.

2.3.1 Theories of non-human threatening vocalisations

Non-human creatures have long been observed to make specific vocalisations in relation to specific actions, including those which have been associated with threat displays (Morton, 1977). These threat displays signal that the creature is a threat to members of their own species, and/or towards creatures of another species. It has also been observed that animals will produce specific threatening vocalisations and noises in response to different creatures (Berthet at al., 2018).

There is research which has suggested that birds and animals use similar means to communicate that they are a threat. Morton (1977) noted that creatures which threaten emit a low-pitch vocalisation, while high-pitched sounds are associated with submissiveness. These differences in the inference of pitch are said to be related to the perception of body size. Physically larger creatures would likely be able to dominate in conflicts. In addition, physically larger creatures are often associated with having large vocal tracts (although this is not necessarily the case across all species of animal; Fitch, 2016).

Therefore, it is possible that by modifying the dimensions of the vocal tract, a creature could create an impression that it is larger than its actual size. This would allow the creature to be perceived as more dominant or threatening than its body size might otherwise suggest. In turn, the creature would be perceived as a greater risk to others. Other creatures have also been shown to modify their behaviour in response to threat displays (MacLean and Bonter, 2013).

However, non-human creatures have been observed making a large variety of sounds when appearing dominant or threatening. For example, some species of monkey have been found to hiss when appearing threatening. Conversely, the threat displays of dolphins involve 'jaw clapping' and leaving the mouth agape (Overstrom, 1983). Macaque monkeys are said to produce 'noisy, aperiodic' sounds (Fitch, 2016) as a means of communicating a threat, and elephants produce increases in f0 and formant locations (Berthet et al., 2018).

These studies are often based on the observation of animals in their natural environment. The threatening behaviour is inferred largely from the reaction of other creatures and the outcome of the 'threat'. For example, whether the animal was able to protect its resources or offspring successfully. This means that the intent of the creature is inferred from the consequence of the supposed threatening vocalisation. This is perhaps analogous to research which infers threatening intent based on the harmful actions later carried out by human threateners (Meloy et al., 2012).

The research on cross-species threat vocalisations, indicate that there is no apparent tendency for animals (across species) to produce threat vocalisations using a particular set of acoustic cues. In addition, creatures can signal threat by producing noises using different parts of their bodies (i.e. not necessarily through an oral tract). As such, these threatening sounds can have markedly different acoustic qualities.

With this evidence in mind, it is not inconceivable that animals are capable of recognising human-produced threats (Hemsworth, 1985). This could possibly be inferred through vocal cues, and/or visual cues such as changes to body language or facial expression. It could also be suggested that, like animals, humans have also developed strategies for sounding threatening. The following section describes the prominent theories regarding human threat vocalisations and speech.

2.3.2 Theories of human threatening vocalisations

One of the most prominent theories surrounding human threat production was popularised by Ohala (1984). In his body of research on the subject, it is proposed that sounds have a relationship to some specific meaning. This concept is widely referred to in the relevant literature as sound symbolism (Ohala, Hinton & Nichols, 1997). One example of meaning inferred from an acoustic channel is body size. In other words, information about body size can supposedly be inferred via an acoustic channel. For Ohala, size sound symbolism is evident in both human and non-human vocalisations.

One such example is the relationship between body size and the vocal expression of body size (known as the 'frequency code'). This was explained by the sexual dimorphism of the human vocal tract. There are observable differences in the size and dimension of male and female vocal tracts. These differences are explained through males 'requiring' the means to effectively communicate that they are a threat. It is an evolutionary advantage for males to be able to convey a larger physical build (irrespective of their actual build), in order to attract a mate and to better protect themselves and their resources. Ohala proposes that to portray a larger physical build effectively, male vocal tracts undergo a number of changes around puberty. These changes are designed to maximise the level of threat or dominance that a male can convey vocally.

Ohala notes the following about the male vocal apparatus:

- male vocal folds are larger in size, denser and longer than female vocal folds.

- males have a longer vocal tract than females. In addition, the larynx can be found lower in the neck in males, than in females.

These changes, combined, result in males producing speech with lower fundamental frequencies and lower resonances than females. This results in males speaking with a lowered fundamental frequency; doing so is said to mimic the vocal output of a larger person, even if the speaker is not comparatively large himself. Further evidence to support this theory can also be seen in the interpretation of smiling, during which, the lips are held in a spread position with the teeth exposed. This facial gesture shortens the length of the vocal tract during speech production. By shortening the vocal tract this in effect mimics the body of a smaller person, resulting in the speaker being perceived as physically smaller than that might be in actuality. In turn, being perceived as smaller is also associated with traits such as submission or appeasement. It should be said that the distinction between concepts such as 'body size' and 'physical dominance' used in this theory is not altogether clear. For example, both dangerous, and small, relatively harmless animals, could produce high-pitched noises when trying to threaten.

In support of this theoretical approach, there has also been empirical research which has linked fundamental frequencies produced by men to their level of testosterone (Dabbs Mallinger, 1999; Evans et al., 2008). Put simply, higher levels of testosterone in men is linked with the production of lower fundamental frequencies. This finding also relates to the expression of aggression and anger, in that increases in testosterone are known to coincide with these feelings. Also of note, is how such biological changes do not result in comparable vocal changes for female speakers.

Xu, Kelly and Smillie (2013) later expanded on these approaches with their own 'bio-informational dimensions theory.' Put simply, this theory centres on how the expression of emotions is shaped by human evolution. It was proposed that the behaviour of other people could be manipulated through the vocal expression of various qualities:

- Size projection: equivalent to the frequency code (Ohala, 1984) or 'size code' (Chuenwattanapranithi et al., 2006)
- Dynamicity: or how 'vigorous' the vocalisation sounds. Sounds which are more vigorous are characterised by a large movement range of active articulators and high velocity f0 and formant movements. Less vigorous sounds possess the opposite acoustic characteristics. More vigorous sounds are said to be associated with 'strength',

while less vigorous sounds are said to be associated with ‘weakness’.

- Audibility: or how far a sound can be transmitted by the speaker. By producing louder sounds, a speaker has the advantage of being able to communicate to others from a distance.
- Association: or the use of sounds that are associated with non-emotional biological functions. For example, a speaker who can produce a sound associated with feeling ill can convey to others that a substance is not safe to eat as such warning them about danger.

For threatening speech, it might be expected that people would opt for voice features which are associated with appearing physically large, being strong, and being able to be heard by others from a distance.

There has also been literature from the biological anthropology and evolutionary psychology domains which relate to the assessment of physical strength based on vocal cues. Sells et al. (2010), found that listeners were able to accurately assess the upper body strength of male speakers from vocal cues alone. This was in spite of the cultural and/or biological differences between the speakers and listeners. In addition, while both male and female listeners were able to accurately assess upper body strength, they could only do so on the basis of male-produced speech. When presented with female voices, listeners were not able to accurately assess the level of upper body strength.

These differences in the perception of physical strength (and by extension perhaps, threat) of male and female speech might derive from biological differences. Females do not undergo comparable physical changes to their vocal tracts around puberty as is the case for males. However, it is reasonable to assume that females would require some strategy to deal with dangerous conflicts. For this study, comparisons between male and female vocalisations of threats will be made. This is in order to explore whether there is any evidence to suggest that females have adopted their own bespoke strategies for threatening speech production. More recent research has also linked low mean f0 with the perception of dominance (Tsantani, Belin and Mcaleer, 2016). Notably, this link was only observed for male speech.

2.3.3 Research on human threatening vocalisations

There has been a small body of research based at The University of York concerned with the perception of threatening speech. Watt, Kelly and Llamas (2013) examined the subjective perception of threat from spoken, neutrally-worded utterances. Here 'neutrally-worded' refers to utterances which do not contain any explicit references to commit harm. The sentence *'I know where you live'* was read aloud multiple times by male speakers of English. For each reading of this utterance, the speakers first read aloud the utterance 'cold' (or without any instruction), then again in a 'threatening way'. The participants were given no instruction on how to sound 'threatening'.

After this, these speakers read aloud a short script where this target sentence was embedded. In the first script, the wording was innocuous and was designed to appear non-threatening. In the second script, the wording was more overtly threatening. Speakers of several foreign languages (Arabic, Swedish, Norwegian and Hebrew) were also recorded performing the same tasks that had been translated into the respective languages. This translation was produced by other native speakers of these languages. These English-language and foreign-language target sentences were then extracted from the recordings. These four utterances acted as stimuli in a subsequent perception experiment.

A panel of 30 native-English speakers participated in this perception experiment. These listeners had been screened to exclude participants with knowledge of any of the foreign-languages recorded for the experiment. After listening to each recording, the listeners were asked to rate the level of threat and 'intent to harm' they felt were conveyed in these utterances (using a 0-100 scale). Overall, these listeners were able to distinguish the utterances which had been produced with a 'threatening tone of voice' from those read 'cold'. However, this pattern was notably more consistent when the listeners rated English-language utterances.

This suggested two possibilities relating to threatening language. Firstly, that listeners appear to interpret threats more consistently to one another when the linguistic content was understandable to them. In other words, inferring a threatening intent seemed to rely to a large extent on whether the listener could understand the meaning of the words which were uttered. Secondly, there was some evidence to suggest that the presence of a 'threatening tone of voice' could have a minor effect on the perception of spoken threats, and, in particular, threats which were uttered in a foreign language.

This study does have a number of shortcomings which should be noted. The perception stimuli only consists of recordings of brief, read speech delivered by speakers who were not screened for their level of acting experience. It is entirely possible that the speakers were not reading aloud these threats in a way which resembled how they (or other people) might deliver threats in real-life situations. As such, the listeners could be basing their impressions of 'threat' on voice recordings which did not feature phonetic cues which listeners use to infer 'threat'.

This research led to further inter-disciplinary investigations into the perception of dominance based on altering particular phonetic features as well as visual cues. As described earlier in this subsection, threats have been associated in related research with the expression and perception of dominance.

Mileva et al. (2017), examined the effect of altering the mean pitch of a series of recordings consisting of the sentence 'I wouldn't do that if I were you'. These utterances were produced by twenty speakers (10 male, 10 female). These audio recordings were later edited to create low and high pitch levels. For the male speakers, the mean f0 of each recording was edited to 90Hz (low) and 140Hz (high). For the female speakers, the mean f0 of each recording was edited to 170Hz (low) and 250Hz (high).

A panel of 36 listeners rated each recording for dominance (on a 1-9 scale). It was shown that dominance had very high inter-rater reliability. Voices with lower-pitched were perceived as more dominant than higher-pitched voices. It was also found that male and female voices were consistent to one another for the effect of pitch alteration on dominance perception. As such, lowering or increasing the mean f0 of male voices had a similar effect on dominance perception compared to female voices. It is not yet known whether male and female voices which are naturally below the mean f0 (compared to the average for their respective sexes) would also be perceived as more dominant. If so, this could raise the possibility that some voices have properties that are perceived by others as inherently more dominant, and as such, perceived as potentially more threatening.

The findings of these studies can be summarised as follows:

- Listeners seem to be able to distinguish between voices where the speaker portrayed a 'threatening tone of voice', from voices where the utterance was read 'cold'. However, this task appeared to be easier when the language spoken was also understood by the listener.
- Listeners associate 'dominance' with low-pitched voices.
- Listeners do not appear to perceive dominance differently for male and female voices when their pitch has been altered.

In summary, there has been a small body of research which has specifically focused on the vocal production of threats. A number of phonetic features have been associated with the production and perception of threatening speech. As such it can be predicted that (human-produced) threats could be characterised by the following phonetic properties:

- Lowered mean f0
- Increased intensity (the acoustic correlate to 'loudness')
- Spread lips and open-jaw articulatory settings

The research discussed in this chapter also provides a basis for a number of the research questions which will be addressed in the current research. **Research question 1** intends to examine the phonetic properties of threatening speech. As listeners have previously been found to infer threat from neutrally-worded utterances (including those spoken in languages which are unfamiliar to the listener) there is some evidence to suggest that listeners assess threats through the speech signal. The decrease in mean f0 has also been seen to lead to impressions of dominance from spoken, neutrally-worded utterances.

Research question 1.e of the current study is concerned with exploring the similarity or difference between male and female speakers during threat production. There have been predictions concerning male threat production. These predictions are based on the physiological changes specific to males who have undergone puberty. This prediction has not yet been tested empirically. In addition, it is unknown whether female speakers would vocalise threats similarly to males despite their anatomical and hormonal differences.

What is far more extensively documented in phonetic research are the phonetic properties of emotional speech. The following section focuses on the findings of research which has examined the phonetic production of anger.

2.3.4 Phonetic correlates of anger production

In lay-person discourse, and in related literature, threats are described in relation to anger or aggression, for example as noted in Chapter 2.0.3, threat assessment protocols often note whether or not the speaker 'sounded angry'. This is despite the fact that 'appearing angry' is arguably not a necessary precondition for making a threat. Instead, what may be important is whether a speaker can successfully *convey* anger (or dominance) to the desired target(s). There is currently a lack of research available to indicate whether producing threats with phonetic features that are associated with anger influences threat perception. For some threat assessment practitioners, speaker motivation appears to be linked with speaker intention. For a speaker to make a 'genuine' threat, it appears to be seen as advantageous to convey that they are motivated by feelings of anger (actual or feigned).

In addition, there have been no studies which have examined the phonetic features of threatening speech. As such, the current study uses data which was collected using a variety of methods that have already been established in emotional speech research. The issues and considerations relating to angry speech discussed in this section are also highly informative to the collection of threatening language.

There have been many phonetic studies into angry speech. These studies can be separated into two broad themes: anger perception and anger production. For anger perception, there have been numerous studies which explored how human subjects distinguish between different emotional states. In research by Scherer and Oshinsky (1977) human listeners based at the University of Pennsylvania listened to synthesised tone sequences; the listeners were then asked to judge the tones along a series of labels (e.g. 'pleasantness-unpleasantness', 'potency-weakness'), along with judging whether or not various emotional states were expressed by these tones. It was found that the listeners associated anger with a 'high pitch level', 'fast tempo' and an 'upward pitch contour'.

There has also been research which has examined the perception of emotional states across languages. Thompson and Balkwill (2006) had English-speaking judges listen to foreign language speech produced by several speakers to represent different emotions. The judges had no formal training in the languages that they were exposed to. It was observed that judges were able to recognise emotional states more easily when the speech was produced in their language (i.e. English). This result appears to be analogous to the findings of Watt, Kelly and Llamas (2013) where threats were more easily identified when the target language was understood by the judge. This again reinforces the idea that recognition of behavioural or emotional states appears to be somewhat dependent on being able to understand the meaning of the words used.

There have also been attempts to develop machine learning tools which can accurately classify the emotions expressed during speech. Shami and Verhelst (2007) divided their emotional speech recordings into a training and a test set (as part of an emotional speech machine learning algorithm). When programmed to examine features relating to 'pitch', intensity and speech tempo, 20% of the 'angry' recordings were classified by the algorithm as 'happy' based on these phonetic features. In contrast, 35.6% of the 'happy' speech recordings were classified as 'angry'. These findings suggested that happiness and anger have similar phonetic properties which could lead them to be incorrectly classified.

There has also been a wide body of phonetic research relating to how speakers produce different emotions (Williams and Stevens, 1972; Banse and Scherer, 1996; Kienast and Sendlemeier, 2000; Yildirim et al., 2004). In these studies the following phonetic features were found to be associated with angry speech:

- increase in fundamental frequency
- increase in intensity
- increase in speech tempo

There has also been research which found that angry speech is produced with a lowered mean fundamental frequency (Chuenwattanapranithi et al., 2006; Schröder et al., 2001), speech tempo (Scherer and Oshinsky, 1977; Zetterholm, 1998) and displacement of vowels in more extreme parts of the speaker's oral cavity (Kienast and Sendlemeier, 2000). These findings along with those relating to emotional speech perception suggest that anger can be expressed and inferred in a variety of ways.

The studies presented in this section also raise the issue of what data can be said to accurately reflect the phonetic expression of emotional states. In many of the production studies described earlier, 'actors' are recorded reading aloud prescripted texts. These recordings are subsequently analysed for various phonetic properties which are thought to be associated with emotional speech production. This approach raises a number of issues: mainly, that actors are assumed to in some way provide more 'representative' or 'higher quality' speech data. This assumption appears to be rooted in the belief that actors (particularly professional actors) can produce a wide range of 'clear and unambiguous' emotions (Williams and Stevens, 1972).

However, selecting only participants with acting experience could lead to the collection of data which represents a 'stereotype' of an emotional state or data which does not relate to what the actor would do in a real life scenario. In other words, an actor might produce speech data which represents how they have been taught, rather than data which reflects their own personal impression of what an emotion 'sounds like'. Also, the label 'actor' in research is often not defined whatsoever. Different levels of training and experience between actors may have an effect on their performance. To examine the potential difference between actor and non-actor speech data, this research has collected data from participants who self-identified as belonging to either of these groups. These participants were also able to report on their level of training and experience.

There are also studies which have attempted to avoid these methodological shortcomings by collecting 'authentic' recordings of different emotional states. For example, Williams and Stevens (1972) included recordings taken from the Hindenburg disaster in their analysis. While such data represents a spontaneous and ecologically-valid representation of distress, these data are heavily reliant on accurate interpretation by the researcher(s). For data taken from, for example, highly-distressing situations, the situational context determines the supposed emotional state of the people involved. However, it is possible (indeed likely) that different people will respond to the same event differently. It is also possible that people in authentic situations are experiencing a multitude of emotions throughout the recording (or simultaneously throughout the recording). For these reasons, the data might not be truly representative of the target emotion. It is also likely that the production of threatening speech would be affected by overlapping emotions and behaviours.

As touched upon in this section, it is possible that speakers from different cultural and/or language backgrounds might vary in their production of angry speech (and threats). This idea has has been written about as early as the 1800's (Darwin, 1872). In his research, Darwin describes how anger is expressed in 'nearly the same manner' across the world. The exception to the consistent expression of anger is said to be clenching of the fists which is said to be used only by cultures where people fight using their fists. There are a number of 'universal' features noted by Darwin which particularly relate to anger speech production, these include:

- A 'firmly compressed mouth'
- Protruded lips
- Increase in breathing rate

The current study examines data collected only from native English speakers. As such, cultural or language-specific expressions of threat and anger will not be explored here. It is hoped, however, that the methodological approach and findings of this research can be applied to more linguistically and culturally-varied data in future studies.

The studies and theories presented in this chapter intended to provide a broad overview of the prevalence of threats (particularly in forensically-relevant contexts), and how threats have been examined by threat assessment and legal practitioners and linguists. A number of gaps in the available research were identified; especially the lack of empirical research which examines the phonetic and linguistic features which relate to expressing a verbalised threat in experimental conditions as well as in a criminal setting. In addition, it is not clear how threat production and/or perception relates to what is known regarding the production and perception of different emotional states through speech. This research intends to address these gaps by examining the following research questions:

2.4 Research questions

The aims of the current research are summarised by the following research questions:

1. What phonetic features can be associated with a threatening tone of voice?
   (a) Are the phonetic features under investigation in this study, comparable to previous descriptions of emotional speech? In particular, how do the phonetic features that are associated with threatening speech in this research relate to the phonetic features associated with angry speech in previous literature?
   (b) Are there phonetic features which differ between the production of direct and indirect threats? Is there a compensatory effect present during the production

of indirect threats?

(c) Do these phonetic features differ between authentic and simulated threats?

(d) Do these phonetic features differ between speakers with acting experience, and speakers with no formal acting experience?

(e) Do these phonetic features differ between male and female speakers?

2. Are there linguistic features present in the simulated threats collected for this study which relate to the linguistic features said to predict the likelihood that violence will be actualised?

   (a) Do the written direct and indirect threats conform to the definitions of 'direct' and 'indirect' threats coined in this study? If not, how could the current definitions of direct and indirect threats be changed to better reflect this difference?

3. Is there a relationship between the typographical features of a scripted threat and its subsequent phonetic production?

These research questions aim to address several gaps in our knowledge about threatening language. Firstly, whether there is any phonetic evidence to support the idea of a 'threatening tone of voice'. This research focuses on the phonetic features which have been previously described in relation to threatening speech, or have been predicted to occur during threatening speech. These phonetic features are: fundamental frequency and intensity (the acoustic correlates of pitch and volume respectively), articulation rate (associated with speech tempo) and vocal tract and voice quality features. As such, this research is focused on providing a broad phonetic analysis of threatening speech. These threatening speech recordings are also compared a sample of the same speaker's non-threatening speech. This is in order to address whether speakers (across the board) produce threats with phonetic cues which deviate from their more typical, non-threatening speech. In addition, this research also considers whether there is a difference in the phonetic production between direct and indirect threats. In particular, indirect threats are typically characterised by a lack of overt references to harm. Therefore, it is possible that speakers might need to communicate this intent via an enhancement (or modification) of phonetic features relative to their typical speech. This enhancement of phonetic features is what is referred to as a 'compensatory effect' in research question 1.b.

As there are no other studies which have already examined the phonetic properties of threats, there is currently no framework available for the collection of threatening speech.

As a consequence, this study borrows heavily from the approaches used in emotional speech research. Previous studies on angry speech have analysed speech data from authentic (or real world) scenarios. In order to examine whether this approach is applicable to forensically-relevant data, this study provides a phonetic analysis of both authentic and simulated speech data. In addition, this research will consider whether speakers under experimental conditions produce threatening speech using similar phonetic cues to that of angry speech. This study will also consider whether indirect threats are delivered using different phonetic features to direct threats, as a means of expressing an intent to harm when there are (potentially) no words which explicitly convey this intent.

Another tendency of emotional speech research is the use of acted data. As previously described in the current study, this approach might lead to data which does not represent how people 'typically' produce different emotions. This study will examine data provided by participants who self-identify as actors and participants who reported having no formal acting training or experience. Finally, because much of the background literature concerning threatening speech is centred on male productions, it was of interest to this study to see whether females produced threats which differed in their usage of phonetic features. Examining female data would shed some light on the extent to which factors such as biological sex and gender identity influence the vocal production of threats.

The linguistic data collected for this research are not intended to compare threatening to non-threatening data. There is instead a focus here on validating these simulated data as representative of threatening language. Therefore, this study examines the presence (or absence) of many of the linguistic features that have previously been associated with threatening language and the actualisation of violence. In addition, this thesis seeks to better understand how lay-people interpret the difference between indirect and direct forms of threatening language through the language used when authoring these types of threats.

Finally, this study will also examine the relationship between typographical features present on the handwritten threats collected for this research and the speaker's subsequent phonetic production of the corresponding text. This research question was formed rather more organically than the others, insofar as its inclusion in this study only came about because some of the authors had used typographical features in their written texts. The findings relating to this research question have applications to a broader range of issues in forensic phonetics and forensic linguistics. These include the phonetic production of written transcriptions in courts, or the inference of phonetic and prosodic features based on written testimony,

The following chapter provides an overview of the methods used in the collection and analysis of the data of focus to this study.

3 Methodology

This chapter details the methodology used in this study, which aimed to collect data which addressed these research questions (and their corresponding sub-questions). Firstly in §3.1, this chapter clarifies this research's position on threats and other related speech acts (i.e. warnings and promises). Following this, in §3.1.1-3.1.2, the framework used in this research to describe direct and indirect threats more specifically is presented.

Later, this chapter will detail the data collected for this study, and their relevance to the research questions of this study. Firstly, by clarifying what is meant in this study by 'authentic' and 'simulated' data (§3.2-3.3). In addition, the issues, limitations, and ethical considerations of these data will also be discussed. §3.4 describes the methods used to analyse the spoken threats collected for this research. The following section (§3.5), details the methods used to analyse the simulated written texts collected for this research, as well as describing the statistical tests that were applied to these data.

3.1 Terminology used in the current study

In this study, the definition of a 'threat' deviates slightly from the descriptions of threats, warnings and promises as discussed in Chapter 2.2.4. These differences are highlighted in **bold font** in Table 2. The main differences between the definition of threat in this study (compared to previous linguistic literature) is that firstly, threats are able to benefit the threatener. This 'benefit' could, for example, take the form of successfully coercing someone into a course of action which assists the threatener. Using a threat to dissuade the recipient from testifying in court against the threatener, would be an example of this type of coercion. Note, that making threats may not *necessarily* result in an improved state of the world for the threatener. However, to maximise this possibility, the threat must be mutually understood by someone who can reasonably be expected to influence the situation.

It is also questionable whether the recipient of a threat can necessarily bring about the threatener's desired outcome. For conditionally-phrased threats, it can be argued that by satisfying the condition, the recipient can avoid a negative sanction. However, for the authentic threats in this study, there is no evidence available that the recipient acted in a way that appeased the threatener as a direct result of the threat itself. It is also unknown whether the 'threateners' included in this study would have kept their word if the condition stated had been satisfied.

In addition, this research considers warnings to be distinct from threats in terms of whether the speaker controls the outcome. In this research, the speaker of a threat is seen to be largely in control of the outcome. In contrast, the outcome of warnings are not regarded as being in control of the speaker. This allows for a distinction to be made between speech acts which are made in response to situations brought about by the speaker, and speech acts which describe the outcome produced by other agents. For example, if a speaker told an interlocutor that the weather was forecast to be dangerous, this could reasonably be interpreted by the interlocutor as a warning. It is less likely that such a statement could be interpreted as a threat, as the speaker cannot be expected to control the outcome (that of dangerous weather conditions).

The differences between these speech acts (as used in this study) are highlighted in bold font in Table 2:

Table 2: Overview of threats, warnings and promises as used in this study

Feature	**Threats**	**Warnings**	**Promises**
To the speaker's benefit	**Yes**	No	No
To addressee's benefit	No	Yes	Yes
To the speaker's detriment	No	No	No
To the addressee's detriment	Yes	No	No
Speaker controls outcome	Yes	**No**	Yes
Addressee controls outcome	?	?	?
Speaker committed to act	No	No	Yes

The following sections describes the working definitions of direct, indirect and conditional threats as well as threats of suicide that are used in this research.

3.1.1 Direct threats

In this study, direct threats are defined as expressions which explicitly state that a target(s) could conceivably be the victim of a harmful act[8]. This harm may be perpetrated by the threatener, or alternatively, by others acting on behalf of the threatener. Therefore, a direct threat requires an overt reference to the commission of physical, sexual, economic, emotional, psychological or other form of harm.

The target of the threat does not necessarily have to be specified, or limited to a single person (e.g. *'I'm going to kill the next person I see,'* or *'I'm going to kill you all'*). However, it is necessary that direct threats reference the intended victim(s) of the harm. This is the case even if the message is communicated to a third party, and not directly to the target of the threat. In this situation, the third party may be fearful for the safety of another person, or an object. This enables examples such as *'I'm going to kill your dog, if you don't pay up'* to be classified as a direct threat. This is despite the stated target (the dog) having no awareness of being threatened. Instead, it is the dog's owner (a third party) who is fearful for the safety of their pet. In effect, the third party is another target of the threat - albeit one which is not explicitly referenced. For this reason, target reaction is not a necessary feature of a direct threat in this study. In legal cases, it is also not always established whether the target(s) experienced fear as a result of the message. It is also not expected that the experimental data collected for this study would ever cause anybody to fear harm. Defining direct threats based on target reaction would impose ethical constrictions, which would impact much-needed research in this area.

A lack of available information on target reaction to threats, shifts the focus of this research towards the actions and intentions of the threatener. For a direct threat to be established, the threatener in communicating their message, must either:

- intend for their message to induce fear in their target
- or be aware that their message could be inferred by the recipient of the message as a threat

These additional factors aim to prevent false, or unsubstantiated accusations from determining whether an utterance is threatening. In addition, these factors also allow for threats which were recorded, but never delivered to their target to be examined. Further-

[8]Note, that this 'harmful act' does not have to be a criminal act

more, attempted threats which might not have induced fear, can also be described as a threat. Rather than excluding these examples from threat speech research, they can be investigated for why they did not realise the intended effect on their target. It is important to emphasise here that the threatener may have no intention to *carry out* a harmful act. Instead, the threatener may intend to induce fear in their target through referencing a harmful act. As discussed in section 2.2, there is currently no method which can accurately infer intent through analysing speech features, or text. Therefore, this research will not use data which relied upon the author's inference of the speaker's (or writer's) intention.

Another requisite for a direct threat is a reference to the threatener's identity (or the identity of those perpetrating the 'harm'). Take an example such as *'You're going to be shot.'* Here, there is a clear target ('you'), and a clear reference to violence ('shot'). However, what is not clear is who - if anyone - is going to commit the violent act. This example could quite as easily function as a prediction, or as a warning *(e.g. 'Don't go outside, or you're going to be shot').* It is a statement that ultimately benefits the recipient. Without a clear reference to *who* will commit the harmful act, this example is better categorised in this framework as a possible indirect threat (see Chapter 2.2).

Direct threats can be expressed using different sentence types, such as declaratives *('I'm going to kill you)'*, or as interrogatives: *('Do you want me to kill you?')* While the second example is worded as a question, pragmatically, it does not function as one. It might be more appropriate to describe this example as a rhetorical question. In this situation, giving an answer would be considered unusual, or ineffectual in changing the threatener's course of action. Direct threats can also be phrased conditionally (see Chapter 2.2).

The requisites for direct threats, as used in this study, are summarised in Table 3:

Table 3: Summary of the definition of direct threats as coined in the current study

Feature	Required?
Overt reference to harm	Yes
Overt reference to perpetrator's identity	Yes
Overt reference to target of harm	No
Target feels threatened	No
Threatener intends to induce fear **or** the threatener is aware that their message appears threatening to the recipient	Yes

3.1.2 Indirect threats

In this study, indirect threats are defined as expressions which subtly communicate that harm could conceivably befall a target(s). Similarly to direct threats, indirect threats here are also defined based on the words used, and the speaker's intention. As these speech acts are phrased in a more subtle way, indirect threats are substantially more difficult to identify from existing data than direct threats. This is largely due to the lack of overt references to the commission of violence. On the surface, indirect threats can appear to be requests (*'Could you bring the money by tomorrow?'*), or even compliments (*'Those are some lovely children you have there'*). However, in certain circumstances, these sentences can communicate an intention to commit harm towards the recipient or another entity.

As with direct threats, indirect threats also do not require a clearly-defined target(s). However, a reference to the identity of the perpetrator would be necessary. For this reason, it is again necessary to establish whether the threatener intended to induce fear. As noted in Chapter 2.2, empirically *proving* a speaker's intent based on words or phonetic features is currently not possible. However, it may be possible to demonstrate that a threatener is aware that their message may be interpreted as a threat.

Phrasing a threat in an indirect way also provides the threatener with plausible deniability. Any suspected 'language crime' can easily be portrayed as something innocuous. For example, as an innocent request for information, or as a compliment. As such, denying any wrong-doing during arguments, or during a legal dispute, is a relatively simple task for a threatener. For this reason, collecting indirect threats from internet sources, or forensic casework, is problematic. In authentic cases, it is unknown whether the threatener truly intended to threaten anybody. It is also extremely difficult to infer whether an intent to harm could reasonably be assumed by the recipient (based on the content of the alleged threat, and the circumstances surrounding the 'threat'). As with direct threats, it is unknown whether the alleged victim truly felt fear.

Therefore, the indirect threats investigated in this study are only collected from recordings made under experimental conditions. This allows for speaker intention to be controlled. This is because the participants' goal is to successfully communicate an intention to harm. In addition, defining an indirect threat is not dependent on target reaction. This avoids any potential to cause distress to participants. Authentic data which were alleged to be indirect threats are also not investigated here for reasons later explained in Chapter 3.2.

The requisites for indirect threats, as used in this study, are summarised in Table 4:

Table 4: Summary of the definition of indirect threats as coined in the current study

Feature	**Required?**
Overt reference to harm	No
Overt reference to perpetrator's identity	No
Overt reference to target of harm	No
Target feels threatened	No
Threatener intends to induce fear **or** threatener is aware that their message appears threatening to the recipient	Yes

3.1.3 Conditional threats

Conditional threats are frequently described in threat assessment literature as 'if...then' clauses. These statements can be expressed using the formula:

If X, then Y - where X denotes a condition, and Y denotes a consequence, should the condition not be met to the speaker/author's satisfaction.
For example, *'If you don't invite me to your party, then I will be very upset.'*

Alternatively, it is possible to construct a conditional statement using 'or,' as expressed below:

X, or Y - where, as before, X denotes a condition, and Y denotes a consequence.
For example, *'Invite me to your party, or I'll be very upset.'*

In the latter condition, 'if' is implied, but not overtly stated. The implication is that *if* the 'threatener' was not invited to the party, then he would feel very upset. It is also possible for other words or phrases to replace 'if,' while retaining the same meaning. For example, *'As long as you invite me to your party, then I won't be upset'* or *'Provided that you invite me to your party, I won't be upset.'* Similarly, 'or' can be replaced with another word or phrase which allows for the same meaning to be expressed. For example, *'Unless I'm not invited to the party, I won't be upset'* is equivalent in meaning to *'Invite me, or I'll be upset.'* Therefore, statements which do not use the words 'if' or 'or' can still function as conditional statements.

Note that it is ungrammatical to use both 'if' and 'or' as a means of distinguishing the condition from the consequence, in the same conditional statement [9] :

**'If you don't invite me to your party, or I'll be very upset'*

**'If you invite me to your party, or I'll be very upset'*

As such, conditional statements can be phrased in multiple ways. However, they require two key elements: a condition, and a consequence that might occur should the condition not be satisfied. These elements can be presented in either order:

Condition	**Consequence**
'If you don't invite me to your party...'	'...then I will be very upset'
Consequence	**Condition**
'I will be very upset...'	'...if you don't invite me to your party'

Conditional, unlawful threats can be structured in a comparable way to the conditional statements featured above:

Condition	**Consequence**
'If you don't invite me to your party...'	'...then I will kill you'
Consequence	**Condition**
'I will kill you'	'...if you don't invite me to your party'

[9] It would be acceptable to use 'if' and 'or' in a sentence such as, *'If you don't invite me to your party or your house, then I will be upset'* - where 'or' allows for additional conditions to be stated.

In these examples however, the consequence is clearly an expression of harm toward the recipient of the message. While conditional sentences may include a negative (or positive) sanction, *conditional threats* require an overt, or implied, negative sanction. Though this negative sanction is not guaranteed to have the desired effect on the recipient (i.e. inducing fear), there is at least an attempt by the threatener to bring about this effect via their message. Conditional threats therefore, can be thought of as a subcategory of conditional statements. These statements require some reference to a negative sanction. Note, that negative sanctions can include acts such as sending a child to their room if their behaviour does not improve, or telling a manager about a romantic affair in a workplace if the affair continues. These examples can be classified as conditional threats, even though they are not necessarily criminal acts.

By modifying the 'consequence' of a conditional sentence, it is also possible to create a conceivable indirect threat that is conditional:

- *'If you don't invite me to your party, you'll see.'* or
- *'Invite me to your party, or else.'*

Such examples of indirect threats, feature no words which denote violence or harm, and the 'consequences' stated are vague. The wording of these consequences could be interpreted as ominous, yet the 'threatener' may not have foreseen this reading. As with non-conditional indirect threats, accurately identifying indirect threats which are phrased conditionally is highly problematic.

Like the consequence, the condition of a threat can also be modified to be explicit or non-explicit. This 'explicitness' is not only limited to overt references to violence or harm. It can be also be thought of as how clearly the condition is expressed (i.e. whether what has to be done in order to satisfy the condition is obvious or not). As illustrated in Table 5, indirect threats can feature either an explicit or non-explicit condition, but require a non-explicit consequence.

Table 5: Explicit and non-explicit conditions and consequences in indirect and direct threats

Condition	Type of condition	Consequence	Type of consequence	Type of threat
‘Give me all your money...’	Explicit	‘...or I’ll kill you’	Explicit	Direct
‘Give me all your money...’	Explicit	‘...or else’	Non-explicit	Indirect
‘You better do this...’	Non-explicit	‘...or I’ll kill you’	Explicit	Direct
‘You better do this...’	Non-explicit	‘...or else’	’Non-explicit	Indirect

Using this framework, conditional statements such as *’If you don’t shoot him dead, you’re going to get it’* would constitute a possible indirect threat. This is despite the condition clearly referencing a violent act. The ‘indirectness’ of this threat stems from the consequence not being clear (based on the threatener’s words alone). As with non-conditional indirect threats, additional information about the situation would be required. This would more strongly suggest that a threat had or had not been intended.

For this study, threats which are phrased conditionally (featuring both a condition and a consequence), are categorised as either direct or indirect threats [10]. This is opposed to a distinct category such as ‘conditional threats.’ This is because it is the *consequence* of a conditionally-phrased threat that communicates (possible) harm to the target. The condition of the statement, does not necessarily express harm to the target. The conditional phrasing is as an alternative means to induce fear in the target(s).

3.1.4 Threats of suicide

Unlike the types of threats discussed here so far, threats of suicide express an intention (whether actual or otherwise) to harm the threatener themselves, and not the recipient of the message. Such threats require that the ‘harm’ stated fits the following criteria:

- the ‘harm’ would reasonably pose a risk to the threatener’s life
- the ‘harm’ would be self-inflicted by the threatener (i.e. not inflicted by another person)

[10] Based on the criteria outlined in section 3.1.1 and section 3.1.2

The life-threatening 'harm' that is expressed, can be stated explicitly; for example, *'I'm going to hang myself, if the money doesn't come through in time,'* or *'Do you want me to kill myself?'* These examples would constitute a direct threat to commit suicide. From the words used, it is clear that there is a harmful, life-threatening act, and that this act will be directed towards the threatener. In contrast, the 'harm' stated in a threat of suicide can be unclear, as in, *'I don't want to live anymore,'* or *'You'd be better off if I wasn't around.'* In these cases, the words do not make it clear what course of action the threatener wishes to take (if any). Nor is there necessarily a clearly defined target. Similarly to indirect threats, more information would be needed in order to establish whether the threatener intended to make a threat.

Crucially, threats of suicide do not require suicidal intention. The threatener does not need to feel suicidal when producing their 'threat,' or have any intention of endangering their own life. As with direct and indirect threats, the threatener must convey to another person, an intention to inflict harm. More specifically, this 'harm' must be both self-administered, and life-threatening. Threats to commit self-harm can also be directly, or indirectly phrased. These threats share mostly the same criteria as threats of suicide; however, the harm specified (or implied) is not immediately life-threatening.

To reiterate, threats of suicide are treated in this study as a subcategory of direct and indirect threats. These threats differ from those specified in §3.1.1 and §3.1.2, in that they are specifying, or implying, and intention to commit harm towards the threatener themselves.

Table 6 summarises the criteria for direct and indirect threats as defined in this research:

Table 6: Summary of direct and indirect threats as defined in this research

Feature	**Direct threat**	**Indirect threat**
1. Overt reference to harm (towards another being, or self-directed)	Yes	No
2. Overt reference to perpetrator of harm	Yes	No
3. Overt reference to target of harm	No	No
4. Target feels threatened	No	No
5. Threatener intends to induce fear **or** is aware that their message appears threatening to the recipient	Yes	Yes
6. Can be phrased using different sentence types (declarative, interrogative etc.)	Yes	Yes
7. Can be phrased conditionally	Yes	Yes

The data collected for this research consists of a set of authentic threatening recordings, and a collection of simulated threatening recordings and texts. These data are described in the following sections.

3.2 Authentic data

In this study, 'authentic' recordings are defined as real-life speech recordings which were not originally created for research purposes. The authentic speech data analysed for this research were collected from two sources: data available on publicly-available video streaming websites, and data originally analysed by staff based at the forensic speech and audio laboratory JP French Associates, a York-based forensic speech and audio laboratory. Ethics approval was sought and obtained from the Department of Language and Linguistic Science at University of York prior to the collection and analysis of these data. These data were stored on an encrypted hard-drive. This hard-drive was locked in a safe at the University of York when not in use.

Table 7 provides a summary of the authentic data collected for this research. The speaker at the centre of each authentic case is referred to as 'A (as in 'authentic') 1,' 'A2,' 'A3' and so on. This is to ensure anonymity.

Table 7: Summary of the authentic data used for phonetic analysis:

Speaker	**Nationality of threatening speaker**	**Description of threatening sample**	**Description of non-threatening sample**	**Duration of non-threatening speech sample - prior to editing for f0 and AR calculations(seconds)**	**Duration of threatening speech sample -prior to editing for f0 and AR calculations (seconds)**
A1	North American	Phone conversation with girlfriend	Telephone interview	227.9	7.48
A2	British	Phone message made to an unknown caller	Radio interview	151.93	25.79
A3	North American	Televised interview between two competing fighters	Television interview	152.32	2.62
A4	North American	Radio broadcast with phone-in segment	Audio recorded instructional video	230.13	2.16
A5	North American	Video message sent to former trainer	Audio recording of speaker	140.65	63.82
A6	North American	Phone message made to unknown caller	Television interview	106.71	22.71
A7	British	Video message uploaded online targeting the police	Audio recording of speaker	162.76	9.66
A8	British	Covert recording of a conversation made using a mobile telephone	Police interview	124.56	2.54
A9	British	Phone call to emergency services	Police interview	122.71	1.33
A10	British	Phone call to emergency services	Police interview	58.04	4.96

All of the speakers described in table 7 are male. However, as shown in this table, the nationality of these speakers varies. These speakers also vary in terms of their accent. This lack of consistent data stems from the difficulty of finding both authentic threats and suitable reference samples which were consistent across the board for speaker accent and nationality.

Three of the cases (Speakers 8-10 in Table 7) were archived materials selected by staff based at JPFA. These recordings were edited to remove personal information (e.g. names of people and places or telephone numbers). This was in accordance with ethical requirements for sensitive data protection.

The remaining cases used in this investigation were sourced online from publicly available sources (Speakers 1-7 in Table 5). As such, these recordings were not edited to remove personal details. Of these publicly-available data, any which were not originally telephone transmissions were band-pass filtered (between 300-3400Hz) before any phonetic analysis. This was to enable all of these data to acoustically resemble telephone calls, allowing for more robust phonetic comparisons to be made across the data. This band-pass filtering was performed using Praat software (Boersma and Weenink, 2014).

These recordings were primarily selected due to the availability of both a 'threatening' speech recording, and a reference recording consisting of (presumably) non-threatening speech. Crucially, this threatening recording could reasonably be interpreted as a direct threat. As discussed previously in this thesis, there is no method available at present which can establish whether or not a speaker intended to threaten someone through the tone of their voice or their choice of words. Instead, these recordings are treated as examples of direct threats. This is because the content of the speech contains overt references to harm another person or group.

As such, the language used in these utterances would be difficult to interpret as anything other than an explicit threat to harm. It should be stressed that there is no information available surrounding the circumstances of the threat itself. Therefore, it is not possible to establish the speaker's exact motive, or their intent at the time the speech was produced. Nor is it possible to know how the message was interpreted by the recipient(s). This lack of information means that these threats do not necessarily meet the criteria for direct threats as provided in previous threat-related literature.

The materials sourced from the forensic speech and audio laboratory were also selected based on their speaker comparison score. Forensic speaker comparisons involve analysing two speech recordings that are of forensic interest for comparable phonetic and linguistic features. The results of this analysis are presented in terms of the likelihood that the speech in both samples was produced by the same speaker. In forensic speaker comparisons performed at the laboratory, this likelihood is expressed with a number ranging from 1-5. [11]

A likelihood rating of 1 would indicate that the speech in both recordings were highly unlikely to belong to the same speaker. A likelihood rating of 5 would indicate that the speech in both recordings was highly likely to belong to the same speaker. For the authentic data compiled for this study, data which was given a speaker comparison score of 1 were not analysed. This is because the threatening speech and reference speech data could have feasibly been produced by different speakers. However, sourcing data with higher speaker comparisons scores is constrained by the difficulties of forensic speech data.

Recordings submitted for forensic speech and audio analysis are often of a poor audio quality. In addition, there are also often issues with the intelligibility of the speech data. As such, using data from this source for research purposes poses numerous problems. The data provided from a known suspect and a disputed suspect recording are inherently somewhat similar to one another. If the recordings were acoustically (or phonetically) different, it would be unlikely that the police (or other agency) would enlist the expertise of a forensic speech and audio analyst, since accessing this expertise costs both time and money. Likewise, recordings with similar sounding voices would presumably not be submitted for forensic analysis either. As such, the recordings analysed in forensic speech and audio laboratories for speaker comparison purposes are likely to be (superficially at least) similar to one another.

Another issue with authentic materials is the variation regarding the circumstances in which recordings were created. One major issue is the non-contemporaneity of the recordings. It is possible that there is a significant length of time between the two recordings.

[11] Alternative methods for calculating and expressing the results of a speaker comparison are currently being researched and developed.

During this time, changes to the speaker's vocal anatomy through the ageing process or changes to his/her vocal anatomy (e.g. loss of teeth) might be reflected in these recordings. This would likely affect the similarity of the comparison. Related to this, is also how speakers behave across different situations. The speech of a suspect in a police interview is unlikely to resemble his/her speech while committing a bank robbery, for instance.

In addition, many forensic speech materials also feature speakers under the influence of alcohol, medications and/or illegal substances. The effects of alcohol on speech, in particular, have been thoroughly documented (Pisoni and Martin, 1989). This includes not only changes to vocal properties, but also changes to motor control on speech co-ordination and the quantity of speech which is produced (Tanner and Tanner, 2004). Other factors frequently seen in forensic data such as voice disguise (Fecher & Watt, 2013), stress (Scherer, 1977), and emotional state (as previously described) are also known to affect the vocal properties of speech.

Another consideration is how the recordings were produced. Authentic materials range from speech transmitted over the telephone, over the internet using a Voice over Internet Protocol (VoIP) service, to speech that is captured on recording devices such as CCTV or body cameras. There is also the possibility that the speech data was passed through multiple recording channels. For example, a forensic case recording could consist of a mobile phone call conversation that was put on loudspeaker, and then recorded by another device. Once again, this could impact the quality of the data used for an analysis. The proximity of the speaker(s) to the recording source(s) would also have an effect on the acoustic signal.

With all of these factors to consider when performing forensic speaker analysis, a likelihood rating score of either 1 or 5 would be improbable. Recordings 7-10 in this study, were given speaker comparison scores of 2 or above (French, personal communication, 2015). This means that while there is acoustic and/or linguistic evidence to support that a suspect and known sample were possibly produced by the same person, there is inevitably room for doubt.

Despite these methodological short-comings, analysing these real-world data provides much-needed ecological validity to this thesis. It also enables comparisons to be made between authentic and simulated samples. **Research question 1.c** is motivated by the extent to which experimental data can be said to resemble real-world language crimes. Through the findings of this thesis, it is hoped that future research in this area will be better equipped to identify further representative samples of threatening language from a variety of sources.

3.3 Simulated data

3.3.1 Participants

Recordings of simulated speech, and written data were collected from native English-speaking participants (n=41). There were 20 male, and 21 female participants (age range = 18-24 years, mean = 19.68 years, standard deviation = 1.49). Just over half of the participants reported having some level of acting experience (10 male and 11 female). The participants who reported acting experience were largely recruited from drama or performance groups based at the University of York via email and word-of-mouth. Most of these 'actor' participants reported that they belonged to amateur dramatics groups, had performed in stage productions, and had studied drama at school and/or college. An advantage of the participants having a small age range, is that their level of reported acting experience (where applicable) seemed to be similar.

Table 8 provides an overview of the participants who were recruited for this research along with their respective identifiers. In this study, each of these participants will be referred to as 'S(speaker)N(number)', for example, 'S32' corresponds to 'Speaker 32' (a female actor).

Table 8: Summary of participants recorded for simulated data

	Male	**Female**
Non-actors	Speakers 1-10	Speakers 11-20
Actors	Speakers 21-30	Speakers 31-41

At the recruitment stage, participants were informed that their voice would be recorded reading aloud text. Participants were also told that the content of some of these texts would involve profanity, as the texts were based on forensic cases. At the data collection sessions, the participants were told that the research was focused on threatening speech in particular. This prevented participants from rehearsing threatening speech, or asking other people about how to appear threatening.

Before the recording session started, participants were provided with an information sheet about the project. After any questions about the research were addressed, participants signed the information sheet and an accompanying consent form. After each data collection session was completed, the participants were given a debrief about the full purpose of this research. As before, participants were given the opportunity to ask any questions or to withdraw completely from this experiment. Upon completion of the debrief, participants received £10 for their time. [12]

3.3.2 Procedure

All participants were recorded in quiet locations across the University of York campus. The equipment and settings used for each recording session was the same, the details of which can be found in the Appendix. Each participant was recorded in a single session performing a series of tasks. These tasks were presented as part of a Microsoft Powerpoint presentation. This allowed the participants to modify the size of the text if needed, and to complete the tasks at his/her own pace undisturbed. While completing these tasks, the participant wore a headset, with a microphone positioned approximately 3cm from their mouth. Once fitted, participants were instructed not to touch or adjust the headset during the recording session. These data were recorded directly in Audacity software (Ash et al., 2015) on a laptop. The sampling rate was 41.kHz. Details of the recording equipment used for this study can be found in the Appendix.

Each task in this experiment was intended to collect specific types of speech and written data. The tasks are described in Table 9.

[12]Self-funded by the researcher.

Table 9: Experimental tasks and corresponding collected data

Task number	Description of task	Data collected
1	Read aloud text	Read, emotionally 'neutral' speech
2	Speaker creates indirect threat based on scenario provided	Read-aloud, 'threatening' speech, and threatening text script
3	Speaker creates direct threat based on scenario provided	Read-aloud, 'threatening' speech, and threatening text script
4	Speaker reads aloud same sentence as instructed	Read speech, portraying different emotional states
5	Speaker reads aloud scripts taken from Authentic speech data in a 'threatening tone of voice'	Read-aloud, 'threatening' speech

Task 1 collected read speech for each participant. These recordings were intended to collect reference, or emotionally-neutral, baseline speech data for each participant. This data served both a reference point for a speaker's phonetic characteristics, and also as a comparison to the other read speech tasks in the experiment (Tasks 2, 3, 4 & 5). For this task, speakers read aloud a copy of the parable 'The North Wind and The Sun.'

Tasks 2 and 3 centred on the collection of what these participants understood to be threatening speech. The participants were provided with a short backstory which was based on a real-world case centred on a farm that was the focus of a dispute regarding animal welfare ('Jail for animal rights extremists who stole body of elderly woman from her grave' (Morris, Ward & Butt, 2006)). The backstory given to the participants was altered to give fictional names and places, and provided a simplified version of these real-life events. In short, the backstory describes a farm that is alleged to be mistreating animals bred for medical research purposes. The participants are told that footage taken from this farm showing this alleged abuse has been uploaded to the internet. In addition, the personal contact information of the owner of the farm (the fictional character 'Tom West') has also been uploaded. This contact information is being shared by a (fictional) pressure group who are campaigning for the farm to be closed down. Using this contact information, the participant is 'able' to telephone the farm owner.

After reading this backstory, the participants were asked to hand write two scripts: an indirectly threatening text (Task 2) and a directly threatening text (Task 3). These texts would form the basis of an indirect and direct threat that would later be 'directed' toward the owner of the farm. The participants were not instructed *how* to write these scripts. This included, for example, the length of the texts or any language that should be included or avoided. However, the participants were informed that their messages should communicate an intention to harm the farm owner. This instruction was included to ensure that the data collected could be said to represent the central topic of this research. After writing down his/her threats, the participants were instructed to read them aloud in a threatening way as if he/she was leaving a phone message on the target's phone. The speakers were allowed to attempt as many renditions of his/her threat(s) as needed, until he/she was satisfied with his/her recording(s).

Task 4 intended to collect examples of emotionally-neutral sentences which were read aloud in whatever the speaker understood as an 'angry tone of voice.' For this part of the experiment, the participants read aloud the following sentences: 'After what happened, I'm going to lose it tomorrow,' and 'After what happened, I'm going to kill him tomorrow'. These sentences were presented individually along with an instruction on how to read them aloud. For example, 'Read this sentence aloud in an angry way' or 'Read this sentence aloud while gritting your teeth.' Only the 'angry' productions of these sentences were analysed for this research. The remaining sentences were intended to act as 'filler' or dummy sentences to disguise the true purpose of this task. For each participant, these sentences were presented in a randomised order.

Finally, participants completed Task 5 of this experiment. For this task, participants were provided with transcriptions of the authentic threat data presented in this research. The participants were instructed to read these scripts aloud in a 'threatening way.' Crucially, the speakers did not listen to the original authentic threat audio recordings. Therefore, his/her production were not imitations or reproductions of the authentic threats. This data was collected to gather vocal productions which contained directly threatening language. This is partly because it was not known whether the data collected for Task 3 would provide data which could be classified as direct threats using the framework described earlier in this research.

All of the instructions for each task, as well as all of the threatening texts written by these participants can be found in the Appendix. It should be noted that Speakers 24, 25 and 30 did not write down any 'threats' in response to Tasks 2 and 3. Written texts were collected under this experimental setting for a number of reasons. Firstly, because threats of a forensic nature have been (and/or could conceivably be) written down before being communicated via, for example, over the phone. However, what is unknown is the extent to which people who produce spoken criminal threats would have originally scripted their threats. Secondly, it was seen as beneficial to analyse spoken material along with a transcription that was produced by the author. This allowed for any differences between the spoken and written threat, and/or interesting typographical features to be examined if applicable. Finally, based on previous pilot experiments ran by the author, participants often reported their difficulty in producing threats spontaneously. Allowing participants to write down their own messages prior to any audio recording, seen to help reduce any anxieties based on the performance of spoken threats.

This method of collecting speech data is limited by the extent it can apply to real-life threats - more specifically, threats which are produced spontaneously. What this data does represent, however, is the speaker's attempt to appear threatening. It is recording and describing these attempts that is paramount to the current study.

The following section provides detail about the phonetic and linguistic methods used to analyse the authentic and simulated data collected for this thesis.

3.4 Methods used for spoken data analysis: authentic and simulated data

The phonetic data analysis in this thesis combines auditory and acoustic approaches to examine the differences and similarities between either:

- what is assumed to be threatening and non-threatening speech samples taken from the same speaker (authentic data) - or,
- the same speaker in experimental conditions producing threatening and non-threatening speech (simulated data)

This auditory-acoustic approach reflects the manner of forensic phonetic analysis currently practised in the United Kingdom (R v Flynn and St John [2008] Crim LR 799). By comparing the vocal productions of the same speaker across multiple experimental conditions, in effect, this thesis presents the findings of 51 speaker comparisons. However, a key difference is that in a typical forensic speaker comparison, the intention is to ascertain the likelihood that the speaker in the criminal recording is also the speaker in a suspect recording. In this study, the analysis of the authentic threat data most closely resembles this practice. However, the intention here is to investigate whether these speakers adopt what could be described as a 'threatening tone of voice' while making what is likely to be a threat.

The numerous difficulties presented when analysing authentic data have been previously described in this thesis. By collecting data in experimental conditions, it is hoped that there will be minimal interference from confounding factors (such as those noted for authentic materials). For the simulated materials, the identity of the speakers is already known, as are the conditions that these recordings were created in. The speaker comparison for these data intends to provide a more robust comparison of the same speakers producing threatening and non-threatening speech.

The following sections details how each phonetic and linguistic feature of interest to this study was analysed. These parameters were chosen due to their connection to threats in previous research or theories centring on threatening language. In addition, this study seeks to describe the broad phonetic and linguistic properties of threats with a view to identifying possible areas of interest for future relevant research.

3.4.1 Fundamental frequency: authentic data

In this study, measurements of fundamental frequency (f0) were made using Praat software (Boersma and Weenink, 2014). Each sound file was edited before the extraction of these measurements to remove instances of background noise or speech which was irrelevant to the current study. As the authentic data consisted of male voices, Praat pitch settings were initially set at a minimum level of 60Hz and a maximum level of 150Hz. These settings were adjusted for voices which exhibited f0 values which fell outside of this range.

Each recording was then converted to a pitch object in Praat. This process allows for the manipulation of the available component harmonics (or 'candidates'). These pitch objects produce synthetic tones which correspond to whatever candidates are selected at the time. By manipulating these candidates, the pitch object can be modified to produce tones which more closely resemble the perceived pitch of the corresponding speech. From these 'corrected' pitch objects, the following measures were extracted from Praat manually:

- minimum pitch (Hz)
- mean pitch (Hz)
- median pitch (Hz)
- upper (75) and lower (25) quartiles (Hz)
- maximum pitch (Hz)
- standard deviation (Hz)

3.4.2 Articulation rate: authentic data

Articulation rate (AR) provides a measure of the tempo or pace of fluent speech. The numeric output represents the number of syllables produced per second. Typically, for speakers of English, AR is between 4-6 syllables per second (Laver, 1994). Figures lower than this range would indicate very slow speech, while higher figures would indicate very fast speech.

In this research, articulation rate was calculated using the methods discussed by Künzel (1997). In keeping with this method, before AR was calculated, each speech recording was edited to remove pauses of >0.1 second. All pauses in each experiment task which were longer than this were removed, with the exception of pauses which occurred within words. After each speaker's task recording was edited, the articulation rate was calculated. This involved manually counting the number of syllables which the speaker actually produced. During fluent speech, speakers regularly omit or reduce syllables (e.g. pronouncing the word 'library' with two syllables as in /laɪ.bri/, as opposed to three as in /laɪ.brə.ri/).

After the number of syllables for each task recording was counted, this figure was divided by the duration of the edited recording. These calculations produced an articulation rate output for each task recording for each of the speakers involved in this research.

3.4.3 Fundamental frequency: simulated data

The simulated data were edited using the same procedure as described previously in relation to the authentic data. For female data, Praat pitch settings were set at a minimum level of 170Hz and a maximum level of 250Hz. Due to the larger number of simulated data (than authentic data) the following measurements were extracted from Praat using a script, rather than manually using Praat's interface:

- minimum pitch (Hz)
- mean pitch (Hz)
- median pitch (Hz)
- Upper (75) and lower (25) quartiles (Hz)
- maximum pitch (Hz)
- standard deviation (Hz)

This pitch measurement extraction script is shown in the Appendix.

3.4.4 Intensity: simulated data

Intensity measurements were extracted from the edited task recordings using a Praat script, rather than selecting each measurement separately using Praat's interface. The previously described Praat pitch extraction script was modified by the researcher to extract the following intensity measurements automatically:

- minimum intensity (dB)
- mean intensity (dB)
- median intensity (dB)
- Upper (75) and lower (25) quartiles (dB)
- maximum intensity (dB)
- standard deviation (dB)

The extracted data were exported as a .csv file and later as a Microsoft Excel file. The intensity script is shown in the Appendix. As reliable intensity measurements could not be guaranteed from the authentic speech materials (see Chapter 3.2), this thesis will only present the intensity analysis of simulated speech materials.

3.4.5 Articulation rate: simulated data

For the simulated data, the method of calculating articulation rate was identical to that of the authentic data. As explained earlier in this subsection, only syllables which were actually produced by the speaker were counted for this analysis. Even though much of the simulated data is read speech, there are a number of false starts, repetitions and improvisations present in these spoken data. For this analysis, these disfluencies or additional linguistic content were not removed from the data. This was in order to maximise the quantity of data available for each AR analysis. Therefore, even though multiple speakers read aloud the same text (in similar experimental conditions), there are differences in the number of syllables produced and the duration of each speaker's reading.

3.4.6 Vocal profile analysis: simulated data

In forensic phonetic casework performed in the United Kingdom, a modified version of the Edinburgh Vocal Profile Analysis Scheme, the VPA protocol (Beck, 2007), is used to catalogue the presence of various features relating to vocal setting and vocal tract features. In this study, vocal setting refers to the configuration of the vocal folds during speech. For example, holding the vocal folds tightly (but not completely) together results in creaky voice. 'Vocal tract features' refers here to changes made to the dimensions of the vocal tract during speech production. For example, describing the lips as spread apart as opposed to rounded, or describing the perceived height of the larynx during speech.

The author of this research has been trained to perform vocal profile analysis of both forensic and non-forensic speech data as part of a Masters level programme in Forensic Speech Science. This programme included extensive ear-training on recordings of speakers who have or are adopting different vocal tract or voice quality settings. For example, inferences of vocal properties (such as a raised larynx or tongue-fronting) present in the speech samples collected for this research, could be likened to recordings which were created or selected to exemplify these properties. The author was solely responsible for the analysis presented in §5.2.8 and §5.3.4.

Using the VPA protocol involves carefully listening to each recording and performing an impressionistic analysis. As such, compared to the other forms of phonetic analysis discussed so far, this analysis is inherently more subjective. A copy of this modified VPA protocol was completed by the author for all Task 1-4 recordings collected for this research under experimental conditions. This scheme allowed for non-neutral (or non-modal) vocal tract features or vocal settings to be recorded. In addition, the extent or degree of these non-neutral vocal tract features or vocal setting was noted on a scale of 1-3. For example, describing a voice sample as 'Creaky voice (3)' would indicate that there is an extreme level of creaky voice. A copy of a blank modified VPA protocol can be found in the Appendix.

However, it should be noted that because of the subjective nature of this analysis, it is common in forensic phonetic casework for multiple practitioners to calibrate their individual analyses. Owing to the large quantity of speech data collected for this research, it was not feasible for this calibration to take place across the data. As such, this analysis should be taken as an initial impression of the vocal tract features and vocal settings used during threatening speech. If any apparent voice quality or vocal setting features of interest is noted in this research, these can be scrutinised in further detail in any subsequent acoustic phonetic research.

Voice quality was not analysed for the authentic recordings sourced for this research, as it was thought that this property of speech would be particularly susceptible to changes that could be accounted for by contextual or environmental changes between the non-threatening and threatening recordings. As previously described in §3.2, the authentic recordings varied in terms of the situation the speech was produced in or the recording device(s) used. For example, some of the non-threatening speech samples (A8-A10) were made during a police interview setting, where it might be expected for speakers to speak differently to their more typical, modal speech, as well as (possibly) their threatening speech. These issues were encountered initially during the examination of fundamental frequency and speech tempo presented in §5.1. The inherent variability of the authentic recordings led to the decision that the results of any further phonetic analyses would also be influenced by contextual or environmental changes, as opposed to changes from non-threatening speech when the speakers make a threat. As such, the collection of simulated data allows for a more robust comparison to be made between threatening and non-threatening speech, and minimises the interference of other factors relating to voice quality features.

3.5 Methods used for written data analysis: simulated data

The following computational methods were used to analyse the written 'threatening' scripts produced by participants during Tasks 2 and 3 of the experiment. As these scripts were handwritten, they were first typed up digitally using Microsoft Word software. These digitised scripts have a number of key differences from the original handwritten documents.

Any spelling errors presented in the original documents were corrected when being typed up digitally. This was to ensure that as much data as possible could be recognised (and therefore analysed) by the computational tools used in this study. In the event that any words or phrases were unintelligible, these would have been excluded from the analysis. However, this was not necessary as these data were legible. In addition, any words or phrases which were obscured or crossed-out by the author were not included in the digital copies of the texts. This is because these digital reproductions aimed to capture (as closely as possible) the words chosen by the author to communicate an intent to harm. As such, any words which were obscured can be assumed not to reflect this intent. It was also not possible for the computational tools used in this research to infer any meaning from underlined text. Therefore, this typographical feature was not expressed in the digitised transcripts.

3.5.1 Tropes software: simulated data

Tropes software (2014) was used to provide a semantic analysis of the digitised texts described above. This software was chosen as it is free to download and has previously been used to analyse the inferences of the emotional content of texts, albeit texts which are not forensic in nature (Piolat and Bannour, 2009; Dzogang et al., 2010). The Word .doc files relating to indirect threats (Task 2) and direct threats (Task 3) were imported into Tropes separately. The following Tropes analyses are presented in this research:

- Text style
- Reference fields 1 and 2:
- Verbs
- Expressions of doubt

The outputs relating to the first two categories listed above are expressed verbally (e.g. ‘this text is mostly narrative.’). The remaining outputs are expressed as a raw figure (e.g. the texts contained a total of 34 ‘stative’ verbs), as well as a percentage (e.g. 45% of the verbs in these texts were ‘stative’). Descriptions of these categories will be provided alongside their respective results sections later in this research.

3.5.2 LIWC2015 software: simulated data

LIWC2015 (Linguistic Inquiry and Word Count; Pennebaker, et al., 2015) is a computational software tool developed by researchers based in forensic psychology that provides a wide range of outputs (N = 94) relating to the grammatical, emotional and cognitive features that are present in a text(s). This software tool compares the content of an imported text(s) to an inbuilt dictionary created by its developers. This dictionary consists of around 6,400 words, word stems and emoticons. Each of these tokens has been allocated to various categories. For example, in the data collected for this research, the word 'torture' is allocated to the categories 'affect', 'negemo (negative emotion)' and 'anger'.

The words were allocated to these categories based on the findings of a series of investigations involving human judges. For a word to be included in the dictionary (and allocated to particular categories), it needed to be consistently agreed upon by these judges. When a document(s) are imported into LIWC, it provides a raw numeric figure for the number of words present in the document(s). All of the other categories are expressed numerically as percentages. For example, a text could consist of '3.40%' words associated with 'anger'. The combined percentages of each category do not total 100%. This is because, as seen earlier for the word 'angry', a text can be allocated to any number of categories.

This computational tool was chosen to analyse the written threat data collected in the research, and to allow comparisons to be made between the results of this research and that those of previous related research (Pennebaker, 2011; Chandler, 2017). In addition, LIWC has been applied extensively to a wide range of research interests. These interests range from how the tweeting about death or harm could be used to predict solidarity with refugees (Smith, McGarty and Thomas, 2018), the analysis of suicide notes (Malini and Tan, 2017) and exploring the written cues to deception in Spanish language data[13] (Almela, Valencia-García and Cantos, 2012).

LIWC2015 also outputs a number of so-called 'LIWC summary variables'. These categories are referred to as 'analytical thinking', 'clout', 'authenticity' and 'emotional tone'. The intended meaning and purpose of these four categories has been written about publicly; however, it is not publicly known how these specific outputs are calculated by the software (Pennebaker et al., 2015). In addition, while this dictionary consists of a large number of words used frequently in English, there will also be words that it cannot recognise (and by extension, cannot categorise). The percentage of the text that could be analysed by LIWC is expressed by the 'Dictionary' output.

The following LIWC outputs are presented in this research:

[13] LIWC software has so far been adapted to 10 languages including Spanish, Arabic, Russian and Turkish.

- Word count
- Dictionary
- LIWC summary variables
- Posemo and Negemo (positive and negative emotions)
- Anxiety, anger and sadness
- Certain and tentative
- Power
- Past, present and future tense
- Death

More thorough descriptions of these categories will be provided later in this study along with their respective results. In addition, the rationale for the selection of these categories will also be discussed.

3.5.3 Statistical tests

As the majority of the data presented were judged to be not normally distributed, non-parametric statistical tests were chosen. These included Wilcoxon signed-rank tests for the comparison of related data groups, for example, the difference between the mean fundamental frequencies (f0s) for the same speakers delivering indirect and direct threats. Mann Whitney U tests were used to compare categorical, independent groups, such as the mean f0 used in the speech of female non-actors and female actors. Friedman tests were used to examine differences between groups. Repeated one-way ANOVA tests were used to compare the means of three or more groups where these groups represented data from the same participants (e.g. the mean f0s of male actors producing baseline speech and direct and indirect threats). Spearman's Rank-Order Correlation tests were used to measure the strength and direction of the non-parametric association between two variables. Finally, Pearson Correlation Coefficients were used to measure the strength and direction of the parametric association between two variables. These tests were performed using SPSS software (SPSS, 2015).

The following two chapters (4 and 5) detail the results of the analysis performed for this study. Firstly, the results of written text analysis for the simulated data collected for this research is presented (Chapter 4). This is in order to show the content of these written threats, before any discussion of their phonetic production. Chapter 5 examines the phonetic production of threatening speech relative to baseline, non-threatening speech. Initially, by describing the authentic data collected for this research, and then establishing a rationale for the collection of simulated threatening recordings. Finally, the phonetic production of a written text is more closely examined in §5.5-5.10. Providing this additional analysis intends to show how data relevant to practitioners based in forensic linguistics and forensic phonetics can benefit from a more multi-disciplinary approach.

4 Results - Written text analysis

The following subsections do not examine the differences between non-threatening texts and threatening texts, in a way which is analogous to the speech materials that will be analysed in the following chapter. Of interest here, is first, to describe various linguistic features found in these threatening texts. The following sections present the results of various linguistic analyses performed on each of the handwritten threatening scripts collected for this research:

- Task 2 (indirect threat scripted by the speaker)
- Task 3 (direct threat scripted by the speaker)

These linguistic analyses aim to address **research question 2**: 'Are there linguistic features present in the simulated threats collected for this study which relate to the linguistic features said to predict the likelihood that violence will be actualised?' For the current study, the following linguistic features were selected for analysis:

- Text style
- Reference fields 1 and 2
- Verbs
- Expressions of doubt
- LIWC2015 software outputs
- Pronouns
- Modal auxiliaries
- Conditional statements
- Profanity
- Promises and warnings
- Assertive statements
- References to time, location, violence or harm, target and author

Most of these features were selected as they have previously been linked with threatening language, or more frequently, are features which are used to predict the likelihood that harm will be actualised. This comparison aims to assess whether the collection of threatening data under experimental conditions can provide data which relates to authentic materials.

This research will also consider whether there is any evidence to suggest that direct and indirect threats can be characterised or differentiated based on the frequency of linguistic features (see **research question 2.a**). If authors of these texts are found to differentiate 'indirect' and 'direct' expressions of threat, then this might suggest that the definitions of these terms (as coined in this study) might not align with their interpretation by lay-people.

Factors such as author sex and level of self-reported acting experience, and their influence on the linguistic production of threatening texts are not explored in great detail here. This is because several authors did not provide a written text, instead opting for an improvised performance. Therefore, direct (and/or statistical) comparisons between male and female authors, and actors and non-actors would not be robust. As such, this thesis will only present impressions about how various linguistic features might vary across different author groups. However, stronger comparisons between direct and indirect threats can be made. This is because all authors who produced written texts, provided both an indirect threat and a direct threat.

It also needs to be stressed that the written data samples presented in this chapter are very short, and as such, are not ideal candidates for computational analysis. Text analysis software packages (including those presented in this study) are sensitive to how much data they require to process reliable outputs. Therefore, the results of this analysis should be interpreted with caution. These approaches are presented in this thesis to allow for more direct comparisons to be made with previous, related studies into threatening or aggressive language, and corpora representing English language use across different modes of communication.

Later in this chapter, a more qualitative approach to data analysis is presented. This consisted of manually tagging each relevant variable present in each written script using Microsoft Word software. This qualitative analysis intended to provide a more detailed and less automated examination of linguistic features which are particularly relevant to threatening language. Scanned copies of the original copy of each handwritten threatening message analysed in this research is provided in the Appendix.

Each participant recorded for this study will be referred to as 'S(speaker)N'. For example, 'S1' corresponds to 'Speaker 1', and 'S2' corresponds to 'Speaker 2'. As described in §4.2, the number of words written for Tasks 2 and 3 ranged from 8-227. The average length of these written threatening messages was 55.55 words.

4.1 Output of Tropes software analysis

The data for all indirect threats and all direct threats were entered separately into the Tropes software package. The primary goal of this analysis was to establish whether indirect and direct threats differed regarding semantic (or pragmatic) inference. That is to say, whether it is possible that these groups of texts could be distinguished from one another in terms of semantic parameters. It was anticipated that direct threats would feature themes concerning violence or death. On the contrary, indirect threats would lack references to these themes.

This analysis also sought to address a key issue with the approach taken in this thesis for data collection. It was anticipated that indirect and direct threats might differ with respect to semantic themes or concepts due to the structure of the experiment itself. As the participants always produced indirect threats first, it might be expected that these threats would be characterised as being more descriptive, or more centred around establishing a clear narrative, as compared to the direct threat which followed. For the subsequent direct threat, it was possible that authors would opt to avoid repeating information already stated in the indirect threat. If this were the case, these texts ought to be classified by Tropes as having different text styles. In turn, this difference could be attributed to the design of the experiment versus inherent differences between indirect and direct threats.

The following outputs of this software will be presented and discussed in this thesis:

- Text style
- Reference fields 1 and 2
- Verbs
- Expressions of doubt

Text style

The first output produced by Tropes is called ‘Text style.’ For this, the texts are analysed by the software by ‘comparing the distribution of the occurrence frequency of the categories observed in the text with linguistic production norms’ (Tropes, 2014). In other words, Tropes associates various linguistic features within the documents provided with the overall style of the text(s). The style of the text is determined using inbuilt data taken from a large number of 'real-world' texts. Details about these inbuilt data (e.g. the exact quantity and content of these texts) do not appear to be publicly available.

Imported texts can be matched by Tropes to one of the following text style categories:

- Argumentative: the speaker involves himself (sic), argues, explains or analyses to try to convince the interlocutor
- Narrative: a narrator states a series of events, happening at a given time, and in a given place
- Enunciative: the speaker and the interlocutor establish a mutual relation of influence, make their standpoints known
- Descriptive: a narrator describes, identifies or classifies something or somebody

For all of the indirect and direct threats collected for this thesis, the data were described by Tropes as having a 'rather argumentative' text style. As such, the texts collected in this study were associated with the author inserting his/herself into the narrative. Also, that the purpose of these texts is to convince the reader about a particular worldview. Logically, this makes sense in the context of the threatening language collected for this study. These threats are being communicated by the person who is likely to be responsible for any (fictional) actions implied or specified by the threat(s). The recipient of the message (the farmer) is following a course of action that the author is instructed to disagree with, to the extent that the author's motivation for sending his/her messages is a desire to close down the farm entirely.

Following this text style analysis, it was observed that some of the threatening texts begin with explanations or justifications for *why* the speaker is contacting the recipient. This can be seen in Speaker 25's indirect threat:

'Good afternoon, Mr West. I am calling to let you know that your awful treatment of the animals on your farm has not gone unnoticed by the general public, and we shall not rest until we have justice.'

Subsequently, these messages move on to instruct or request the recipient to close down the farm.

'Rest assured that if you do not discontinue this disgusting practice, it will no longer be your livestock's screams that are heard for miles around. Good day.'

These communications then can be seen as a means of convincing someone to do something he/she would have not otherwise done. This is achieved through overtly or covertly referencing some harmful act that might be brought about by the author. That said, this is a form of convincing which can be unethical, and potentially illegal.

A major limitation of this study is that the data collected represent only a specific 'type' of threat. It is probable that collecting threats using a different experimental design would have resulted in alternative text styles. For example, if participants were instructed that they were threatening to fire an employee, it might be expected for the discourse to centre on formally describing a series of events which has led up to the threat itself. This might have led to a more descriptive, formal, or enunciative text style than the data discussed here. Alternatively, it is possible that threats directed towards a person familiar to the participant would result in differences in language use. In addition, their relationship with that person might also affect the words chosen and, subsequently, the semantic and/or pragmatic information that is communicated.

For the data collected for this study, there was no difference identified by Tropes between the text styles of indirect and direct threats. This lack of variance in text style also suggests that the experimental design did not have as much as of an effect on text style as initially feared.

Reference fields 1 and 2

This subsection presents the output of semantic analysis generated by Tropes. Of interest here is whether indirect and direct threats have different semantic themes: in other words, whether readers of these texts might interpret different overarching themes or a difference in 'gist' between these text groups. It is expected that direct threats (overall) will include references to violence which are absent from indirect threats.

To provide a sense of the semantic themes expressed in these groups of texts, Tropes groups together frequently-occurring words which have similar or related meanings or themes. These groups are referred to as ‘equivalent classes’. These equivalent classes are broken down into references based on the words found in the texts. These references are then grouped together to form what is known as ‘reference field 1’ and ‘reference field 2’. Reference field 2 represents the most fine-grained and detailed groups of references based on the words featured in the texts. These finer-grained references are then grouped together to form a broader, more general group known as reference field 1. In other words, this analysis is hierarchical, whereby equivalent classes are used to form groups of references. These groups of references are further broken down into levels which show the general themes covered in the texts (reference field 1), and more detailed themes (reference field 2).

For example, in the data collected for this study, one of the main references identified by Tropes is ‘Farm.’ This is based on frequent use of the word ‘farm’ across the texts, which refers to the location of the target’s workplace (a farm). These words also form a broader semantic context of ‘agriculture’ (reference field 1) and a narrower semantic context of ‘farming.’ This relationship is shown in the table below. Other data from this study are also included to help provide a clearer picture of this semantic classification process.

Word(s)	**Reference**	**Reference field 1**	**Reference field 2**
Farm	Farm	Agriculture	Farming
Mr, man	Mister	Man	Man
Fuck, fucking	Fuck	Sexuality	Sexuality
Police	Police	Law	Police
West	West	Location	Location

There are a number of observations that can be made about the data presented above. Firstly, it is possible for references and reference fields 1 and 2 to have the same name. However, the token frequency of words in these groups might lessen from reference field 1 to 2. This is because the classification becomes more detailed as the analysis progresses. In other words, related terms which were allocated to reference field 1 might be divided into more specific groups for reference field 2 (e.g. ‘Law’ to ‘Police’). Secondly, that terms such as ‘fuck’ and ‘West’ appear to be inaccurately classified. In these data, ‘fuck,’ and ‘fucking’ do not appear to be intended as references to sexuality as such, but as profane or offensive language which is only obliquely related to sexual acts. Similarly, the word

‘West’ appears to be used exclusively by authors in this study to refer to the surname of the target, not as a reference to direction or location.

The following table (10) shows the reference fields of the indirect threat data and the corresponding direct threat data. The numbers enclosed in brackets refer to the frequencies of words present in the data that are associated with these semantic themes.

Table 10: Overview of Tropes reference fields 1 and 2

Indirect threat (Ref. field 1)	**Indirect threat (Ref. field 2)**	**Direct threat (Ref. field 1)**	**Direct threat (Ref. field 2)**
Location (31)	Location (30)	Fight (28)	Farming (27)
Animal (31)	Farming (30)	Agriculture (27)	Weapon (22)
Agriculture (30)	Animal (29)	Animal (24)	Animal (21)
Man (21)	Man (21)	Man (18)	Man (18)
Law (10)	Way (10)	Location (15)	Location (15)
Time (9)	Time (9)	Feeling (15)	Body (14)
Sexuality (8)	Sexuality (7)	Body (14)	Time (11)
Social group (6)	People (6)	Time (11)	Death (7)
Communication (6)	Right (5)	Death (7)	Pain (7)
Business (5)	Person (5)	Behaviour (6)	Way (4)
Person (4)	Business (4)	Sexuality (6)	Building (4)
Life (3)	Justice (3)	Family (6)	Family (4)
Family (3)	Police (3)	Device (4)	Nuclear weapon (3)
	Family (3)	Housing (4)	Person (3)
	Telecommunication (3)	Health (3)	Life (3)
		Life (3)	Sex (3)
			Fight (3)
			Mammal (3)
			Sexuality (3)

These outputs suggest that overall, indirect threats lacked words which could be classified as violent or harmful themes. On the other hand, direct threats expressed semantic themes relating to violence, for example 'Fight,' 'Weapon,' 'Death' and 'Pain.' These themes were ascertained based on words such as 'corpse,' 'murder,' 'torture,' 'gun,' 'bombs' and 'explosives.' These words were only present in the direct threats collected for this study. Also, terms referring to 'Body' were often used when specifying the type of harm that would befall the target. This can be seen in Speaker 17's direct threat: 'I want to see every last drip of **blood** leave your **body**' as well as in Speaker 29's direct threat: '...we're going to pour it (shampoo) down your **mouth**.'). There were individual texts which did not appear to conform to the generalisation that direct threats exclusively featured overt references to harm. These inconsistencies will be discussed further in this subsection (references to harm).

Another difference between the text groups is the frequency of words relating to 'Location'. As mentioned earlier in this subsection, the target's surname 'West' formed most of the words referring to location. From this, it appears that indirect threats featured more references to the target's surname (or full name) than direct threats. However, there are similar numbers of references to 'Man' in both texts. This might suggest that authors chose to refer to their target differently between indirect and direct threats, rather than not address him again during their direct threats.

Verbs

This subsection details the verb categories identified by Tropes, and the frequency of their usage between indirect and direct threats. Currently, there are no predictions made in relevant literature specifically about the use of verbs used in threatening language. However, verbs can convey the harmful or violent acts that are of interest to threat assessment practitioners. Tropes identified 535 verbs across the indirect threat scripts collected for this study and 587 verbs in the direct threats.

The following figures (2 and 3) show the different verb categories in these data. The number within each section of the chart represents the frequency of the verb type across the texts. For example, there are 135 reflexive verbs across the indirect threat texts, and 140 in the direct threat texts.

Figure 2: Verb categories used in indirect threats

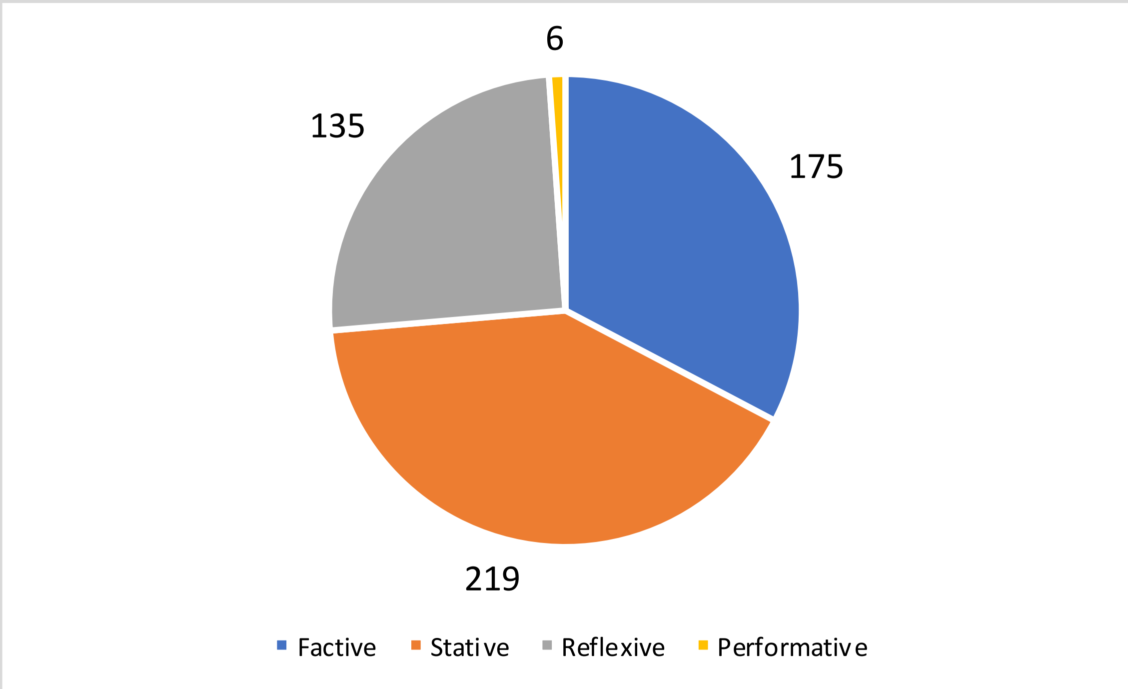

Figure 3: Verb categories used in direct threats

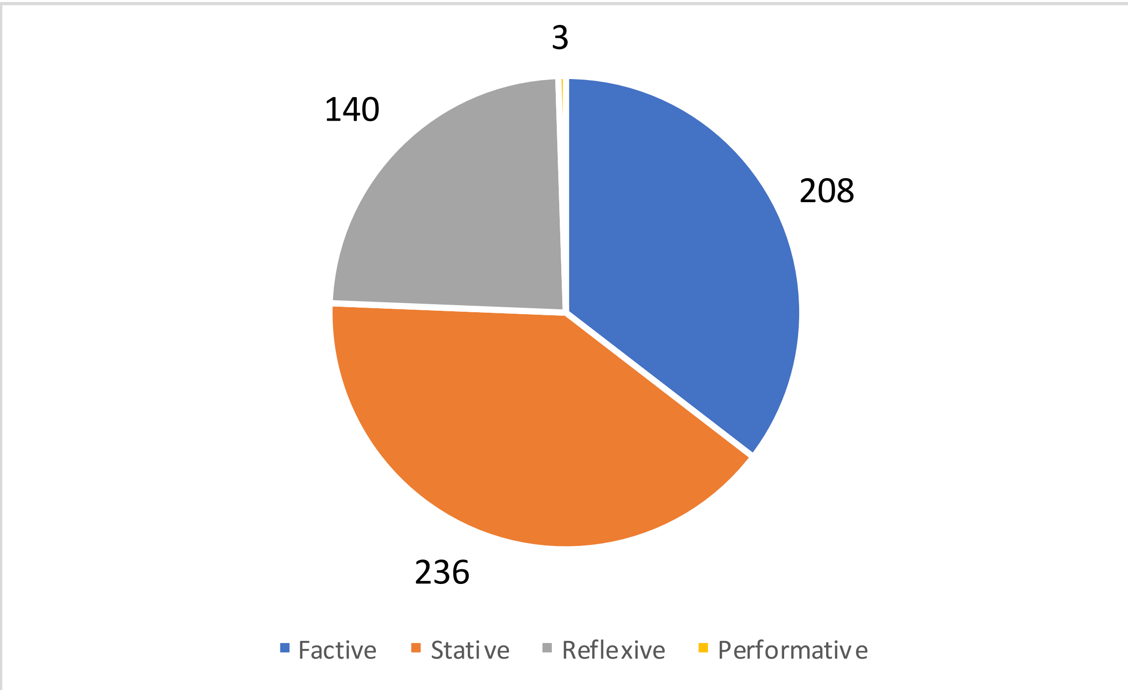

According to these figures, there is a similar proportion of each verb category in both indirect and direct threats. Both groups of texts feature very few examples of performative verbs. The Tropes definition of this category is highly similar to the performatives discussed in Chapter 2. These verbs then, express an action that is carried out upon uttering it aloud. For the indirect threats collected for this study, there are six instances of performative verbs identified by the software: 'respond,' 'hope (2),' 'bet,' 'want to' and 'ask.' For the direct threats, the verbs 'hope,' and 'want to (2)' were flagged up as performative. These findings suggest that authors did not opt to express actions committed solely through language use. As previously discussed, authors do not have access to performative verbs which are specific to communicating a threat (e.g. *'I threaten you'). Authors also did not appear to use other related words such as 'promise' or 'warn.' Instead, the (hypothetical) actions stated in these data were mostly either physical or conceptual.

'Factive' in this context refers to verbs which express physical, 'real-world' actions. Words such as 'leave,' 'scream,' 'mate,' and 'cut' are all examples of factive verbs recognised by Tropes in the data collected for this study. The direct threat data feature a higher frequency of factive verbs than indirect threats do. However, the percentage of verbs which were factive are highly similar across these two groups of texts. Also of interest here is whether the nature of these actions differs between indirect and direct expressions of threat. It was predicted that direct threats would be characterised by words which express violent or harmful actions. In contrast, indirect threats were expected to avoid these expressions. This issue will be more thoroughly addressed in Chapters 4.2 and 4.11 using the output of LIWC software analysis.

The frequency of stative verbs was also similar between the indirect and direct threats. These verbs express either states ('are,' 'will') or concepts relating to possession ('have'). As before, while direct threats feature a higher frequency of stative verbs than indirect threats, the percentage of total verbs which are stative are markedly similar for both groups of texts. Likewise, there are striking similarities between the percentage of verbs from these text groups which are identified by Tropes as reflexive. These verbs are defined by Tropes as verbs which express 'statements about facts, actions, people, objects, (and) feelings.' Reflexive verbs included in these data are 'like,' 'think,' 'know' and 'must.'

These findings suggest that authors overall, do not appear to include these verb types at differing frequencies based on whether a threat is directly or indirectly expressed. However, the nature of these verbs might vary. This topic will be explored later in this chapter.

Expressions of doubt

It was predicted that threats would overall be characterised by the author presenting their actions (actual or otherwise) as highly likely or certain to occur. As such, words which suggest doubt ought to be absent from these data, or at any rate highly infrequent. This is the case in both the indirect and direct threat data collected for this study. In the indirect threat data, two terms were found that relate to doubt: ‘perhaps’ and ‘pretty sure.’ In addition, both of these terms featured in Speaker 31’s indirect threat. For direct threats, three words were identified by Tropes as expressing doubt: ‘perhaps,’ ‘probably,’ and ‘maybe.’

The lack of data in this corpus relating to doubt seems to be in keeping with previous descriptions of, or predictions about threatening language. Also associated with the expression of doubt, is the use of modal auxiliary verbs relating to uncertainty. This topic will be discussed in more detail in Chapter 4.4.

4.1.1 Summary and discussion of Tropes software analysis

To summarise, indirect and direct threats appear to be characterised by similar argumentative text styles. As such, through these texts, the author involves him/herself into the narrative by arguing, explaining, or analysing something to convince the reader. However, there seem to be differences between these text groups regarding semantic themes. Direct threats, as predicted, feature more references to violence or harm than indirect threats (although this was not the case for all authors). Also in keeping with previous predictions was the lack of terms in these data which relate to doubt.

This subsection also addressed whether there were differences in the type of verbs used in indirect and direct threats, in particular whether these authors used verbs which expressed other related speech acts (i.e. verbs which were performative). Overall, verb usage did not seem to show much difference between these text groups. Performative verbs also appeared to be infrequent in these data. However, these verb categories will be explored in more detail in subsequent subsections. This is because there may be differences between indirect and direct expressions of threat and verbs which express negative or harmful actions.

4.2 Output of LIWC2015 software analysis

This subsection will detail the results of an automatic analysis performed by LIWC2015 software. This software reports the analysis of 42 categories. For this thesis, the output for the following categories will be reported:

- Word count
- LIWC dictionary recognition
- LIWC summary variables (Analytic/Clout/Authentic/Tone)
- Posemo and negemo
- Anxiety, anger, and sadness
- Certain and tentative
- Past, present, and future tense
- Death

The LIWC summary variables Analytic/Clout/Authentic/Tone were selected for discussion, as these are outputs which are proprietary to LIWC software. Each of these categories (and the motivation for their inclusion) will be explained in more detail in their corresponding subsections. The remaining categories listed above have been previously associated with threatening communications. With some exceptions, the majority of these outputs are reported as percentages relative to the remaining text. For example, if a text scored '2.33' for the 'Power' output, the software has detected that 2.33% of the text contains words which are associated with power. As such, texts with higher scores would have a larger proportion of words which are recognised as part of a given category relative to the remaining text.

Word count

Of interest to this study is whether the length of a text could indicate evidence of the organisation or planning of a harmful action. It is possible that direct threats which specify the location, timing and nature of harm intended would use more words to convey this, than messages which were vague (indirect threats). Longer messages, in general, have been associated with increased risk of harm (Spitzberg and Gawron, 2016). Figure 4 shows the word frequency for the indirect and direct threats collected for this study.

There is a large amount of variation between the lengths of these texts. The shortest document consisted of just 8 words (Speaker 16); the longest was 227 words (Speaker 31).

Figure 4: Word count frequency of all authors writing indirect and direct threats

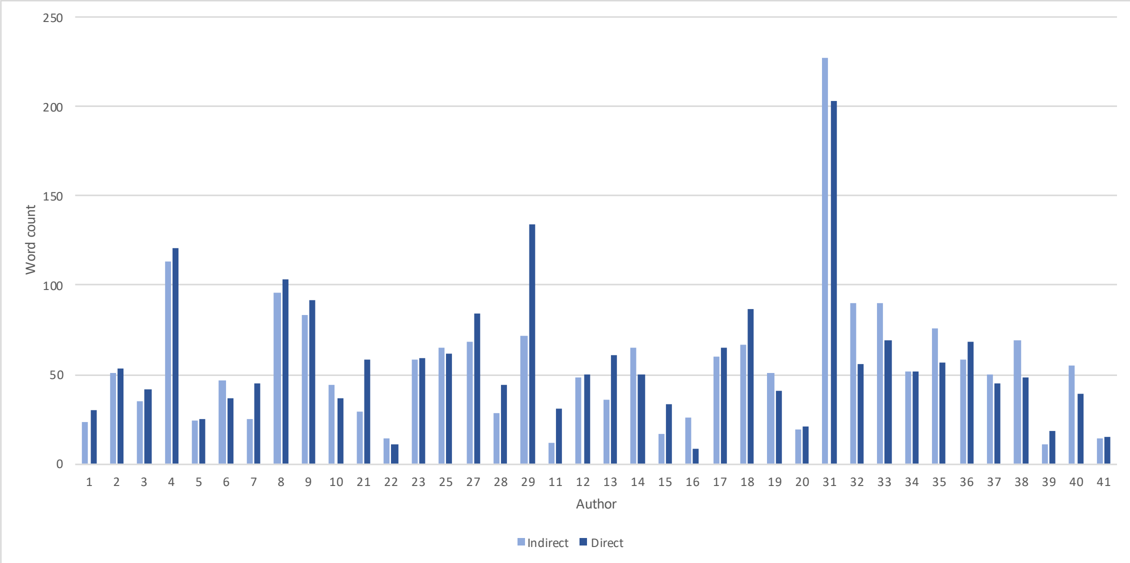

Overall, the average length of these texts was 55.55 words. However, even the longest texts collected here are not necessarily suitable for automatic approaches, such as those used in this study. It is likely that the reliability of any results derived from these methods will be strongly affected by a lack of linguistic data. That said, these data show that there is considerable variation in the amount of linguistic content available for analysis. The development of any approaches which aim to (for example) analyse, or automatically detect 'threatening language', should be mindful of this degree of variation between authors.

A Wilcoxon signed-rank test did not show a statistically significant difference between the word count of indirect and direct threats across authors (Z = -0.70, p = 0.49). The figure below (5) shows that instead, there is a strong positive correlation between the word counts of these documents ($r(38) = 0.89$, p = <0.05).

Figure 5: Correlation of word count frequency of all authors writing indirect and direct threats

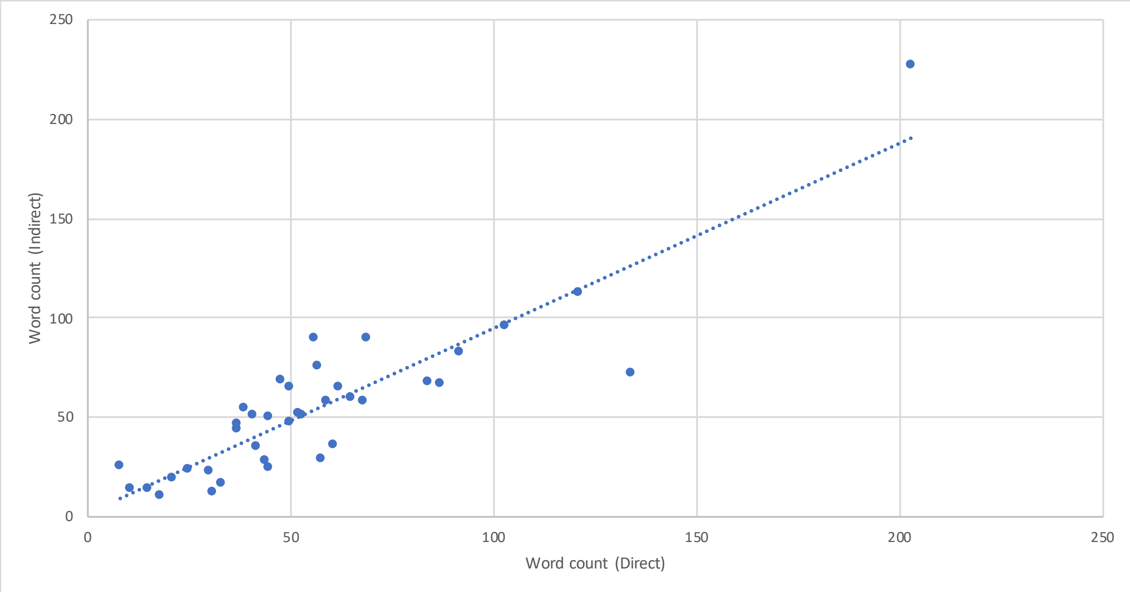

This suggests that the authors of these texts do not necessarily use more words to detail their threats, as might be expected for direct threats (in particular). Instead, word count seems to be author-specific, and not necessarily related to the expression of threat.

Dictionary recognition

As described in Chapter 3.5.2, LIWC analyses written texts based on an extensive, pre-built dictionary. For the data collected for this thesis, the software was able to recognise the majority of the written content. Put simply, most of the words in the corpus are also featured in LIWC2015's inbuilt dictionary. LIWC reports the percentage of a text which is recognised by the software under the category 'Dictionary (Dic).' On average, 90.18% (s.d. = 4.9%) of the words included in the analysis were recognised. Speaker 22's indirect threat contained the fewest number of recognised words (71.43%), and was among the shortest texts collected for this study. Speaker 20's direct threat consisted entirely of words which were recognised by the software.

This finding boosts the reliability of the remaining LIWC analysis. However, there are a range of words found in these data which are not recognised by LIWC. This includes frequently-occurring terms such as *farm, animal(s)* and *rights*. There are also words which convey violent acts which are not included in this analysis. These include words such as: *guns, bombs, stab, explosion, gouge, shoot(ing)* and *shotgun*. As such, any inferences that

LIWC generated about each text were not based on these terms. It is plausible that human readers might associate such words with expressions of aggression or death. It is also possible that LIWC analysis of related themes (such as death) would also be affected.

LIWC summary variables

Analytic

This output provides information about an author's use of either logical or more narrative thinking patterns. Higher scores (above 50) are said to suggest more formal thinking processes. Lower scores (below 50) are said to suggest more 'here and now' thinking, or thinking based on personal experience rather than logic. There are no predictions made in this study about the level of analytic thinking present in threatening language. However, it might be expected that threats would show evidence of a more formal, planning process, detailing the author's motive and the actions he/she 'intends' to take.

For the data collected in this study, the mean analytic scores for all indirect threats is 25.58 (s.d.= 26.37). Direct threats scored similarly, with an average score of 19.41 (s.d.= 13.44). These scores suggest that the authors were expressing thought processes that were less formal, and more centred on the author's personal experiences. That said, there was a fairly wide range of scores across authors - particularly for indirect threats. A Wilcoxon signed-rank test did not show a statistically significant difference between the analytic scores for either indirect or direct threats ($Z = -0.66$, $p = .49$).

Figure 6 shows the variation of analytic scores for all of the simulated indirect and direct threats collected for this study.

Figure 6: Analytic scores of all indirect (light blue) and corresponding direct threat data (blue)

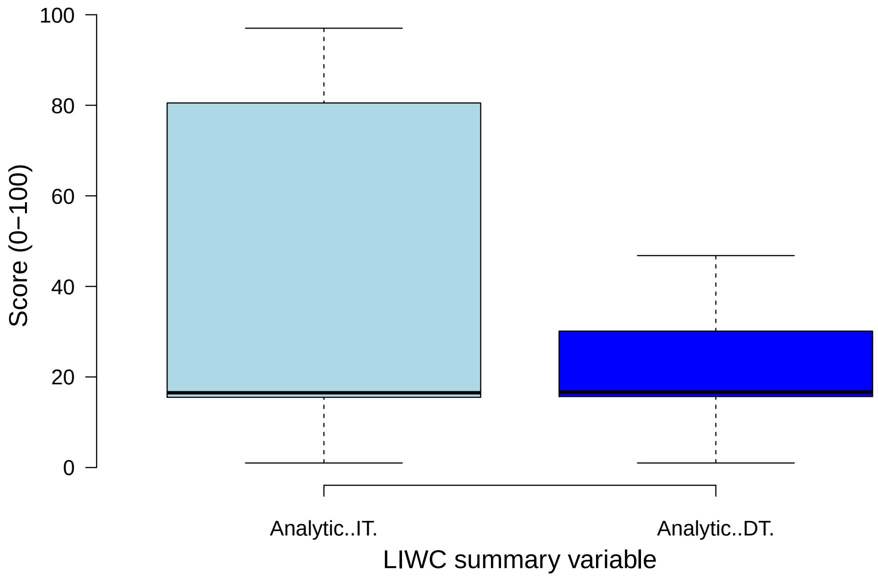

Clout

This output calculates a score regarding an author's social status, confidence and/or leadership based on the linguistic content of their text(s). As before, higher scores (50-100) indicate that the author used language which was characterised as conveying a relatively high social status, confidence, and/or leadership qualities. Lower scores would suggest the opposite. No predictions are made in this study concerning the degree of clout that these texts convey. However, it might be advantageous for an author to portray him/herself as confident - and perhaps as such more competent - when writing a threat.

The indirect threats collected for this study scored 88.18 (s.d. = 14.75) on average. The corresponding direct threats scored 84.21 (s.d. = 23.56) on average. A Wilcoxon signed-rank test did not show a statistically significant difference between the clout scores for indirect and direct threats (Z = -0.25, p = 0.80). These results suggest that overall, the authors wrote threatening texts which could be associated with language use which conveys confidence, leadership, and/or higher social status. In addition, indirect and direct threats could not be distinguished from one another based on this metric alone.

Figure 7 shows the spread of clout scores for all of the simulated indirect and direct threats collected for this study.

Figure 7: Clout scores of all indirect (light blue) and corresponding direct threat data (blue)

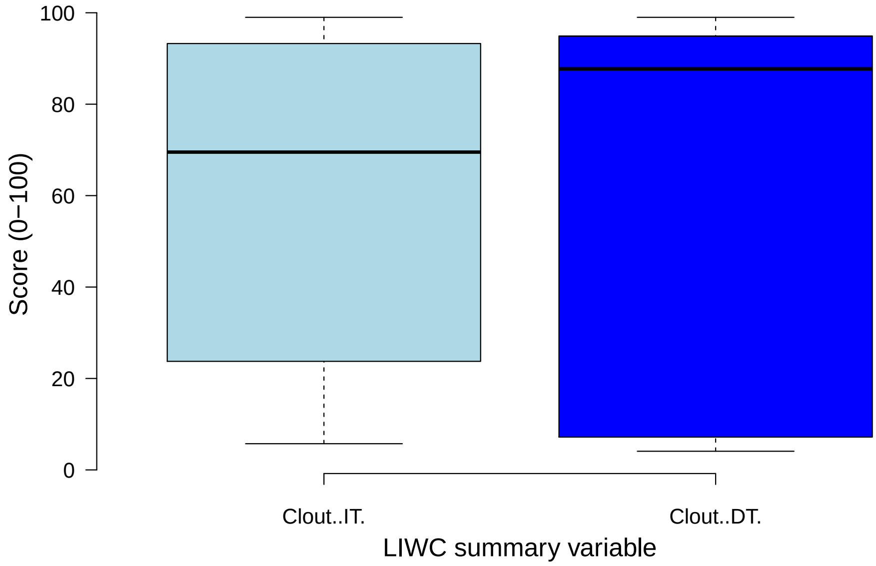

Authentic

'Authentic' here refers to texts which convey honesty, sincerity, and/or vulnerability. Texts which score highly (50-100), are associated with stronger degrees of these qualities. Once again, there are predictions made in the previous literature about how threatening language might score using this metric. It should be emphasised, however, that the texts collected for this study are essentially bluffs. The authors are instructed to threaten a target that does not exist, and he/she does not have the means to carry out many of the harmful acts described in his/her texts. Put simply, while these texts might convey a sincere *intention* to cause someone to fear harm, they do not represent actions which are grounded in reality. As such, it might be expected that these texts will score low for this output.

For the indirect threats collected for this thesis, the mean authentic score (overall) is 51.51 (s.d. = 30.16). The corresponding direct threats scored 57.00 (s.d. = 29.38) on average. A Wilcoxon signed-rank test did not show a statistically significant difference between the authentic scores for indirect and direct threats (Z = -0.84, p = 0.40). This suggests that overall, these authors wrote threatening texts which were slightly associated with expressions of authenticity. However, as seen in Figure 8, these scores varied considerably across the authors.

Figure 8: Authentic scores of all indirect (light blue) and corresponding direct threat data (blue)

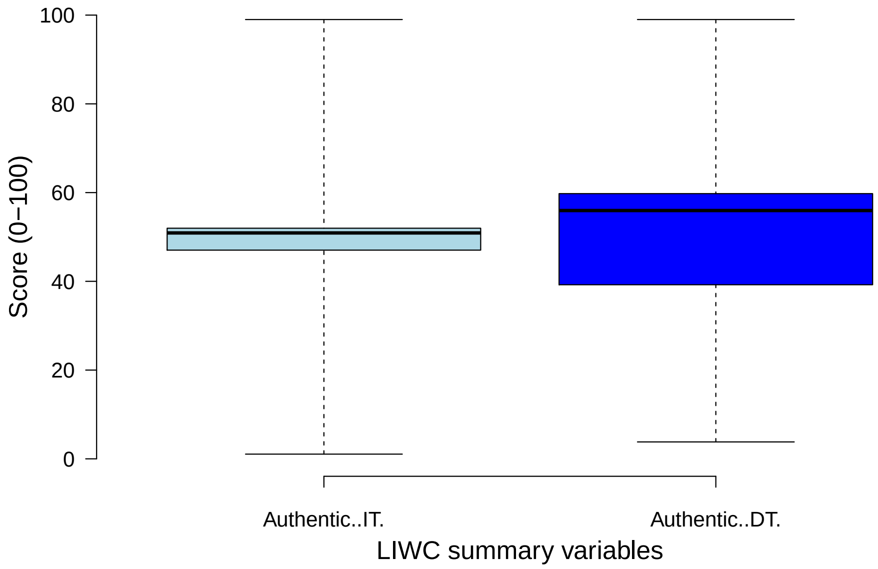

Tone

This output combines the scores reflecting the 'positive' and 'negative' emotional tone of the texts. The individual scores will be presented later in this subsection. Higher scores (between 50-100) are associated with more positive emotional tones, while lower scores (between 0-50) are linked with more negative tones. There are no predictions available specifically about the emotional tone of threatening language. It might be anticipated that more direct expressions of threat might be more associated with negative emotions, while indirect threats would remain somewhat emotionally neutral, or more variable.

For the indirect threats collected for this thesis, the mean tone score is 20.99 (s.d. = 24.93). For the corresponding direct threats, the mean score is 17.00 (s.d.= 28.70). Statistical testing did not show a significant difference between the tone scores for indirect and direct threats across the board (Z = -0.72, p = 0.47). These results indicate that these threatening texts as a whole were associated with a more negative tone. However, a number of threatening texts scored '1.00' for this metric (n=31). This suggests that these texts possibly lacked the linguistic content needed to calculate a more robust score. While all of the outputs from LIWC analysis should be interpreted cautiously due to the length of the text samples provided, this output appears to be particularly problematic.

Figure 9 shows the spread of tone scores for all of the simulated indirect and direct threats collected for this study.

Figure 9: Boxplots for the tone scores of all indirect (light blue) and corresponding direct threat data (blue)

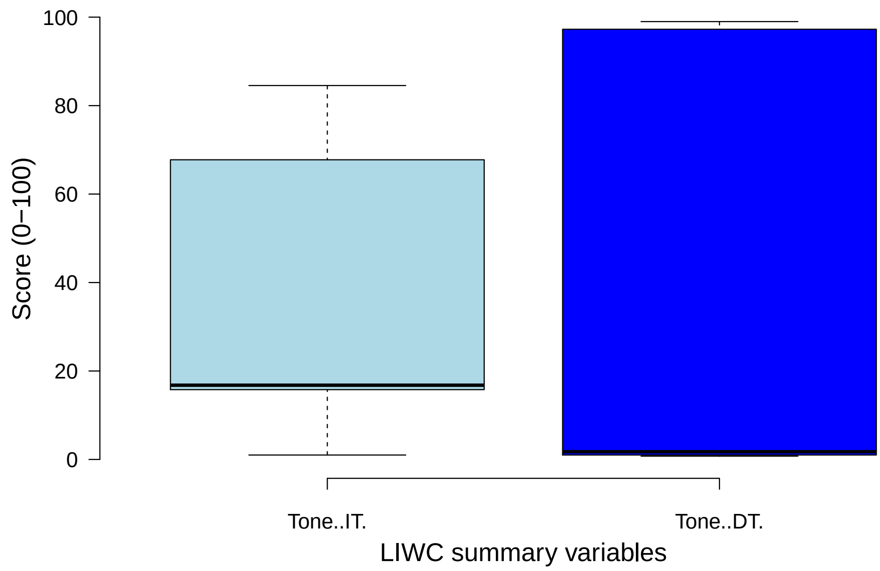

Summary of LIWC summary variables analysis

The figure below (10) visualises the results of a principal component analysis which combines the results of these LIWC summary variables. The red dots represent each author's direct threat text (DT). The blue dots represent each author's indirect threat text (IT). There is considerable overlap between the two text groups. This again suggests that overall, the indirect and direct threats collected for this thesis cannot be differentiated from one another based on these summary variables alone. However, as shown in the axes of this figure, these LIWC summary variables were able to explain a decent proportion of the variance in these texts.

Figure 10: Principal component analysis of the LIWC summary variables for indirect and direct threats

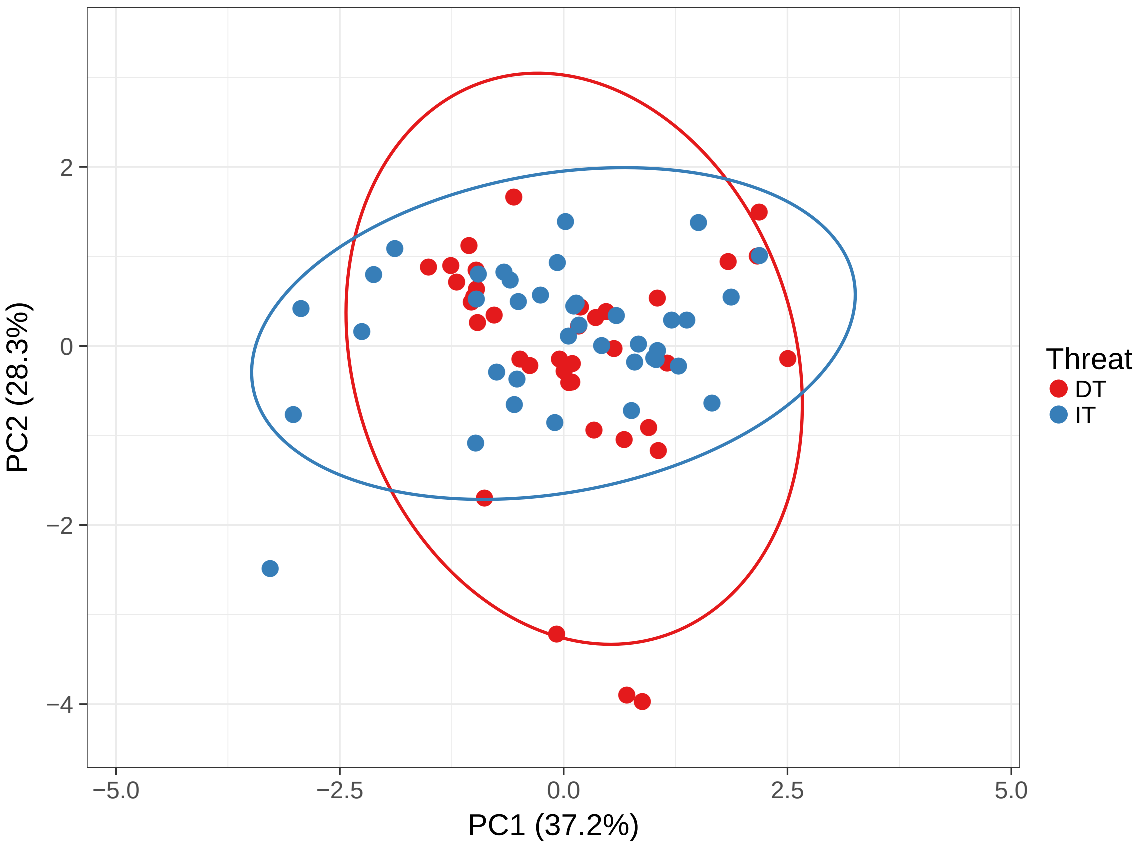

The results of the LIWC summary variable analysis can be summarised as follows:

- None of the summary variable outputs were statistically different between indirect and direct threats.
- The texts overall scored relatively low for analytic thinking and tone, relatively high

for clout, and moderately high for authenticity. However, it is not known how these scores relate to other, non-threatening texts written by these authors.

Posemo and negemo

The following LIWC outputs calculate scores based on the presence (or absence) of words relating to positive or negative emotions ('posemo' and 'negemo,' respectively). A high score for 'posemo' (50-100) would suggest that the text included a high percentage of words which can be associated with positive emotions. Likewise, for 'negemo,' a high score would indicate that the text contained a high percentage of words relating to negative emotions. These outputs differ from the LIWC summary variable 'tone' described previously - where the score reflected whether a text (overall) could be described as more positive or more negative in emotional tone. For the posemo and negemo outputs, the score reflects the degree to which a text is associated with either positive or negative emotional tones.

There are no predictions about the emotional connotations of words used in threatening texts. That said, it might be assumed that the threatening texts would include a higher percentage of negative emotion words than positive emotion words. This effect might be more pronounced in direct threats, as authors can explicitly refer to harmful actions.

For the threats collected for this study, the posemo and negemo average scores can be summarised in table 11 as follows:

Table 11: Posemo & negemo scores

	Indirect threat	**Direct threat**
Posemo	2.08 (s.d. = 2.17)	2.58 (s.d. = 2.55)
Negemo	4.97 (s.d. = 3.86)	5.41 (s.d. = 3.30)

While these texts appear to mostly score more highly for negative emotions than positive emotions, the scores are rather low. A Wilcoxon signed-rank test showed no statistical difference between indirect and direct threats regarding the percentage of words which convey either positive ($Z = -0.80$, $p = 0.42$) or negative ($Z = -0.47$, $p = 0.64$) emotions. Words which LIWC identified as positive from these texts include: 'dear,' 'safe,' 'care' and 'better'. Negative words included: 'ignored,' 'warning,' 'fucking,' and 'kill'.

Anxiety, anger, and sadness

These outputs are again percentages which are generated based on the frequency of words associated with anxiety, anger and sadness, relative to the remaining text. For threatening language, there are no predictions about the use of words which convey these emotional or behavioural states. However, it might be expected that threatening language would contain a higher use of words which convey anger than is the case for non-threatening language. As this study did not collect any examples of non-threatening texts, this issue cannot be addressed directly.

The following table (12) shows the average scores for anxiety, anger and sadness as generated by LIWC.

Table 12: Average LIWC scores for for anxiety, anger and sadness

	Indirect threat	**Direct threat**
Anxiety	0.22 (s.d. = 0.64)	0.25 (s.d. = 0.74)
Anger	1.85 (s.d. = 2.34)	3.21 (s.d. = 3.31)
Sadness	1.03 (s.d. = 2.13)	0.43 (s.d. = 0.83)

Overall, it appears that these texts contain higher percentages of words associated with anger than anxiety or sadness. As with the posemo/negemo results, however, these scores are low. A Wilcoxon signed-rank test showed that there is a statistically significant difference between the percentage of 'anger' words detected by LIWC in indirect and direct threats (Z= -2.04, p = 0.041). For the percentage of anxiety and sadness words between these text groups, there is no statistical difference (Z = -0.25, p = .80, and Z = -1.20, p = 0.23 respectively).

Certain and Tentative

These LIWC outputs again represent the percentage of a text which consisted of words which relate to a particular concept - in this instance, certainty and tentativeness. It is expected that threatening language will, overall, be characterised by the use of words which express certainty. In addition, direct threats overall are expected to contain a higher percentage of words which express certainty, relative to indirect threats. This is because authors might choose not only to be clearer about their target and any harm which might befall him/her, but also about the likelihood that this harm will be realised.

For the threats collected for this study overall, LIWC detected a higher percentage of words associated with tentativeness than certainty. These results are displayed in the table (13).

Table 13: Average LIWC scores for tentative & certain

	Indirect threat	**Direct threat**
Certain	1.16 (s.d. = 1.58)	1.26 (s.d. = 1.49)
Tentative	3.75 (s.d. = 2.88)	2.94 (s.d. = 2.81)

For many of the texts (n=41), there were no words detected by LIWC which were associated with certainty. Such words included 'all,' 'every,' 'sure,' and 'nothing.' Notably, this output does not include auxiliary verbs which can denote certainty or likelihood. Words associated with tentativeness were overall more frequent; these words included: 'something,' 'seem(s), 'trying,' and 'or.' It is possible that the relatively high frequency of 'or' usage (used in conditional statements) might contribute to the greater percentages of words in these texts associated with tentativeness. There is no statistically significant difference between the percentages of indirect and direct threats for certainty (Z = -0.47, p = 0.64) or tentativeness (Z = -1.57, p = 0.12).

Power

This output concerns the authors' 'need for power', and in particular, awareness of his/her social status as conveyed through his/her writing. This category is distinct from the previously described LIWC summary variable 'clout.' There are no predictions about the expression of power in threatening language. This category was selected for discussion here because authors of threatening texts might attempt to convey that they have greater social status than his/her target (or him/herself in other situational contexts). The results presented in this subsection cannot shed any light on how authors express social status in non-threatening texts. The purpose of this analysis is to consider the use of 'power' as a variable in further research on threatening language.

For indirect threats, the average power score is 4.28 (s.d. = 2.79). For the corresponding direct threats, the average score is 3.09 (s.d. = 2.43). There is a statistically significant difference between indirect and direct threats regarding power (Z = -1.681, p = 0.04648). Overall, these results indicate that authors write indirect threats using a higher percentage of words which are associated with power (compared to direct threats). These words included: 'force,' 'strongly,' 'greed,' and 'threat.'

Past, present, and future tense

These outputs show the percentage of words in a text which is associated with the expression of past, present and future tense. In LIWC, these groups are called 'focuspast,' 'focuspresent,' and 'focusfuture.' It has previously been shown in related research that threatening texts are primarily written using the present tense (Chandler, 2017).

The following table (14) shows the overall average LIWC outputs for each of these three categories for indirect and direct threats:

Table 14: LIWC output for tense analysis

	Indirect threat	**Direct threat**
Past tense	1.88 (s.d. = 2.41)	2.17 (s.d. = 2.42)
Present tense	15.25 (s.d. = 5.30)	14.32 (s.d. = 5.21)
Future tense	3.63 (s.d. = 3.50)	5.08 (s.d. = 2.91)

These results seem to support previous findings, asserting that threatening language largely focuses on present events. A Wilcoxon signed-rank test did not show a statistically significant difference between the frequency of words associated with past (Z = -0.68, p = 0.50) or present (Z = -1.09, p = 0.28) tense in indirect and direct threats. There is a significant difference between these text groups, however, for the frequency of words conveying the future tense (Z = -2.03, p = 0.04). The averages displayed in the table suggest that direct threatening texts consist of a greater frequency of future tense words, though the degree and direction of this difference is not consistent across the authors.

Death

This output represents the percentage of words in a document that are associated with death. It might be expected that threats would feature a higher percentage of these words as compared to non-threatening texts. For this study, whether there is a statistical difference between the death scores of indirect and direct threats cannot be calculated. This is because most of the scores returned by LIWC were zero, mainly for indirect threats.

For indirect threats, the average death score overall is 0.04 (s.d. = 0.20). Only two of these texts featured any words relating to death (written by Speaker 9 and Speaker 31).

Both of these texts include the word ‘die.’ For direct threats the average death score is 0.50 (s.d. = 1.21). There are seven texts for which LIWC returned any score above zero: Speakers 3, 4, 11, 12, 17, 28, and 36. These scores are based on the following words: ‘corpse,’ ‘kill (n=4),’ ‘death (n=2),’ ‘murder,’ and ‘alive.’ On the surface, these results seem to suggest that authors seldom use words which can be associated with death. As discussed earlier, these results do not include words such as ‘shoot’ or ‘gun’ which might be associated with death by human readers. It is also possible that these authors chose to avoid making overt death threats, and instead threatened some other form of harm. This will be addressed later in §4.11.

4.2.1 Summary and discussion of LIWC2015 analysis

This analysis examined a variety of LIWC outputs which are pertinent to threatening language, in particular, whether there are differences between the indirect and direct threatening texts collected for this study. For the most part, these text groups appeared to be quite similar to one another regarding these variables. No statistically significant differences were found between indirect and direct threats for the following LIWC categories:

- Word count
- LIWC summary variables (analytic, clout, authentic, tone)
- Posemo/negemo
- Anxiety and sadness
- Certainty and tentative
- Past and present tense

Previous research on the threatening messages written by animal rights extremists has also examined and reported the outputs relating to these LIWC variables (Chandler, 2017). Compared to this previous research, the data collected for the current study has notably lower percentages of the LIWC summary variables. These data are presented in Table 15.

Table 15: LIWC summary variable output for direct and indirect threats compared to Chandler (2017)

	Current study (Indirect threats)%	**Current study (Direct threats) %**	**Data examined by Chandler (2017) %**
Analytic	25.58	19.41	46.71
Clout	88.18	84.21	96.24
Authentic	51.51	57.00	38.23
Tone	20.99	17.00	3.83

The data from these studies share high percentages for 'Clout'. Therefore, there is evidence that threatening communications of both authentic and simulated data are written using linguistic features which suggest confidence or leadership. Also, in Chandler's data the texts have a notably lower score for 'Tone', suggesting that authentic threats possess linguistic features which are associated with negative emotions. As described earlier in this research, it is not known how these LIWC variables are calculated. As such, it is difficult to assess which linguistic features from these data could be contributing to this low 'Tone' score.

For LIWC outputs relating to affect, there appear to be similarities between the data from these two studies. Table 16 shows the outputs for the categories posemo, negemo, anxiety, anger and sadness.

Table 16: LIWC affect outputs for direct and indirect threats compared to Chandler (2017)

	Current study (Indirect threats)%	**Current study (Direct threats) %**	**Data examined by Chandler (2017) %**
Posemo	2.08	2.58	1.92
Negemo	4.97	5.41	4.17
Anxiety	0.22	0.25	0.24
Anger	1.85	3.21	2.21
Sadness	1.03	0.43	0.43

For the data from these studies, each of these affect outputs are highly similar. This similarity could be taken as evidence that authentic and simulated threats relating to animal rights extremism are highly similar. The following table (17) shows the remaining LIWC outputs recorded for these two studies. These outputs relate to the expression of time in these groups of texts.

Table 17: Remaining LIWC outputs for direct and indirect threats compared to Chandler (2017)

	Current study (Indirect threats) %	**Current study (Direct threats) %**	**Data examined by Chandler (2017) %**
Past tense	1.88	2.17	1.30
Present tense	15.25	14.32	12.02
Future tense	3.63	5.08	4.92

The outputs of these studies again seem to be very similar in terms of how time is expressed. Notably, there is a tendency for the authentic and simulated data to primarily express the present tense. It should be stressed, however, that these studies did not examine any 'non-threatening' texts. Therefore, it is difficult to assess how the outputs shown in Tables 12-17 relate to non-threatening texts as written by the same authors.

The following LIWC categories were found to be significantly different between the indirect and direct threats collected for the current research:

- Anger
- Power
- Future tense

It was also observed that the majority of the words were recognised by the software. However, there were also words which were not included in the software's dictionary. As such, these words could not contribute to the scores provided in this subsection. Other limitations of this approach have also been considered. These included the comparison of percentages of word categories and the lack of words available in these data to enable a more robust analysis.

The findings of this chapter suggested that these authors possibly distinguished his/her indirect and direct threats using language which conveys power, anger and future tense. Future research could explore this phenomenon in more detail. In particular, how laypeople and threat assessors understand and/or use these categories to distinguish between indirect and direct threats. Also of interest, is how these findings relate to threatening and non-threatening texts that are written by the same author. A similar analysis to that presented in this subsection could be run on such data. In turn, this would provide the basis for research which assesses the potential risks or usefulness of these LIWC categories for the development of automated threat detection software.

The subsequent subsections provide a more qualitative analysis of a number of factors which have been associated by researchers such as Smith (2006) and Simons and Tunkel (2014) with an increased or decreased risk that violence or harm will be realised. These include the use of pronouns, modal auxiliaries, conditional statements, profanity, and performative verbs.

4.3 Pronouns

This subsection will present the frequency of pronouns used across the indirect and direct threats collected for this study. In previous research, it has been suggested that pronouns can be used to infer whether the author focuses more on him-/herself or on the target of a threat (Smith, 2006). Of particular interest here is whether these authors write threats which centre on their own actions or on the actions of their target. For example, frequent usage of the first person pronoun 'I' would indicate that a text centres on the actions to be committed by the author. On the other hand, 'me' functions as the object of a verb. This phrasing might suggest that the author is not the subject that is committing the action stated in the verb.

Also of interest is whether there are differences in pronoun usage between indirect and direct threats. It is possible that participants interpreted the word 'indirect' to mean 'vague' or 'non-specific.' In other words, 'direct/indirect' might also be inferred by authors to mean *directing* the actions specified in a threat towards a particular target. This could serve not just to make the actions specified in indirect threats less clear, but also to express uncertainty concerning who will carry out the action or who is the recipient of an action. As such, direct threats might, overall, feature more phrases which centre on the author (and their actions towards the target) than do indirect threats.

There also appear to be no published predictions or data concerning the use of indefinite, possessive or reflexive pronouns in threatening communications. It stands to reason that authors of threatening texts might use fewer indefinite pronouns (e.g. 'somebody') if they are *directing* their threats to a particular target. Also, the use of first person possessive or reflexive pronouns could also be used to articulate that the author is accepting responsibility for the actions specified in the sentence (e.g. 'You will suffer my wrath').

For this research, the indirect and direct threats were tagged manually (using the highlight function of Microsoft Word software) for pronouns belonging to the following groups:

First person personal:

- singularly as the subject of a verb (I) or as an object of a verb (me)
- plurally as the subject of a verb (we) or as an object of a verb (us).
- singularly as a possessive (my and mine) or as a plural (our, ours)
- singularly as a reflexive (myself) or as a plural (ourselves)

Second person personal:

- singularly or as a plural as the subject of a verb (you)
- singularly or as a plural as a possessive (your, yours)
- singularly or as a plural as a reflexive (yourself, yourselves)

Third person personal:

- singularly as the subject of a verb (he, she, it, they) or as an object of a verb (him, her, it, them)
- as a plural as the subject of a verb (they) or as an object of a verb (them).
- singularly as a possessive (his, her(s), its, their(s)) or as a plural (our, ours)
- singularly as a reflexive (himself, herself, itself, themselves) or as a plural (themselves)

In addition, the following indefinite pronouns (which refer to a person or people) were tagged:

- someone/somebody
- everyone/everybody
- anyone/anybody
- no one/nobody

In total, the indirect threat data contained 344 pronouns. By comparison, the direct threat data contained 450 pronouns. The following sections show the overall frequency of words corresponding to these pronoun groups across the indirect and direct threat data.

First person pronouns

Figure 11 shows the first person pronouns found in these data.

Figure 11: Frequency of first person pronouns in indirect and direct threats

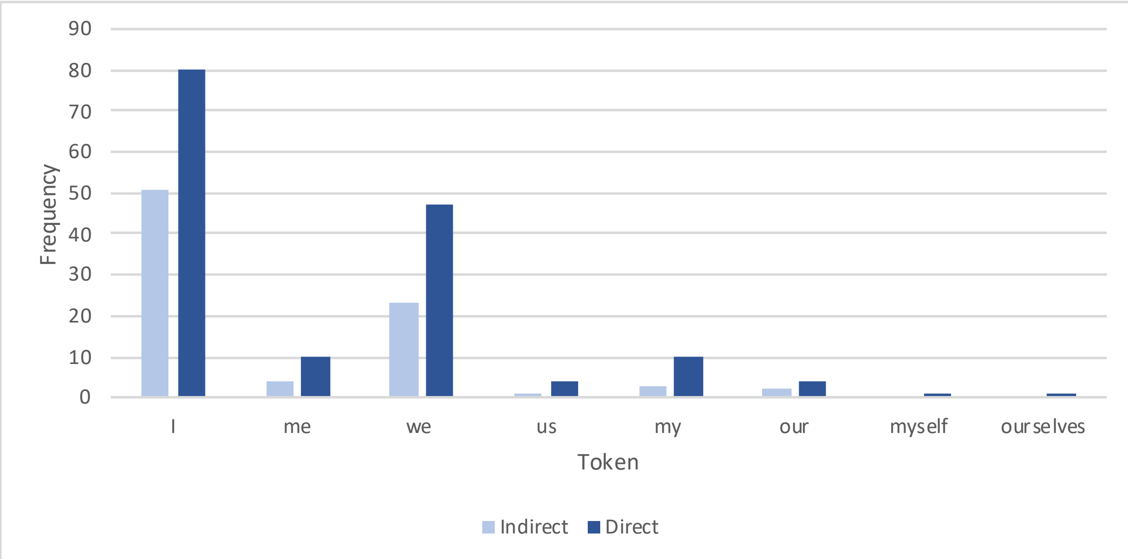

Overall, direct threats used a higher frequency of first person pronouns than indirect threats. This is perhaps unsurprising given that direct threats typically featured more pronouns. The pronoun 'I' is the most commonly-occurring across these data. The higher frequency of 'I', along with the usage of 'me,' suggests that these authors focused more on describing actions that they are performing or would (hypothetically) realise. This is opposed to referring to the authors referring to him/herself as an object of a verb. However, the relatively high occurrence of 'I' in these data is also seen in other, far more substantial, non-threatening English-language corpora. This will be discussed in greater detail later in this subsection.

Another observation regarding these data is the fairly high occurrence of 'we'. This pronoun appears to refer to (hypothetical) members of the 'Stop Wests Farm' campaign group alluded to in the task briefing, or to other individuals who are unhappy with the target's behaviour. This can be seen in the following examples:

- *'Right now we're working our way into your factory to blow it up.'* (Speaker 21's direct threat)
- *'I am not one man, Mr West and I swear if you don't stand down, we will blow you and your filthy little farm apart.'* (Speaker 25's direct threat)

While participants were not instructed to assume the role of a campaign member, some opted to do so. The exact reason for this is not known. However, it can be speculated that these authors obliquely refer to multiple people in their threats to convey that the author is not the only person who is unhappy with the target's behaviour, or the only person who can bring about a harmful action. By doing so, the target might feel at risk from multiple potential sources of harm as opposed to 'just' one. However, contrary to this theory, Speaker 4 appears to have crossed out tokens of the word 'we' in favour of 'I' in his direct threat. This can be seen in his script in the Appendix. For now, it is also not known whether readers would be likely to infer changes to perceived risk based on the use of plural pronouns.

Finally, these data also exhibited very few uses of 'me,' 'us,' 'my,' 'our,' 'myself' and 'ourselves.' In addition, there are no tokens of 'mine' or 'ours' in these data. The lack of these words suggests that authors tended to avoid alluding to their own performance of an action.

Second person pronouns

Figure 12 shows the frequency of second person pronouns across the indirect and direct threats collected for this research.

Figure 12: Frequency of second person pronouns in indirect and direct threats

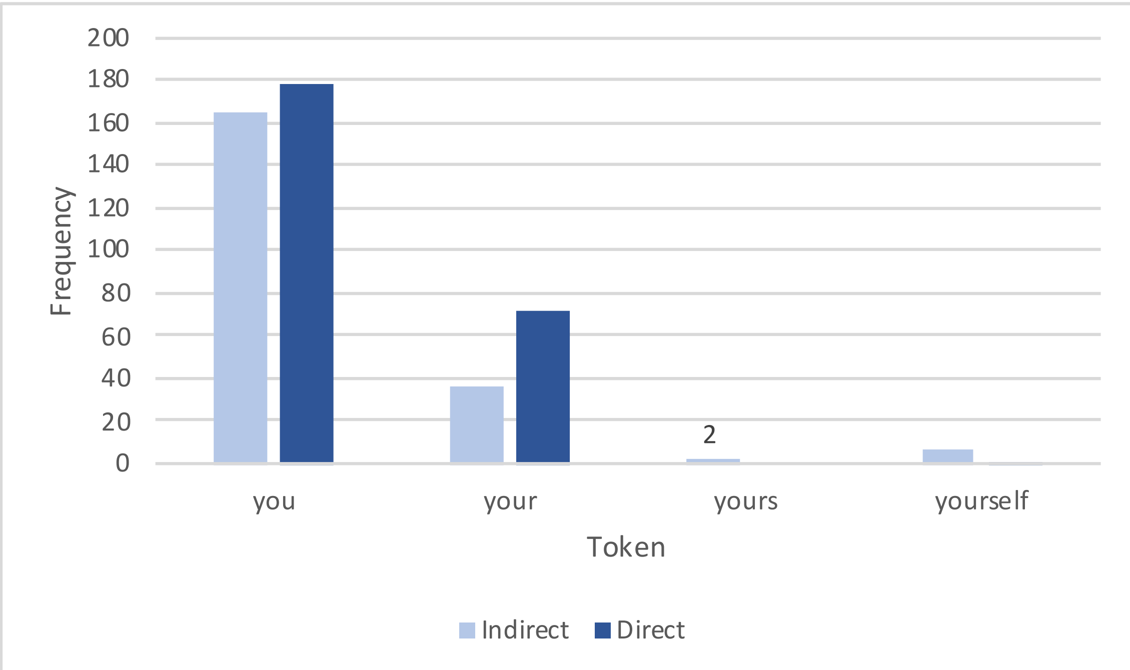

Based on this figure, 'you' is by far the most frequently featured second person pronoun in these data, followed by 'your.' These pronouns also occur slightly more frequently in direct threats than indirect threats. As before, these observations can be linked to the overall higher frequency of pronouns in the direct threat data. The prevalence of 'you' in these data, however, is not exclusive to this corpus of English-language threatening texts. This observation will be discussed in greater detail later in this subsection.

The higher rate of 'you' relative to other pronouns can also be attributed to its use as either the subject or object of a verb in a sentence. The following figure (13) shows the frequency of 'you' across these data as a subject or an object of a verb.

Figure 13: Frequency of 'you' functioning as the subject or object of a verb in indirect and direct threats

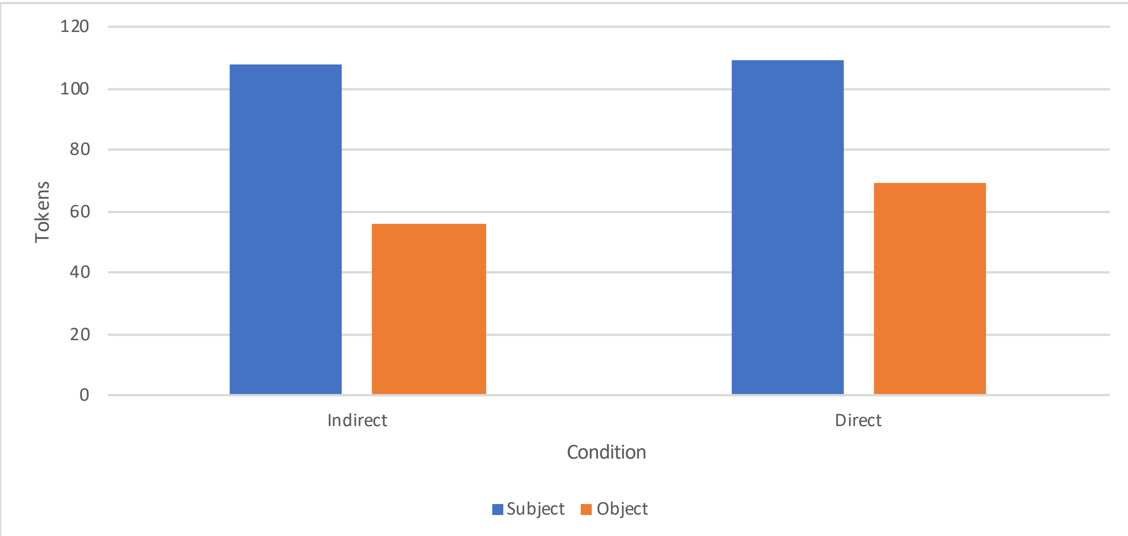

Based on this figure, the majority of 'you' tokens in these data function as the subject of a verb in a sentence. This usage can be seen in the following examples:

- *'I know where you live...'* (Speaker 4's direct threat)
- *'You have two days to close down the farm...'* (Speaker 13's direct threat)

This tendency for 'you' to act as a subject of a verb, combined with the lower rates of 'I/me', indicates that authors tend to focus more on describing the actions of their target than the actions performed by themselves. In addition, there appears to be little difference between indirect and direct threats in terms of whether 'you' functions as a subject or an object of a verb.

Third person pronouns

Figure 14 shows the frequency of third person pronouns found in the indirect and direct threats collected for this research.

Figure 14: Frequency of third person pronouns in indirect and direct threats

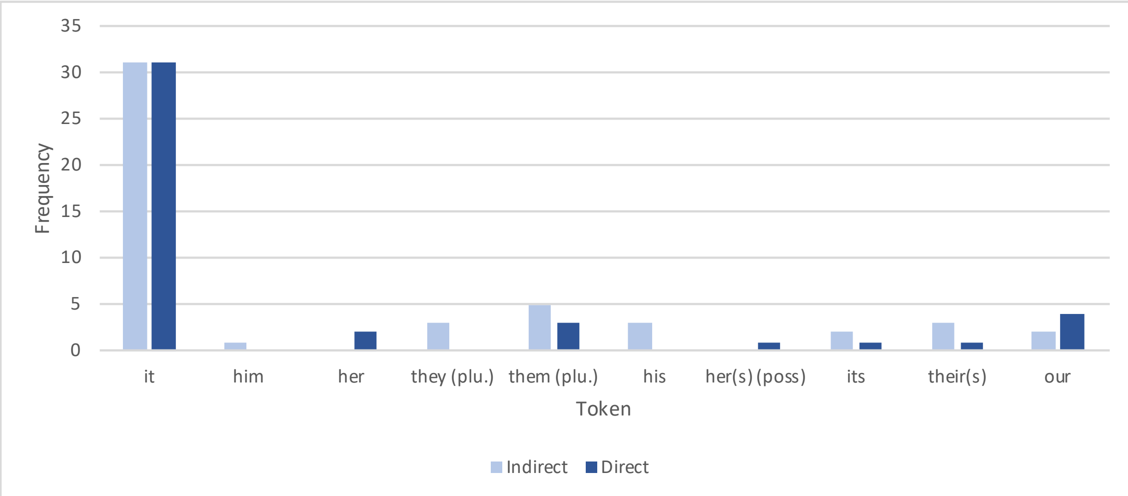

There are comparatively few third person pronouns across these data relative to first and second person pronouns. This might be because the threats were written to address a specific person who is not known to the threatener. As such, it is unlikely that the target will be referred to as the subject of an action using singular third person pronouns ('he/she/it/they'). In addition, using plural third person pronouns ('they/them') would not be viable to refer to this target. That said, 'they/them' is used several times in these data to refer to non-specified others who are unhappy with the target's behaviour.

Another observation is the low frequency or absence of possessive and reflexive third person pronouns in these data. The only tokens of possessive 'his' were found in Speaker 4's indirect threat: *'I would like to remind Mr West that I and other like-minded, compassionate individuals have more than just his phone number, and if necessary would be ready to move our protest to his home and his family.'* In this example, 'his' refers to possessions belonging to the target of the threat. In contrast, 'their' referred to the animals based at the farm (*'Why don't you think about the animals. Their excruciating pain, living in their own filth...'*(Speaker 38's indirect threat)). 'Their' also referred to the imaginary 'other people' who disagree with the target's behaviour (*'You should know that there are a lot of people who are very angry about this, and aren't afraid to take matters into their own* *hands.'* (Speaker 34's indirect threat)).

The most frequently occurring third person pronoun in these data is 'it.' This word occurs 31 times each in both the indirect and direct threat data. The fairly high usage of 'it' in these texts can be attributed to the number of entities that can be referred to using this pronoun. In these data, 'it' is used to refer to:

- the farm (Speaker 6's indirect threat: *'It is unethical beyond means...'*)
- the previous indirect threat (Speaker 13's direct threat: *'I warned you but you chose to ignore it'*)
- the previous actions/behaviour of the target (Speaker 33's indirect threat: *'It's inhumane, cruel and unnecessary, there are various other ways to undergo medical research without testing on animals'*)
- a harmful action directed toward the target (Speaker 17's indirect threat: *'And if you manage to survive it then I can assure you'*)

Despite the relatively high frequency of 'it' in these data (especially in relation to other third person pronouns), this finding is consistent with corpora documenting English-language speech or texts. These data will be discussed in relation to other non-threatening language corpora in the following section.

Indefinite pronouns

Overall, there are very few tokens of the indefinite pronouns listed earlier in this subsection in the data collected for this research. There is one token of 'someone' (Speaker 3's indirect threat), one token of 'everybody' (Speaker 33's direct threat), and two tokens of 'anyone' (seen in both Speaker 4's direct threat and Speaker 20's indirect threat). There are nine tokens of 'everyone' present in these data (seen in Speaker 13's direct threat, Speaker 31's indirect threat, and Speaker 38's indirect threat). However, this number is bolstered by Speaker 31, who included six instances of 'everyone' in her indirect threat. There are no tokens in these data of 'somebody,' 'anybody,' 'no one,' or 'nobody.'

4.3.1 Summary and discussion of pronoun analysis

This subsection presented the frequency of pronouns which occurred in the indirect and direct threat data collected for this research. While direct threats collectively feature more pronouns than indirect threats, there appears to be little difference between these two conditions relating to how pronouns are distributed. This finding indicates that these authors overall did not distinguish indirect and direct threats through pronoun usage.

In addition, this subsection showed that the pronouns 'I', 'you' and 'it' were the most commonly occurring across these data. The extensive use of 'you', in particular, appears to be in keeping with previous research relating to threatening texts (Gales, 2010). Threatening language which describes violent behaviour, deadlines, and the threatener portraying him/herself as a victim (along with many other factors) and the actualisation of violence has been previously examined (Turner and Gelles, 2003). The use of 'I,' 'me,' and 'you' was linked to the expression of these types of predictors of violence.

The data in this research generally used 'you' as a subject of a verb at a higher rate than 'I' or 'it.' This tendency has possible implications for threat assessment. If these simulated data can be considered representative of authentic threats, it is possible that threateners do not necessarily refer to specifically *who* is responsible for sending a communication or for the potential actualisation of violence. Instead, threatening messages might appear to be fairly passive or lack an overt target. However, as mentioned previously in this thesis, it is likely that by calling a person's phone number, the identity of the target is clear.

The frequency of these commonly-occurring pronouns in these data appears to differ from non-threatening, English-language corpora. The following table shows the frequency ranking of each of these words across a variety of publicly-accessible spoken and written English-language corpora. For example, the word 'I' is the 10th most common *word* reported in the General Service List (GSL) Corpus (i.e. not the 10th most frequent pronoun in this corpus). The final corpus listed in this table (Table 18 - 'Simulated threat speech corpus') refers to the simulated threats collected for this research.

Table 18: Simulated threat speech corpus - 'I', 'You', 'We' ranking

Corpus	**'I'**	**'You'**	**'We'**
GSL - Written English (West, 1953)	14th	24th	10th
WFWSE - Written English (Leech, Rayson and Wilson, 2001)	17th	21st	10th
WFWSE - Spoken English (Leech, Rayson and Wilson, 2001)	2nd	3rd	5th
WFWSE - Pronouns only (Leech, Rayson and Wilson, 2001)	2nd	3rd	1st
Simulated threat speech corpus	3rd	1st	8th

In the data collected for this study, the most common word is 'you.' However, for the corpora listed in the table above, 'we' is typically more common in written texts than 'you' or 'I' by a notable margin. By comparison, the spoken English corpora listed above rank 'you' and 'we' pronouns as occurring more frequently. This provides some evidence to indicate that the scripts collected for this thesis are written in a tone that is more similar to spoken materials than as typical of written texts. As such, the authors possibly made some conscious effort to write a script that is more representative of a threatening utterance.

Finally, this subsection also showed that the authors of these threatening texts used very few tokens of indefinite pronouns which refer to a person or people (such as 'someone' and 'everybody'). This finding, along with the previously described pronoun data, suggests that these authors generally opted to refer to specific people when writing a threatening text. Perhaps unsurprisingly, the author and his/her target are the most common referents of the pronouns presented in this subsection. With a greater quantity of threatening language data, the relationship between the target of a threat and the use of indefinite pronouns could be more thoroughly examined. It is possible that threats which are not directed toward a target directly (e.g. a written post on an online forum) might feature vague language, such as indefinite pronouns, which does not clearly describe the target or other people who support the threatener.

4.4 Modal auxiliaries

A number of predictions were made about the use of modal auxiliary verbs. Firstly, that there would be a higher frequency of modal auxiliary verbs which express certainty across the threatening messages, in relation to modal verbs which are associated with uncertainty or possibility. This is because authors of threatening texts would presumably want to appear as though his/her actions are highly likely or certain to happen. Secondly, it is expected that indirect threats might feature a greater proportion of modal verbs which express possibility or probability in comparison to direct threats.

Figure 15 shows the frequency of a range of modal auxiliary verbs by authors producing indirect and direct threats. The frequency of each modal auxiliary token shown in this figure is comprised of its multiple variations. For example, the category 'will' contains the words 'will,' 'won't,' 'I'll,' 'you'll' and so on.

Figure 15: Frequency of modal auxiliary verbs in indirect and direct threats

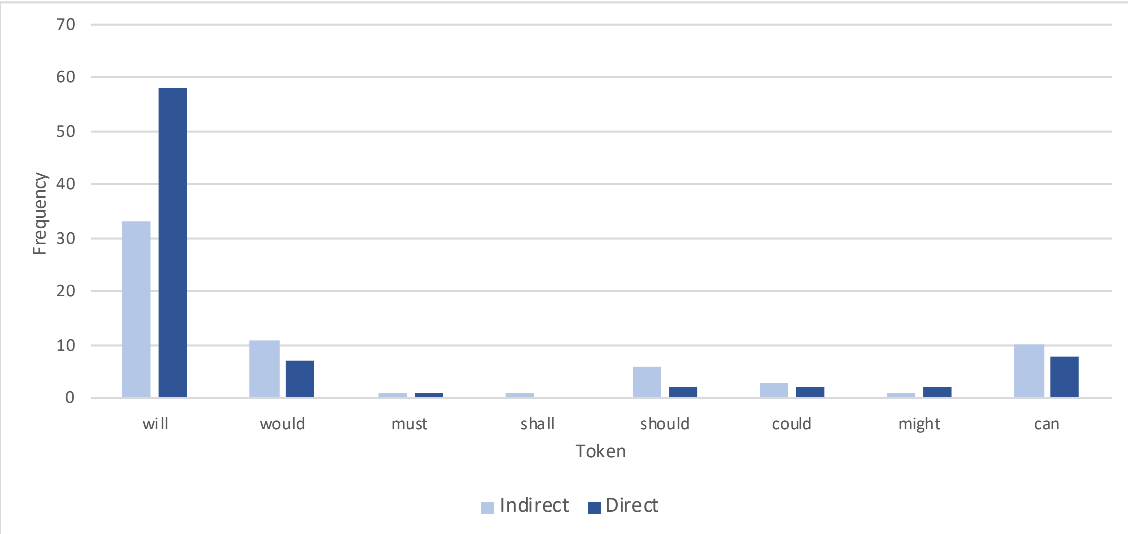

The most commonly occurring modal auxiliary verb in these data is the word 'will.' As predicted earlier in this research, this modal auxiliary occurs more frequently in direct threats than indirect threats. For the indirect threat data, 'will' is the 22nd most frequent word. By comparison, in the corresponding direct threat data, this word is ranked as the 10th most frequently occurring. However, this pattern is not seen for other modal auxiliaries which denote certainty or orders. This includes the words 'would,' 'must,' and 'shall.'

Modal auxiliaries which indicate requests, seek permission or express uncertainty are also infrequent in these data. While these modal auxiliaries seem to occur more frequently in the indirect threats (compared to direct threats), the scarcity of these words across these data does not allow for robust comparisons to be drawn. The exception to this pattern is the word ‘might’ which occurs slightly more frequently in the direct threats than the indirect threats.

The following modal verbs do not occur in these data:

- (s)he’ll/they’ll
- (s)he’d/we’d/they’d
- may

The absence of these modal verbs is likely due to the lack of other characters (both male and female) in the narrative provided to the participants. Also, the modal verb ‘may’ is associated with permission. It is possible that these authors avoided using this word in his/her threats because of this permissive connotation.

4.4.1 Summary and discussion of modal auxiliary verb analysis

The high frequency of the word ‘will’ in these data is consistent with previous descriptions of threatening language, and, in particular, threats which are considered to be higher risk. However, other modal verbs which relate to certainty or likelihood are comparatively infrequent. Modal verbs which suggest uncertainty or possibility are also infrequent in these data. There also seems to be no tendency for modal verbs to be distributed differently across indirect or direct threats. The exception to this generalisation is the modal verb ‘will’ which features notably more in the direct threat data than the indirect threat data.

These data suggest that authors appear to express certainty in the actions specified in his/her threats, more so than uncertainty or possibility. This finding seems to be in keeping with previous research examining the linguistic features of threatening texts (Cohen et al., 2014; Hammer, 2016; Nini, 2017). Also, as previously discussed, the higher frequency of ‘will’ is associated in threat assessment with high level threats or threats which are likely to be followed by the actualisation of violence. That said, some authors of these threatening texts do use modal verbs relating to uncertainty or possibility.

It is possible that threatening texts including these more uncertain modal verbs might be considered to be a lower risk, and by extension, not as likely to be realised. Many publicly-available threat assessment materials discuss the importance of assessing the likelihood that violence will be actualised. However, it is not known whether certain modal auxiliary verbs, in particular, would contribute to the assessment of high or low risk. For the threats collected in this research, there was no possibility of any harm that was described by the authors to be actualised. Despite this, many of these threats assert that harmful actions will occur. It is possible that lay-people or threat assessors might evaluate such texts as being a higher risk (or likely to happen). Further research is needed to better establish the role or 'usefulness' of modal auxiliary verbs during threat assessment.

4.5 Conditional statements

This thesis did not make any predictions (based on previous literature) about the use of conditional statements based on the type of threat, author sex, or author background. There were also very few predictions made in the thesis about how conditional statements might be phrased. It was expected that indirect threats would avoid overt references to harm. As such, it was possible that authors would tend to phrase the consequences of their conditions more vaguely than direct threats.

From the data collected, there was an overall tendency for authors producing threats of either type to use at least one conditional statement (56.58% of the total threats collected for this thesis). In addition, only two authors used more than one conditional statement in a threatening message (Speakers 31 and 35). The phrasing of these conditional statements is somewhat varied. However, the majority (66.66%) were phrased as 'if-clauses.' How authors chose to word their conditional statements (across both sexes and threat types) is shown in Figure 16.

These data suggest that authors tend to opt for 'if X (then) Y' clause structures when forming conditional statements in threatening texts. However, as shown in Figure 8, there are also alternative ways that authors writing threatening texts can express conditionality. This suggests that literature on threatening text assessment should avoid focusing on a particular structure of a conditional threat (for the language in question). Procedures to identify or assess threats should consider that authors can use a variety of structures to

Figure 16: Expressions used to construct conditional statements in indirect and direct threats (%)

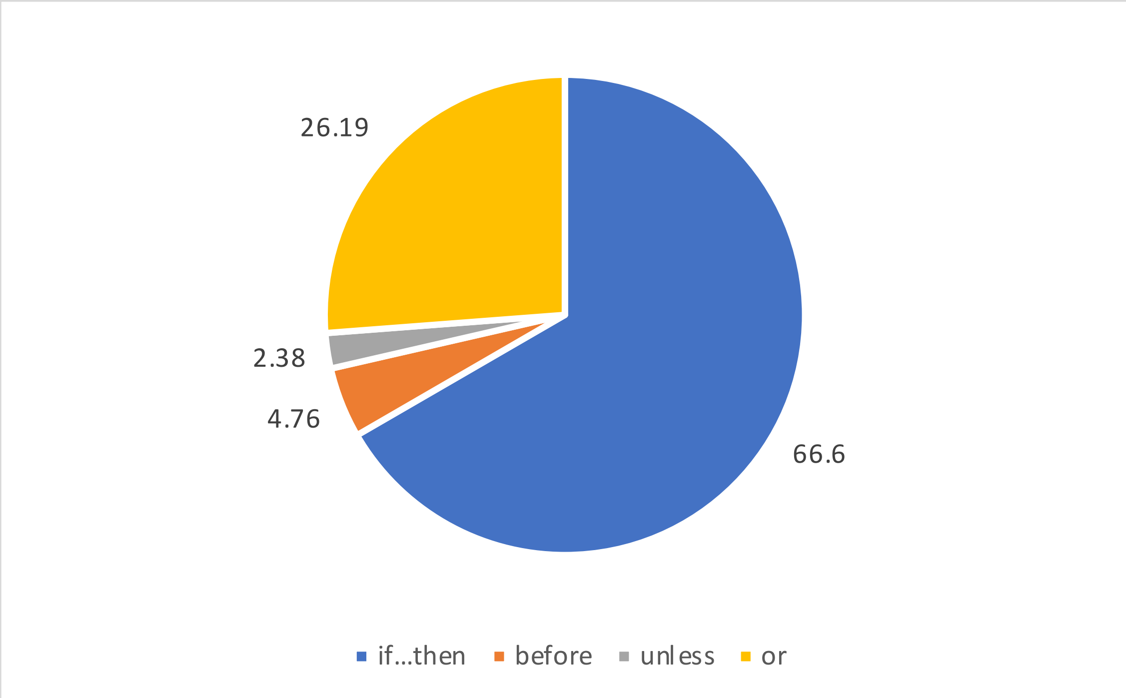

express conditionality, some of which might not occur in the data collected for this thesis. It should also be emphasised that this thesis is not able to establish whether this tendency for 'if' clauses also applies to samples of non-threatening texts provided by the same authors. It is possible that the frequency of 'if' clauses found in the conditional statements of threatening texts reflects their overall high usage in the author's non-threatening written and/or spoken language.

Some threatening messages consisted solely of a conditional statement. This is the case for Speaker 39 (*'If you don't shut down the farm, I've got a gun that will force you to do so.'*) and Speaker 6 (*'Close the farm, or I'll fucking stab you'*) in their direct threats. Messages which consisted of more than a single conditional statement often featured a conditional statement toward the end of the author's message. This can be seen in Speaker 8's message, where the conditional statement is emphasised in bold font:

'You've got to be fucking kidding me. You are one brave cunt. So your not going to shut down your farm? Then I'll fucking do it for you, starting with its toxic owner. I'll make sure you never feel safe again mate because every night you'll be fucking thinking about me blowing up your fucking home up [sic] *Protect yourself and your family while you can*

mate because those abused bloody animals have more use on this earth than you cunts. ***If you decide not to shut I promise you'll live in constant fear of everything you own going to rubble, wherever you hide.'***

A Wilcoxon signed-rank test showed no statistical difference between the number of conditional statements found in indirect and direct threats across the board ($Z = -0.92$, $p = 0.36$). Therefore, the frequency of the production of conditional statements was not associated with one type of threat over another.

References to explicit or vague consequences

It was predicted that for indirect threats, the consequences of conditional statements would tend to be vaguely-worded. This is because indirectly-worded threats would presumably avoid making clear references to harm. Therefore, if authors chose to use conditional statements when writing their indirect threats, the consequences specified would presumably avoid making references to violence or other forms of harm. The condition itself should be clear enough for a target to understand and follow. This condition should also be something which ultimately benefits the threatener, likely at the target's expense. Therefore, it is logical that authors would phrase their conditions in a way which clearly communicates what it is they want.

In addition, the consequence would also need to be clear enough for a target to recognise and interpret it as such. In this thesis, 'vague' consequences are defined as consequences where it is not possible to determine the action to be taken based on the words used. For example, Speaker 6 produces the following phrase as part of his indirect threat: *'I hope you make changes soon or we will have to take further action.'* Here, the 'action' to be taken is unspecified. It is also unclear when this action will occur, or who/what it will be directed towards. Compare this to Speaker 2's direct threat, in which he says: *'Stop this now or I will be taking this to court.'* In this example, there is an obvious consequence should be farm's practices continue (i.e. the commencement of legal action).

Figure 17 shows the overall distribution of vague and explicit consequences found in the written indirect and direct threats collected for this thesis.

Figure 17: Total frequency of explicit and vague expressions used to express the consequence of conditional statements (%)

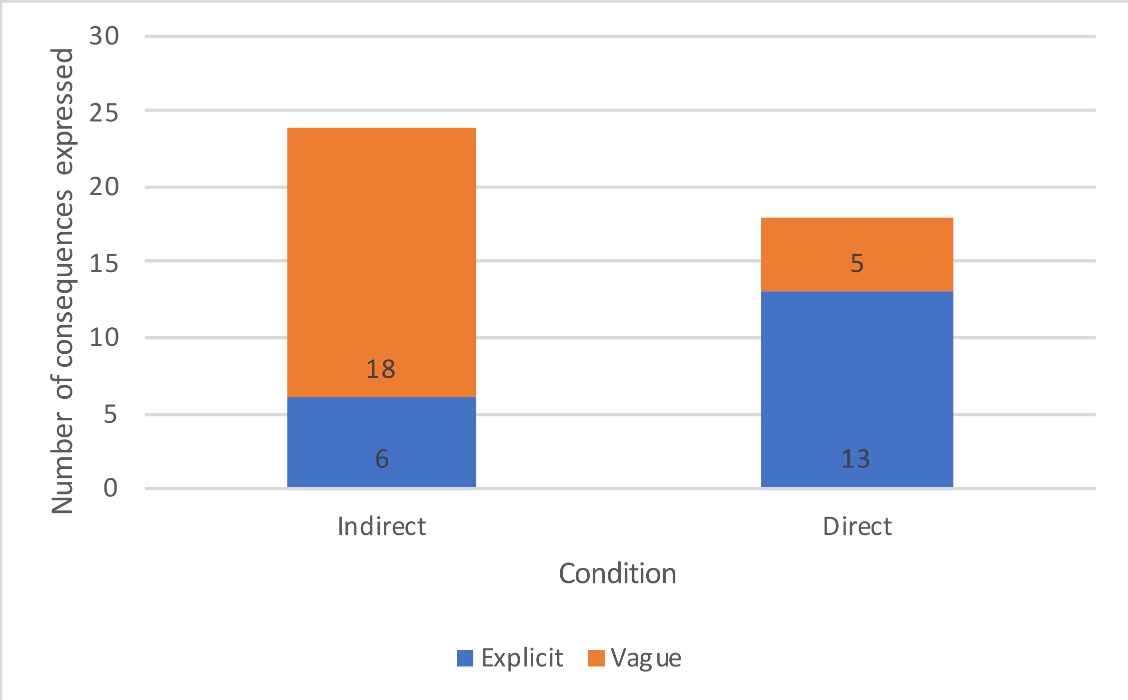

This figure shows a tendency for conditional statements found in indirect threats to be expressed with a vague consequence. Conversely, there is a tendency for direct threats to be communicated using an explicit consequence. Table 19 outlines the types of explicit consequences found in both indirect and direct threats.

Table 19: Explicit consequences occurring in direct and indirect threats

Type of harm specified	Indirect	Direct
Closing down the farm	2	2
Damage to property	2	3
Legal action	1	1
Continuation of protest	1	0
Physical harm against person(s)	0	9
Physical harm against unspecified target	0	1

Note that some conditional statements included more than one explicit reference to harm. For example, Speaker 25 makes the following conditional statement in his direct threat: *'I am not one man, Mr West and I swear if you don't stand down, we will blow you and* *your filthy little farm apart.'* Here, there are references to damage the target (Mr West) and his property. As such, this message would be tagged as using an explicit reference to harm the target, and as using an explicit reference to damage property. Table 22 also shows that the explicitly-phrased consequences made in direct threats were predominantly threats of physical violence. For indirect threats, explicit references to harm centred on either having the farm closed down, or damaging property.

These results are in keeping with the findings of Chapter 5.5.17, in which indirect threats were shown to lack overtly-violent expressions of harm. Also, direct threats featured a higher frequency of apparent references to physical harm or violence compared to indirect threats. These findings are shown in Figure 18.

Figure 18: Frequency of explicit or vague consequences used across author groups

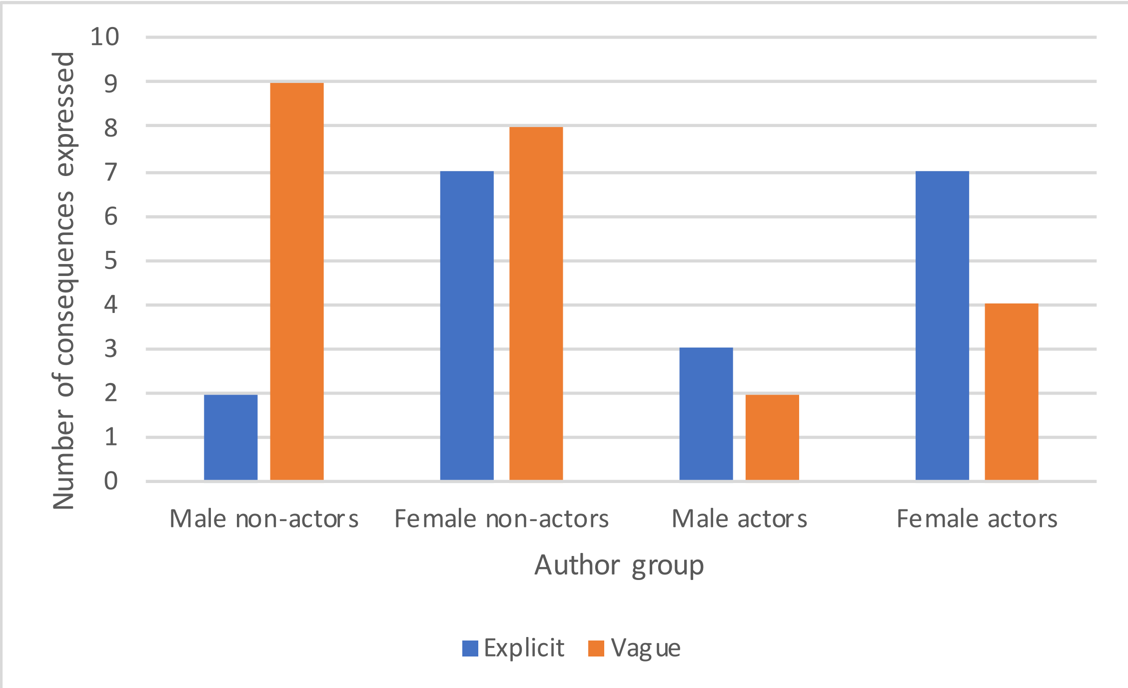

In total, the non-actor data contained 26 consequences in their threatening messages. This compares to the actor data containing 16 consequences. Figure 18 also shows that overall, non-actors tended to express a higher frequency of vague consequences than explicit consequences. For actors, there was a tendency to produce more explicit consequences than vague consequences. As previously reported, differences between males and females across these groups are non-significant.

4.5.1 Summary and discussion of conditional statement analysis

This subsection provided an overview of the frequency of conditional statements within these simulated threatening text(s). In particular, there was a focus on the differences between how authors express vague or explicit consequences when writing indirect and direct threats. The results of the analysis of conditional statements found in the threatening texts collected for this thesis can be summarised as follows:

- Just over half of the threatening messages featured at least one conditional statement. There were only two authors who used multiple conditional statements during their threats.
- Most conditional statements were phrased using 'if...then' structures. Other words associated with conditional statements include 'or,' 'before,' and 'unless.'
- Some threatening messages consisted solely of a conditional statement. In messages where conditional statements are present, it seems to be common for the conditional statement to occur around the end of the message.
- There was no statistical difference in the number of threats produced for indirect and direct threats.
- The conditional statements of indirect threats tended to feature a vague consequence. On the other hand, direct threats tended to feature explicit consequences.
- For explicit consequences, direct threats predominantly specified committing an act of physical harm against the target. For indirect threats, the harm specified varied between somehow having the farm close down, damaging the farm itself, taking legal action, or continuing his/her protests.

As predicted, there is evidence to suggest that the conditional statements found in indirect threats tended to express consequences which are vague or lacking in detail. Direct threats were, in general, more likely to feature conditional statements which included an explicit consequence.

What is not known yet is the effect of vague or explicit conditions on targets of threatening language. It is reasonable perhaps to assume that clear references to violence (within a conditional statement) would cause a recipient to feel fearful. However, the expression of vague consequences might also be an effective means of inducing fear in a target. Firstly, providing little or no detail about the consequence of a conditional statement could allow a threatener to assert plausible deniability. In other words, by not expressing something overtly harmful, a threatener could claim that he/she was misinterpreted. Secondly, vague conditions could reasonably cause a target to feel anxiety. This is because there is uncertainty about, for example, when or where a harmful action might occur. In addition, threat assessors might also face difficulties when advising targets on how best to respond to a threatening communication. The effect of conditional statements on threat assessment by lay-people and threat assessors does not seem to have been researched and/or presented publicly.

What is publicly-available, are numerous examples of threat assessment materials which define indirect, direct and conditional threats as separate categories. As discussed earlier in this research (see chapter 4.1.3), the condition of a conditional threat does not seem to convey something that is (explicitly or implicitly) harmful to the target. Harmful actions are expressed during the consequence of a conditional statement. This pattern can be seen in the data presented in the Appendix of this research. It is also worth noting that the conditions expressed in this research are easy to interpret. In other words, the condition itself is not vague. Understandably, it is in the interest of the authors in this experiment (and threateners more generally) to clearly articulate their goal. After all, this might be the primary (or sole) reason for sending a threatening communication.

4.6 Profanity

Chapter 2 briefly discussed the association of threatening or aggressive behaviour with profane language. In particular, the idea that males swear more frequently than females, both as a means to communicate anger and in emotionally neutral or non-angry speech. The expectation would be for the male participants in this research to produce a higher frequency of profane words than females. There is also literature where profane or offensive language is said to intensify the offensive of language such as threats (Culpeper, 2011). This subsection will also consider whether there is evidence to characterise direct threats as containing greater uses of profane words than indirect threats. This research cannot compare these results with a sample(s) of the author's non-threatening text. As such, it is not known whether the findings for profanity presented here are in-line with the author's 'typical' authorship style.

For this thesis, the threatening messages collected were tagged by the author for instances of profane words. Although LIWC2015 is able to provide a percentage of a text which contains swearwords, it is unable to recognise various examples of profane words which occur in the messages collected for this thesis. For instance, profane terms such as 'shit-faced' and 'prick' are not recognised by LIWC software. Related to this, terms such as 'fucking' were categorised as a verb which describes a sexual act, instead of as an adjective or adverb. The latter word classes appear to better reflect the use of the word 'fucking' in the data collected for this study.

The table below (20) shows the profane words observed in the data of this thesis, and their occurrence in indirect and direct threats. The frequency of each word listed in this table also includes any morphological variations. For example, the category 'fuck' includes variations such as 'fucking' and 'fucker.'

Table 20: Profanity occurrences found in data

Profane word	**Indirect**	**Direct**
Fuck	13	24
Shit	3	4
Cunt	2	3
Prick	1	0
Bastard	1	2
Ass	0	1
Balls	0	1
Bloody	0	1

By far the most frequent profane word in these data was 'fuck' (and its corresponding variations). The most common variation of this word, 'fucking,' was used as an adjective to intensify the meaning of a verb. For example... The words 'fuck,' 'cunt,' 'prick,' 'ass,' 'bastard,' and 'shithead' are all used pejoratively to refer to the target of the threat. One term referred to the location of the target ('shithole'). The only instances where profanity referred to another entity was the word 'bloody' written by Speaker 8 in the following direct threat message: *'...because those abused bloody animals have more use on this earth than you cunts.'*. Here, 'bloody' refers to the animals on the farm. Also, the word 'cunts' referring to other people who might commit harmful actions against the reader (S8 direct).

Overall, there were 56 uses of profane words throughout the threatening messages collected for this thesis. Of these, 20 (35.71%) occurred in the indirect threat messages, and 36 (64.29%) occurred in the direct threat messages. A Wilcoxon signed-rank test did not show a significant difference between the use of profanity across all indirect and direct threats (Z = -1.96, p = 0.05).

Figures 19 and 20 show the use of profanity by male and female authors. As previously described, not all authors produced a written threat. As such, their data is not displayed in these figures. Nor will any direct comparisons be drawn here between the frequency of profanity usage of male/female authors or non-actor/actor data.

Figure 19: Frequency of profane words used by male authors in threatening messages

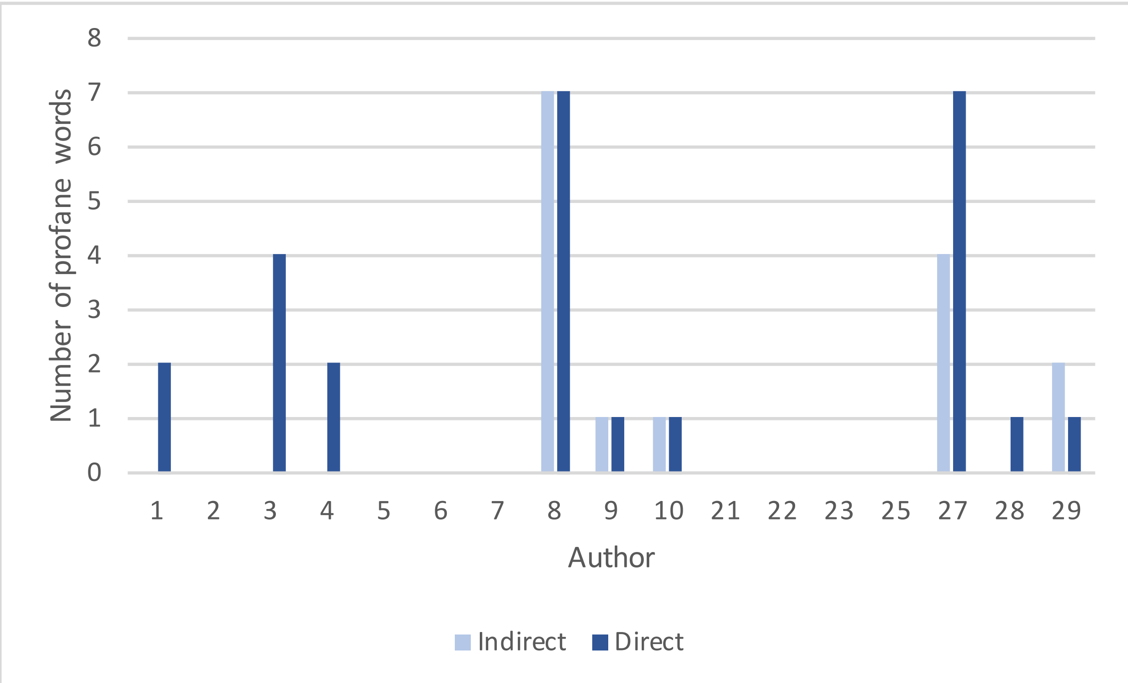

Just over half of authors (57.89%) do not use any instances of profane language when writing a threatening message. Of these authors, the majority (68.18%) are female. However, some authors produced notably more profane words than the remaining participants (Speakers 8, 27 and 32). Of these, Speakers 8 and 27 are male. Among the female participants, only Speaker 32 produced profane language in an indirect threat. The remaining female participants who used profanity only did so during direct threats.

Based on Figures 19 and 20, no apparent differences were observed visually between the

Figure 20: Frequency of profane words used by female authors in threatening messages

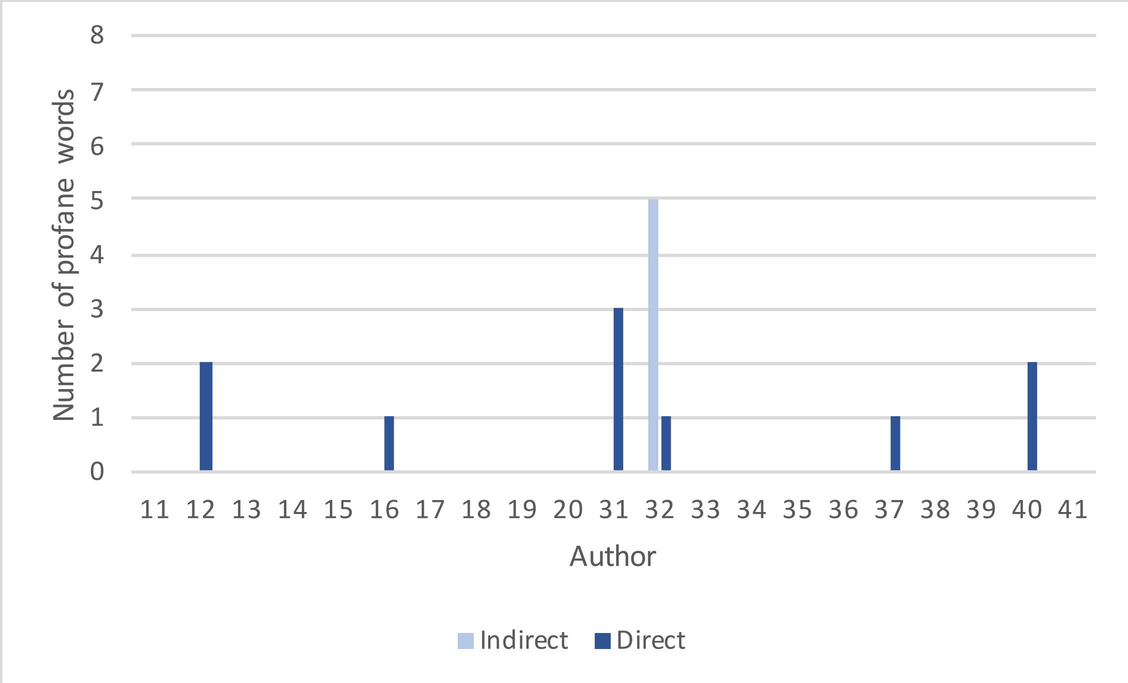

non-actors and actors. As there were fewer actors who provided any written texts, it was not possible to collect robust statistical evidence to show any differences in frequency between the two author groups.

4.6.1 Summary and discussion of profanity analysis

The results of the analysis of profane words found in the threatening messages collected for this thesis can be summarised as follows:

- Just over half of the written threats contained at least one instance of profanity. Of these, the word 'fuck' (and its variations) was the most frequently used.
- There was no statistical difference between indirect and direct threats regarding the frequency of profane words. However, in these data, profanity most frequently appeared in direct threats.

As noted by Gales (2010), threats are not necessarily characterised by unpleasant or offensive language. This overall lack of profanity for threatening language is evident in the current data. There also appears to be some evidence available to suggest that male and female authors might differ in their profanity usage when authoring threatening language. From these data, male authors seem to be more likely to use profanity (and with greater frequency) compared to female authors. This finding relates to previous literature which has examined male and female use of emotional words in written texts. There is currently uncertainty as to whether or not males or females use more words in their (non-threatening) texts which are associated with anger or aggression. In some studies, men use anger words more frequently (Pennebaker, Mehl and Niederhoffer, 2003), while in other studies this pattern is less clear (Newman et al., 2008).

4.7 Promises and warnings

Chapter 2 discussed whether threats can be considered as a distinct type of speech act. Central to this debate is the lack of a performative verb available to authors which can specifically express a 'threatening' intent or action. This issue was also previously addressed in Chapter 5.5.1 in which Tropes software was used to identify performative verbs across these data. However, threats have been associated with (or defined as) promises and warnings. As discussed in Chapter 2.2.4, promises and warnings were also linked with a commitment to carrying out an action. By stating a promise, speakers are bound to a commitment that they will carry out the action stated. If the speaker does not deliver on his/her promise, this failure is often perceived negatively by the recipient. This is, understandably, unlikely to be true for threats. For threateners, framing their threat as a promise could help to iterate that he/she sincerely intends to carry out a harmful act. As such, the target might be more likely to perceive the threat as more likely to be actualised.

In the data collected for this thesis, there are a few instances of the verbs 'promise' or 'warn' being used in threatening messages. Only Speaker 2's indirect threat included some variation of the word 'threat.' However, in this case 'threat' is used as a noun, rather than a verb (*'Which are a threat to everything you care about.'*). Also similarly uncommon in these data were variations of the word 'promise.' Only Speaker 8's direct threat featured this verb (*'If you decide not to shut I promise you'll live in constant fear of everything you own going to rubble, wherever you hide.'*)

More frequent in these data were variations of the word 'warn.' These tokens can be seen in Speaker 12's indirect threat, and the direct threats produced by Speakers 1, 12, 13 and 38. Speaker 12 was the only author of these data to express both of her communications as a warning. Overall, the majority of the threats do not include any words which overtly refer to their previous communications (or future actions) as a threat, warning or promise. These verbs were also not identified by Tropes as tokens of performative verbs.

These data suggest that authors seldom framed their threats as other types of speech acts. As suggested earlier in this thesis, it was uncommon for these threats to include any performative verbs whatsoever. In other words, the expression of the threats collected for this research seemed to rely very little on verbs which communicate a speech act of any kind. Instead, there appears to be a tendency for these texts to use verbs which express the author's potential future performance of a physical action, or verbs which refer to facts, actions or objects. These findings may be taken as evidence which supports the treatment of these speech acts in linguistic research as being distinct (but related) to one another.

The instances of 'warn' referred to above mostly occur in the direct threat data. It is plausible that some of these authors were framing their previous indirect threats as a warning rather than as a threat. This study did not investigate the reasons why this might be the case, or how authors differentiate between these related speech acts. As such, it is not clear whether the participants in the present study equated threats with warnings. There is also the possibility that the participants deliberately chose to refer to their contact with the target as warnings rather than threats. Exploring how authors and readers infer the meanings of these words pragmatically would shed light on overlap in terms of how these words function.

4.8 Assertive statements

Threat assessment materials also link the use of assertive statements to the increased risk of the actualisation of harm (Threat Assessment Procedures Manual, 2017). This includes statements such as 'I really mean it,' and/or 'This is a genuine threat.' These statements explicitly express the author's intention to carry out some kind of harmful action or assert that the communication is a 'sincere' or 'genuine' threat. In other words, the author might include such statements to convey that his/her threat is not 'empty.' For the messages collected for this thesis, there were two tokens which were tagged as clarification statements. These consisted of the following phrases included in Speaker 29's indirect threat and Speaker 31's direct threat respectively: *'...you can trust me on that.'* and *'...and you better pay attention because I'm not messing around now!'*

This scarcity of data suggests that these authors, overall, did not opt to include assertive statements in their threats. While such statements might raise the level of severity of a threat in the judgement of a trained threat assessor, it is not known how lay-people perceive these statements of this type (whether lay-people are broadly in agreement that including a statement like those shown above would make an author seem more likely to perform a harmful action). Further analysis of the data would reveal the presence and functions of other assertive statements.

The following subsections detail how, overall, authors made references to the time, location, nature of threatened harm, and the target of a threat, as well as to themselves as the authors of the threatening messages. These features have been selected for analysis because they have previously been associated in this research with the likelihood that the actions stated or implied in a threat will be actualised. As before, these data were tagged manually by the author using the text highlighting function of Microsoft Word. This is opposed to using an automated tool such as those shown earlier in this chapter. As such, this analysis offers a more subjective and qualitative insight then previously described in the current study. This subjective analysis, however, is perhaps more similar to current, non-automated, threat assessment practices. This said, as described in Chapter 2, it is not clear exactly how threat assessment practitioners infer this information based on a written text.

4.9 References to time

Of interest to threat assessors is *when* an alleged harmful act is likely to occur. Knowing this can help the target and law enforcement or security experts to plan an appropriate response. For example, if a threatener said that a bomb would be detonated in the next hour, an appropriate response might be to evacuate anyone in that area as quickly as possible. This subsection focuses on how the authors of these simulated threats place their harmful actions in time. In particular, there is a focus in this subsection on whether they inform the target of when a harmful act might occur. Also, in time, in particular whether they inform the target of when a harmful act might occur and whether this information is embedded within a conditional statement.

Indirect threats

Overall, it seems that the indirect threats collected for this study tend to make somewhat vague references to the future occurrence of harm (implied or overtly stated). This vagueness is typically seen in the condition stated as part of a conditional statement (or request), for example through the use of words such as 'soon':

- *'If you keep up what you doing there, I know that crazier cunts than me will be* *knocking at your family's door very soon.'* (Speaker 8)

Another relatively common word used to express time in these threats is 'immediately:'

- *'...and begin the immediate ceasing of all its medical trials and for the safe release of all non-human animals contained on site.'* (Speaker 4)
- *'If I were you I would seriously consider close your farm immediately* (sic)' (Speaker 34)

Many of these indirect threat messages do not contain any reference to when a harmful action might occur if the farm does not close down. In particular, those written by the female authors appear overall to not mention any time frame whatsoever. As explained previously, this thesis will not directly compare male and female authors due to differences in the amount of data available for each of these groups. Of the female participants who do refer to the timing of a harmful action, it is again expressed vaguely:

- *'Close the farm before the campaign gets any worse.'* (Speaker 15)
- *'if* (sic) *not, we're not going to have a choice but to hurt your shit until you realise what you're doing is wrong'* (Speaker 27)
- *'We won't stop calling, West. Not until we decide we need approach this with other means anyway.'* (29)

In the latter two examples above, the timing of the 'harmful action' is dependent on the target's response to the situation at the farm. As such, the target is given no specific deadline to change his behaviour. In these examples, it is claimed that the adverse action(s) of the threatener will continue (or escalate) unless the target adjusts his behaviour.

Direct threats

For direct threats, the timing (and/or conditionality) of a harmful act is typically more specific. This was generally expressed in terms of the number of hours or days that the target is being given to change his behaviour, as can be seen in the following examples:

- *'Your business will burn and your family will stare down the barrel of a gun unless you stop in the next two days.'* (Speaker 2)
- *'You have two days to close down the farm - or I will.'* (Speaker 13)
- *'You have twenty four hours to shut down that fucking joke of a farm.'* (Speaker 37)

As with indirect threats, many of the authors in this study did not write their direct threats with any reference to when harm would occur. There also did not seem to be any noticeable tendency for authors who had previously expressed the timing of an indirect threat, to express time in their subsequent direct threats in any particular way.

4.10 References to location

Also of interest in this thesis is whether the authors made references to the location of any harmful action. Providing this information might help to convey to the target (or threat assessors) that the threat is likely to be actualised. For both indirect and direct threats, the most frequently stated location was the farm (i.e. the business being targeted). This subsection provides an overview of how authors of indirect and direct threats refer to the place where they 'intended' to bring about harmful acts.

Indirect threats

In indirect threats, in particular, there was a tendency for references to the farm to be embedded within a conditional statement, rather than named as a target itself. In other words, the authors of indirect threats usually made overt references to the farm while expressing their aim for the business to close down. This can be seen in the following examples:

- *'If you don't shut down West's farm immediately, there will be consequences'* (Speaker 1)
- *'If this farm isn't shut down before I get there, I'm going to make you regret it.'* (Speaker 17)

The farm itself was also identified as a clear target of a harmful action. However, the farm is identified as a location where a harmful act may take place in the following example:

- *'If you don't stop, we'll storm the farm with guns, and won't be stopping for anyone in the way'* (Speaker 20)

In general, stating *where* any potential harm might occur appears to be less frequent than is the case of direct threats. Indirect threats which specify a particular location at which a harmful act may occur tend to refer to the target's home:

- *'...and if necessary would be ready to move our protest to his home and his family. '* (Speaker 4)

- *'If you keep up what you doing there, I know that crazier cunts than me will be knocking at your family's door very soon.* (Speaker 8)

As such, authors are making claims that they knows where the target (or his family or property) are based, in spite of the authors not having access to this (fictional) information.

Direct threats

For direct threats, the farm (and possibly its inhabitants) is far more frequently an obvious target of violence:

- *'You and your farm are gonna burn...'* (Speaker 7)
- *'If you don't shut down your farm, we'll blow it up.'* (Speaker 14)
- *'...then I will make very sure that neither you nor your business come out the other side.'* (Speaker 19)
- *'Right now we're working our way into your factory to blow it up.'* (Speaker 21)

In addition, other locations are described in reference to actions which are in some way harmful (or disadvantageous) to the recipient:

- *'Stop this now or I will be taking this to court.'* (Speaker 2)
- *'Ok, Let's* (sic) *see what you'll say when I arrive on your doorstep with my shotgun fully loaded.'* (Speaker 35)
- *'I will gun you down and blow your house into fucking bits.'* (Speaker 40)

Similarly to indirect threats, these authors also claim to know where the target (or his family or property) is based:

- *'We know where you are and were not happy, if you dont stop we might have to pay you a visit.'* (Speaker 20)
- *'How would you like me to come to your house with a gun and force you to mate for medical research.'* (Speaker 32)

However, as before, these locations are lacking in detail. While some authors attempted to bluff about their awareness of where the target lives (or details of his daily routine), no author in this experiment provides a more specific location in his/her messages (e.g. a specific home address). Had participants been given a fictional address as part of their brief, it is possible that this information would have been included in their messages.

4.11 References to violence or harm specified

This subsection presents an overview of how violent or harmful acts are expressed in the indirect and direct threats collected for this thesis. It is expected that indirect threats will be characterised by a lack of overt references to harm. By contrast, direct threats are expected to include explicit references to actions which would be adverse to the target. As previously described, the terms 'indirect' and 'direct' are not always used across the relevant literature to refer to whether or not violence is explicitly stated. As such, there might be differences across the indirect and direct threats collected for this thesis in how authors understood the terms 'indirect' and 'direct'.

Indirect threats

The indirect threats collected for this thesis generally lacked any overt references to violent or harmful acts. For the threats which did refer to harm, there are several which appear to threaten legal action, sometimes within a conditional statement:

- *'If you do not close down the farm we will have to take legal action defending the rights of the animals and get the police involved moving these animals to a safer place and shutting down your farm for good'* (Speaker 14)
- *'If you do not shut down the farm I will be forced to report you to the police and all local animal rights groups.'* (Speaker 35)
- *'The coppers will get on to you and believe me you will get punished. they (sic) dont take this sort of thing likely (sic) you know. In fact, I've heard they can get quite passionate in ensuring justice is served if you get what I mean.'* (Speaker 40)

There were also a few threats which explicitly state a violent action, or include words which relate to harmful actions:

- *'If you don't stop, we'll storm the farm with guns, and wont be stopping for anyone in the way'* (Speaker 20)
- *'How about everyone in the town locks you up in a cage and start testing on you?'* (Speaker 31)

For the majority of the indirect threats, there is a tendency for the extent of any 'harm' specified to be difficult to gauge fully. Many of the 'harmful actions' referred to in indirect threats might be best described as 'ominous'. The following examples seem to rely on the reader to be able to infer an adverse consequence of a conditional statement or action:

- *'I hope you make changes soon or we will have to take further action.'* (Speaker 6)
- *'Find a better way to make money you fuck or else.'* (Speaker 10)
- *'If this farm isn't shut down before I get there, I'm going to make you regret it.'* (Speaker 17)
- *'Rest assured that if you do not discontinue this disgusting practice, it will no longer be your livestock's screams that are heard for miles around.'* (Speaker 25)
- *'Many people are against what you are doing so for your own sake, do the right thing or there will be consequences.'* (Speaker 33)
- *'We'll keep calling. We have your address. This will just keep getting bigger and bigger.'* (Speaker 38)

Direct threats

The direct threats collected for this thesis are characterised by words and/or phrases which detail violent or harmful actions. As discussed in Chapter 5.5.9, many of the conditional statements included a consequence which made an explicit reference to something which would be disadvantageous to the recipient. The harm specified consists of references to physical harm, sexual violence, legal action, and causing damage to the target's property. The following excerpts show how the authors describe a range of violent acts when instructed to write something 'directly threatening.'

- *'You are going to give us a guarantee of Wests Farms closure or I am going to kill you.'* (Speaker 4)
- *'How about I come over here with my mates and take turns with your wife?'* (Speaker 8)
- *'So get ready to see your revolting farm blitzed to pieces.'* (Speaker 18)
- *'We wont play nice man, Im fucking military trained, and Im going to break your shit-faced jaw wide open with my sledgehammer if thats what itll fucking take'* (Speaker 27)
- *'I can kill you in a matter of seconds and bomb your house as well.'* (Speaker 36)

There are also examples of direct threats which do not appear to have any clear references to violence, or specify how any violence might occur:

- *'So it seems you dont care about what the rest of us think. You dont want to make too many enemies Mr. West. Worse comes to worse well just take matters into our own hands. Be careful.'* (Speaker 6)

- *'I have access to some incredibly harmful and painful resources that I dont want to have to use. But if you force me to, then I will make very sure that neither you nor your business come out the other side.'* (Speaker 19)
- *'We know where you are and were not happy, if you dont stop we might have to pay you a visit'* (Speaker 20)

These findings suggest that, broadly speaking, authors appeared to distinguish indirect threats from direct threats by including language which clearly refers to harmful actions. There is variation among the direct threats regarding *who* is committing these acts. There appears to be a tendency among the authors of these threats to indicate, or imply, that they personally would perform the specified action(s). As previously shown, there are also references in these data to other people potentially committing these acts:

- *'Well come with guns and explosives and show you what it feels like to be treated like your animals.'* (Speaker 15)
- *'If action is not taken from you soon, myself, the campaign, and everybody else who opposes your cruelty will use our weapons'* (Speaker 33)

4.12 References to target and author

Another consideration when assessing a threat is identifying the target(s) and the threatener. Currently, it is not known whether authors of threatening messages frequently identify themselves or their targets in the message itself. After all, this information is also inferred from the communication itself. For example, if a mobile phone receives a threatening message, it can be assumed that the owner of the device is the intended target. As such, the authors in this study might not refer to their target at all. This is because the authors were directed to leave an audio message on the target's phone. It might, therefore, have been though redundant by the authors to refer to the target explicitly by his name. Also, due to the experimental design, there might be differences in how the target is addressed between indirect and direct threats. Put simply, the target could be identified in the first message (i.e. the indirect threat), but not necessarily in the follow-up message (i.e. the direct threat).

As previously described in this research, the target of the threats is often referred to using the pronoun 'you.' The following table (21) shows every other word or phrase used to describe the target by each author in their indirect and direct threat:

Table 21: Target descriptors used by authors over direct and indirect threat cases

Speaker	Indirect	Direct
1	Tom West	Tom West
2	Mr West	Mr West
3	Mr West	fucking bastard
4	Mr Tom West, Mr West (x2)	N/A
5	Mr West	Mr West
6	Mr West	Mr West
7	N/A	N/A
8	mate (x2), twat, pig	toxic owner, cunts, one brave cunt
9	Tom West, Nazi white privileged scum, oxygen thief, prick	animal bating cunt
10	N/A	stupid fucking man
11	N/A	N/A
12	Mr West	Mr West
13	N/A	N/A
14	Mr West	Mr West
15	Mr West	N/A
16	N/A	N/A
17	N/A	Mr West
18	Mr West	Mr West, disgusting unethical swindler
19	N/A	N/A
20	N/A	N/A
21	Beast, despicable man	filth
22	N/A	N/A
23	Mr West (x2)	piece of filth
25	Mr West	Mr West (x2)
27	piece of shit, arrogant selfish bastard	Mr West, man (x2), piece of shit
28	Mr West	N/A
29	you coward, West, cunt	fuck
31	N/A	sick, twisted bastard; Mr West
32	West (x2), fucking bastard	fucker
33	Mr West	N/A
34	Mr West	N/A
35	N/A	N/A
36	N/A	N/A
37	Mr West (x2)	Mr West
38	sick human being	N/A
39	N/A	N/A
40	N/A	mate, big man
41	Mr West	Tom West

Indirect threats

Of the indirect threats, 14 of the messages (around 34% of the indirect threat data) contain no explicit references to the target. The remaining indirect threats typically refer to the target by his name, most commonly, this is expressed using his title 'Mr West.' It is comparatively uncommon for the target to be referred to by his given (or full) name. There are no examples of the target being identified using only his given name. These data also include profane and derogatory language which seem to refer to the target. This includes phrases such as 'Nazi white privileged scum', 'pig' and 'sick human being.' In addition, there are two instances of the target being referred to as 'mate', a typically friendly form of address. These data indicate that for some of the authors, there is no specific target identified in the wording of the message itself. Rather, the language referring to him might offer some insight into the threatener's feelings towards him.

Other than the use of 'I' and 'we,' these messages feature very few details concerning the author him/herself or his/her identity. One script was tagged for additional references to the author during analysis: 'crazier cunts than me' (Speaker 8). This example, however, does not provide any detailed information about the threatener.

Direct threats

For the direct threats, 16 of the messages (around 39% of the direct threat data) contained no explicit references to the target. Similarly to the indirect threats, these data also consist of a mixture of the target's name and derogatory language referring to him, for example, phrases such as 'fucking bastard,' 'one brave cunt' and the word 'fuck' (used in Speaker 29's direct threat as a noun). The target is also referred to as 'mate' and 'big man' by Speaker 40. Once again, these words are often used as positive forms of address. As with indirect threats, there are also no instances of the target being addressed by only his given name (i.e. 'Tom') in these data.

Also, as with indirect threats, there are very few examples where the author of the message is referred to in any way. The two examples of tagged in these data were the phrases 'I'm fucking military trained' (Speaker 27) and 'big man' (Speaker 40). Concerning the former example, such information would presumably be useful for the police or threat assessment practitioner(s) to consider. This information might also influence whether a target would be likely to report the message to the relevant authorities in the first place.

4.12.1 Summary and discussion of references to time, location, violence or harm specified, and the target and author of simulated threats

This subsection provided an overview of how the authors in this study referred to the timing, location, nature of the harm specified and the target of the threat. Of particular interest was whether authors generally differentiated indirect and direct threats using these parameters. On the basis of the data available in this research, the following observations can be made:

- Details about the time-frame of a threat appeared to differ between indirect and direct threats. For indirect threats, authors tended to not provide a concrete time-frame for when a harmful act (implied or explicit) might occur. Direct threats more frequently included more specific time-frames. These time-frames were mostly rather short-term ('immediate', hours, days), rather than longer-term (e.g. months or years). However, across the data, the authors did not necessarily provide any details about time, for example, when the target should address his actions, or when the author would initiate a harmful action against the target.

- Details about the location where any harmful action was threatened to occur were seemingly less frequent and lacked detail. The farm itself was often mentioned within the conditional element of a conditional statement (i.e. the farm should be closed, or something adverse will happen to the target). Occasionally in these data, the farm itself was also a target of threatened harm. Other locations were also referred to in these data, in spite of no other locations being provided in the briefing or instructions for the participants in this experiment.

- These data include a variety of implied and explicitly-stated references to committing harmful actions. Overall, the indirect threats lack any clear references of violent acts. Instead, it appears that many of these messages convey vague or ominous expressions of harm. For direct threats, there are numerous expressions of harm through the use of a variety of verbal strategies. The words used can be associated with different kinds of violence or harm. However, most of the harm specified appears to centre on committing physical violence against the target.

- Explicit references to the target of the threat mostly refer to the target's formal title (i.e. 'Mr West'). In addition, the target also seems to be referred to using derogatory language. Throughout both the indirect and direct threat data, however, numerous authors make no explicit reference to the target's identity whatsoever. There also appears to be little evidence of a relationship between how an author refers to the target in an indirect threat and his/her subsequent direct threat.

- Aside from the use of 'I' and 'we' pronouns, the authors generally make very few references to themselves in these threatening messages. No author provides a name (real or fictional), or other personal information that could positively identify him/her.

4.13 Summary and discussion of written threat data analysis

This chapter has aimed to provide a broad overview of the linguistic features present in threatening language. The purpose of these analyses was twofold; to gauge how similar these simulated data are to previous descriptions of authentic threats (**research question 2**), and to explore the linguistic expressions of 'direct' and 'indirect' threats (**research question 2.a**).

There were a number of linguistic features which appeared to be consistent with previous descriptions of threatening language and/or features which are said to predict the actualisation of violence. These features included the presence of words associated with negative affect and present tense, pronoun usage, and modal auxiliaries. This indicates that the simulated data appeared to use linguistic features in a manner which resembles the authentic written materials described in previous research.

Table 22 provides an overview of the usage of linguistic features in the simulated data collected for this study which have previously been linked with threatening language in other literature:

Table 22: Linguistic features in simulated data compared to previous studies

Findings consistent with previous studies	**Findings inconsistent with previous studies**
Few words used to express doubt or uncertainty (Gales, 2010; Nini, 2017)	Few instances of assertive statements (Gottschalk and Betchel, 2001)
Expression of tense and affect (Chandler, 2017)	Vague details across the data regarding the time, location, manner of harm (O'Toole, n.d.)
High frequency of 'you' (Smith, 2006)	Low rates of profanity (Gales, 2010)
Low rates of profanity (Gales, 2010)	

In addition, there were also very few instances of performative verbs in these data. In particular, there were only a few token of the performatives 'promise' and 'warning'. This finding has significant implications for how lay-people conceptualise the expression of threats. It is possible that while threats are likened to warnings and promises in academic literature (Storey, 2005) or in threat assessment resources (de Becker, 1997), lay-people might infer these speech acts differently to one another. Also, it could be the case that the threatening language collected in this study was characterised by the use of verbs which express physical actions (rather than those 'performed' through speech). This impression is supported by the comparatively high number of factive verbs in these threatening language data relative to performative verbs (Chapter 4.1).

What is lacking from the current study (and previous research), is an examination of how threatening texts differ from an author's 'typical' usage of linguistic features. Examples of each participant's 'typical' written data were not collected for this study because, due to the lack of research on threatening language, it was not known what written data would provide a suitable comparison.

However, what seemed clearer was that the authors wrote indirect and direct threats using mostly similar linguistic features. This impression is supported by numerous statistically non-significant findings for the differences between indirect and direct threats in terms of the frequency of the following linguistic features:

- Word count
- LIWC summary variables
- Posemo/Negemo

- Anxiety and sadness
- Certainty and tentative
- Past and present tense

This said, overall, indirect and direct threats seemed to be distinguished through the presence or absence of references to violence. Indirect threats in this study were characterised by a lack of overt references to commit an act of violence. The phrasing of these threats were often ominous and seemed to imply that some further action could occur. In contrast, direct threats typically included clear references to commit an act of harm and/or violence. These findings seem to corroborate the framework presented in Chapter 3.1, at least in terms of how lay-people seem to conceptualise these types of threats. Further perception research would provide greater insight into how lay-people (or threat assessment practitioners) parse these types of threats, and whether indirect and direct threats are associated with increased or decreased risk of harm.

As it stands, it is entirely possible that some number of these messages might be considered by threat assessment practitioners as being relatively low-risk. This is because the authors are not necessarily expressing a clear plan regarding where the harm is alleged to be taking place, when this plan will be enacted, or what this plan actually entails. In addition, the identity of the target and author are not always clearly stated. By not providing this level of detail in their threats, threat assessors (expert or otherwise) might infer that the threatener does not have an actual plan to bring about harm. As such, the threat might be viewed by assessors as being unlikely to be realised. Lay-person assessment of risk from these written texts will be explored in more detail in upcoming research, but it would also be useful to compare lay-person assessments with those from practitioners who are trained in threat assessment.

This chapter examined the form a threat can take, that is to say, the words that are used to convey a threatening intent (albeit in an experimental setting). The following chapter intends to address how such threats are phonetically realised.

5 Results - Spoken data analysis

As noted in §2.3, there is currently very little known (or predicted) regarding the phonetic production of threats. As such, this research adopts a broader analytical approach, with a view to making generalised observations about threatening speech. This chapter will begin by examining the phonetic production of threats taken from authentic data relative to non-threatening speech samples. After this, the results of phonetic analysis using simulated threatening speech recordings is then presented, along with the rationale for analysing both authentic and simulated data. Finally, the phonetic productions of sections of text which had been written using additional typographical features will be examined.

5.1 Authentic threat data

This subsection details the results of phonetic analysis of the authentic materials collected for this thesis. A transcript of each of these recordings can be found in the Appendix. To ensure anonymity, each speaker in this authentic data group will be referred to as 'A(authentic speaker)N.' For example, 'A1,' corresponds to 'Authentic speaker 1' and 'A2' corresponds to 'Authentic speaker 2.'

These data were analysed for two phonetic features: fundamental frequency (f0) and articulation rate (AR). These results are presented as comparisons between baseline speech and direct threats. Intensity was not analysed due to this feature being so heavily-effected by the speaker's proximity to the recording source. The analysis presented in the following sections intended to address the following research questions:

- **Research question 1**: What phonetic features can be associated with a threatening tone of voice?

Following the presentation of results relating to these phonetic features, these results will be summarised and discussed in relation to the above research questions. The following section presents the results of fundamental frequency (f0) analysis on the authentic threat data collected for this research.

5.1.1 Fundamental frequency

Figure 21 shows the results of f0 analysis for each authentic speakers' direct threat and his corresponding baseline speech. The central horizontal line in each boxplot represents the median f0 value, while the lower and upper horizontal lines for each box represent the interquartile range. The lower and upper whiskers represent the absolute minimum and maximum f0 values recorded for each speech sample.

Figure 21: F0 values extracted from each authentic recording during direct threats (grey boxplots) and baseline speech (white boxplots)

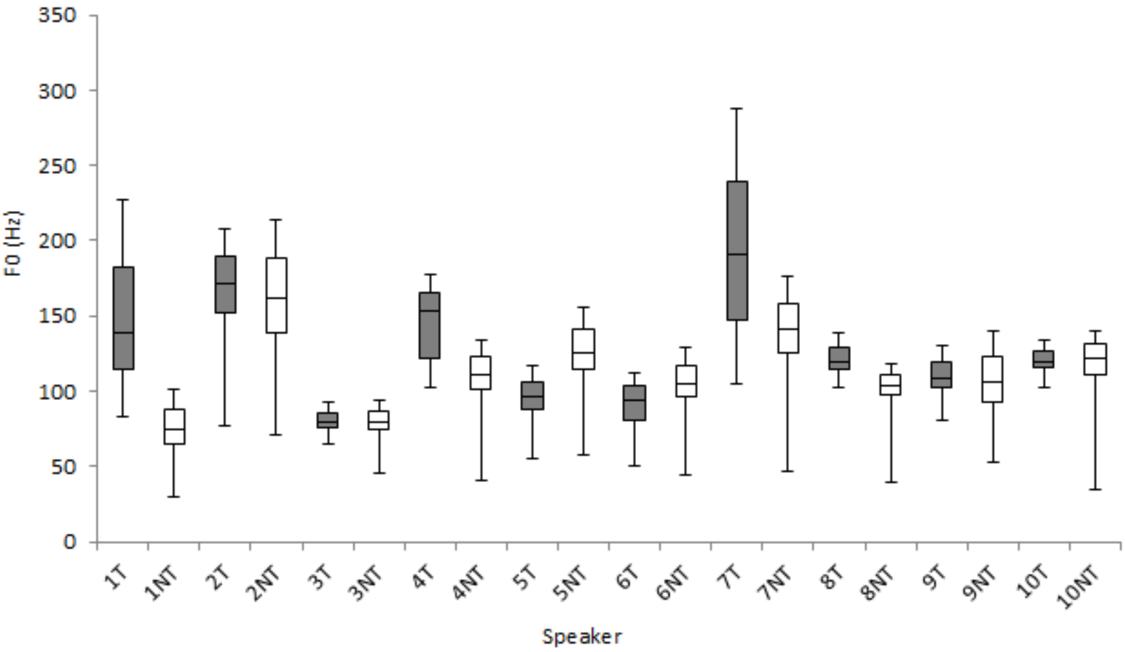

Across all 10 speakers during both speech conditions, there appears to be substantial variability across median, minimum and maximum f0 values. The baseline speech samples for all of these speakers is within the expected f0 mean range for male English speech (around 90 to 130Hz; Takefuta, Jancosek & Brunt, 1972). A2 and A7 both produced their baseline speech with a higher than expected mean f0 value for a male speaker. A1 and A3 produced their baseline speech with a lower than expected mean f0 value for a male speaker. Only A3 produced a threat with a mean f0 which decreased from his baseline speech.

Figure 21 indicates that overall, direct threats were produced with a slight increase in median f0 relative to the corresponding baseline speech. That said, a Wilcoxon signed rank test revealed no statistically significant difference between the mean f0 of the direct threat and baseline speech data ($Z = -1.27$, $p = 0.20$). This finding suggests that there is no overall tendency for these speakers to realise direct threats with a notable increase or decrease in mean f0 relative to their baseline speech.

However, f0 range was found to be statistically significant (Z = -2.50, p<.013). With the exception of Speaker 7, the f0 range was generally larger for non-threatening speech, in relation to the corresponding direct threat. This increase in f0 range is likely to be due to the generally longer length of the baseline speech samples. Longer sections of speech might allow for more opportunity for variations in volume and/or voice quality to occur. These factors in turn, can affect measurements relating to f0.

5.1.2 Articulation Rate

Figure 22 outlines the results of articulation rate (AR) analysis across all authentic threat materials.

Figure 22: Articulation rate for the authentic direct threat data (black) and corresponding baseline speech data (white)

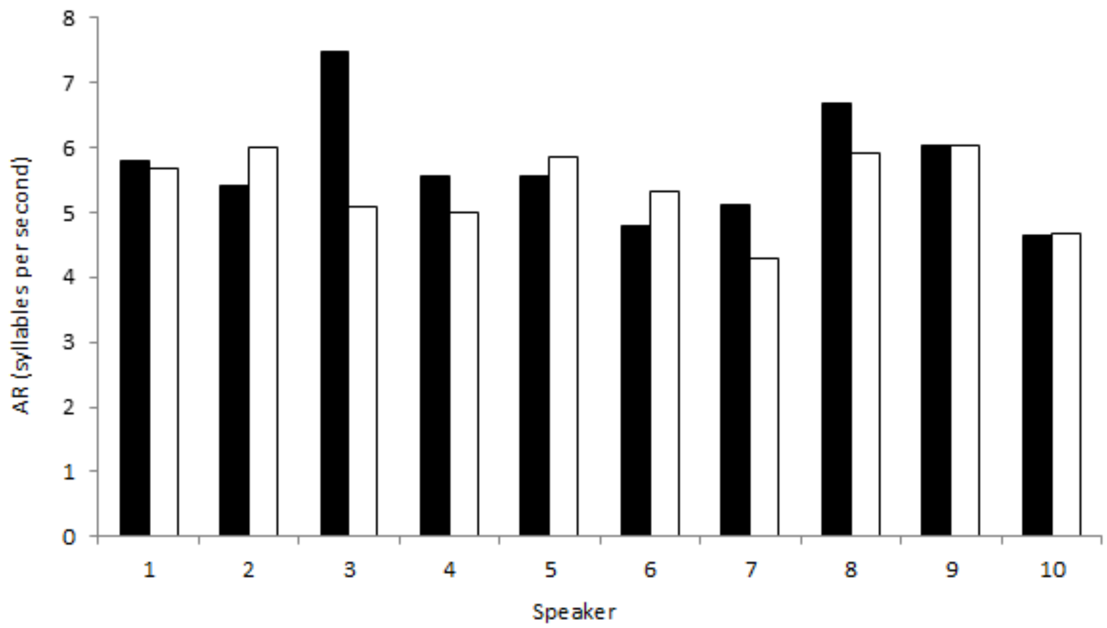

Generally, these speakers produced threats at a slightly higher raw articulation rate than their baseline speech sample. This indicates that these speakers produced their threat with a small increase in tempo relative to their baseline speech. However, statistical testing between the direct threat and baseline speech data did not show any significant result (Z = -1.07, W = 13.5). As such, there was no tendency for threats to be produced with a significant increase or decrease in AR relative to non-threatening, baseline speech.

Instead, most speech samples (irrespective of whether they were threatening or not), fell within expected values of AR for fluent, English-language speech (around 4-6 syllables per second). This excludes A3, who produced his direct threat with a notably higher AR than the other speakers (7.48 syllables per second). Due to the brevity of the threatening speech samples analysed, these results should be interpreted with caution. As with other phonetic measures, more speech data would allow for more robust generalisations to be made.

5.1.3 Summary and discussion of authentic threat analysis

This subsection sought to examine whether there was any phonetic basis for a 'threatening tone of voice.' There was a focus on examining changes to fundamental frequency and articulation rate between baseline, non-threatening speech and direct threats. These phonetic features were selected as they have been previously linked to threatening vocalisations. In addition, these voice features are known to be salient to lay-people (Banse and Scherer, 1996). As such, speakers might consciously (or unconsciously) modify these aspects of his/her vocal production in order to convey a threat.

These data suggest that speakers produce direct threats in 'real-world' interactions inconsistently with one another. That is to say, these data do not lend much in the way of support in favour of a specific 'threatening tone of voice,' in terms of f0 and articulation rate (i.e. broad phonetic cues).

5.2 Spoken data analysis: simulated data

The following subsection details the results of phonetic analysis of baseline speech, indirect threats and direct threats for the simulated speech materials. The following phonetic features were selected for analysis on the basis of previous relevant research or threat assessment materials:

- fundamental frequency (f0)
- intensity
- articulation rate (AR)

- voice profile analysis

For each of these phonetic features, male speech data is presented first followed by female speech data. Within each section, further comparisons will be made between the speech of non-actors and actors. The results presented in these subsections intend to address the following research questions:

- **Research question 1**: What phonetic features can be associated with a threatening tone of voice?
- **Research question 1.b**: Are there phonetic features which differ between the production of direct and indirect threats? Is there a compensatory effect present during the production of indirect threats?
- **Research question 1.d**: Do these phonetic features differ between speakers with acting experience, and speakers with no formal acting experience?
- **Research question 1.e**: Do these phonetic features differ between male and female speakers?

Following discussion of these findings, 'angry' speech data recorded as part of Task 4 of the experiment will be presented. These results will contribute towards addressing **research question 1.a**: ('Are the phonetic features under investigation in this study, comparable to previous descriptions of emotional speech?).

After presenting an analysis of these angry speech data, a comparison of the authentic and simulated threats will be presented. More specifically, the male Task 5 simulated data collected for this study will be discussed in relation to the authentic threats. In this task, the participants read aloud scripts transcribed from the authentic data presented in this research. These authentic and simulated data are compared as they were produced by male speakers producing the same linguistic content. The findings of this subsection aim to address **research question 1.c**: 'Do (these) phonetic features differ between authentic and simulated threats? The phonetic similarities and differences between these data will be discussed.

The following section presents the results of fundamental frequency (f0) analysis performed on the simulated threat data collected for this research. This section will present f0 data from the following experimental tasks:

- Task 1 (baseline read speech)

- Task 2 (indirect threat scripted by the speaker)
- Task 3 (direct threat scripted by the speaker)

Each participant recorded for this study is referred to as 'S(speaker)N'. For example, 'S1' corresponds to 'Speaker 1', and 'S2' corresponds to 'Speaker 2'.

5.2.1 Fundamental frequency analysis

An initial repeated measures ANOVA test showed no statistical significance for the mean f0 values recorded for all of the speech produced by the speaker groups recorded for this study (male non-actors, male actors, female non-actors, female actors - with the model design: Intercept + Condition + Sex + Acting + Condition*Sex*Acting using SPSS software). This was reported at $F(1.42, 52.62) = .47$, $p = .56$. There was no significant main effect of any of the individual factors, nor any significant interactions (see Appendix item 'Tests of Within-Subjects Effects'). The following section details a further analysis into each of these groups, beginning with the f0 results pertaining to the male participants. This was in order to ascertain whether there were any within-group tendencies for the production of mean f0 across different the experimental conditions.

Male speakers

Chapter 2 discussed how physical changes to the male vocal anatomy around puberty have an effect on mean f0 for male speakers. Connected with this is the idea that these changes occur so that a male speaker can suggest to others that he is physically imposing. Theories related to the so-called 'frequency code' suggest that large entities are associated with loud, low-pitched noises, and smaller entities are associated with quieter and/or high-pitched noises. In turn, larger creatures are perceived by members of the same (or different) species to be more of a threat. This is due to physical size being linked with the perception that violence could be actualised (Ohala, 1984). As such, human male speakers who adopt threatening behaviour might be expected to produce lower-pitched speech. By doing so, males can apparently make themselves appear physically large to others and therefore, be perceived as a greater threat by others.

The figure below (23) details the median f0 and f0 range for all male speakers (Speakers 1-10 and Speakers 21-30) for each of the three experimental conditions (baseline speech, indirect threats and direct threats).

Figure 23: F0 for each condition produced by all male speakers (n=20)

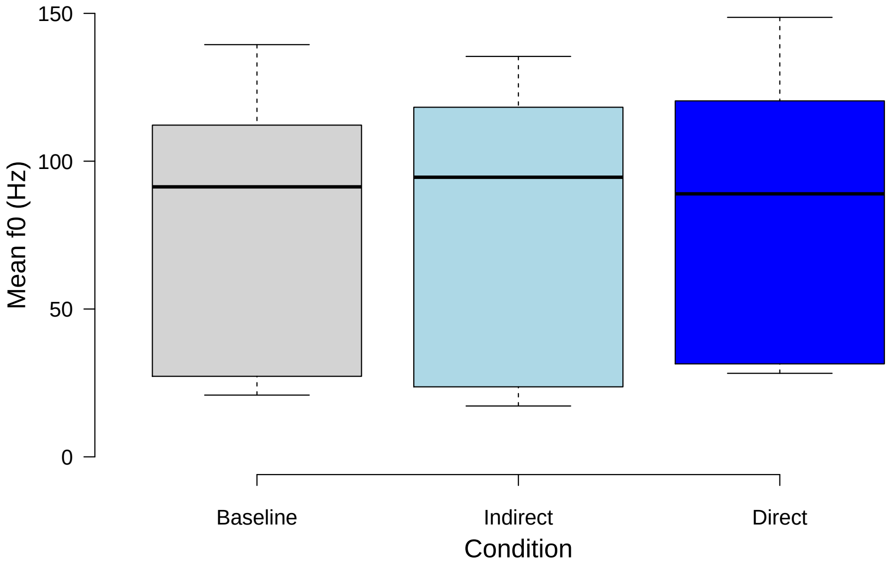

Contrary to these theories, this figure appears to suggest a slight increase in median f0 for the indirect threat condition (relative to baseline speech). However, a Wilcoxon signed-rank test for these data yielded no significant differences between the mean f0 of male speakers producing indirect ($Z = -1.00$, $p = 0.31$) or direct threats ($Z = -1.31$, $p = 0.19$) and their baseline speech. As described in **research question 1.d**, another aspect of this research is whether there are differences between non-actors and actors regarding their phonetic production. To this end, Figures 24 and 25 show the differences in mean f0 produced by the male non-actors recorded for this study (Speakers 1-10). These figures show the mean f0 for each of the three experimental conditions (baseline speech, indirect threats and direct threats).

Figure 24: Mean f0 for each condition produced by male non-actors

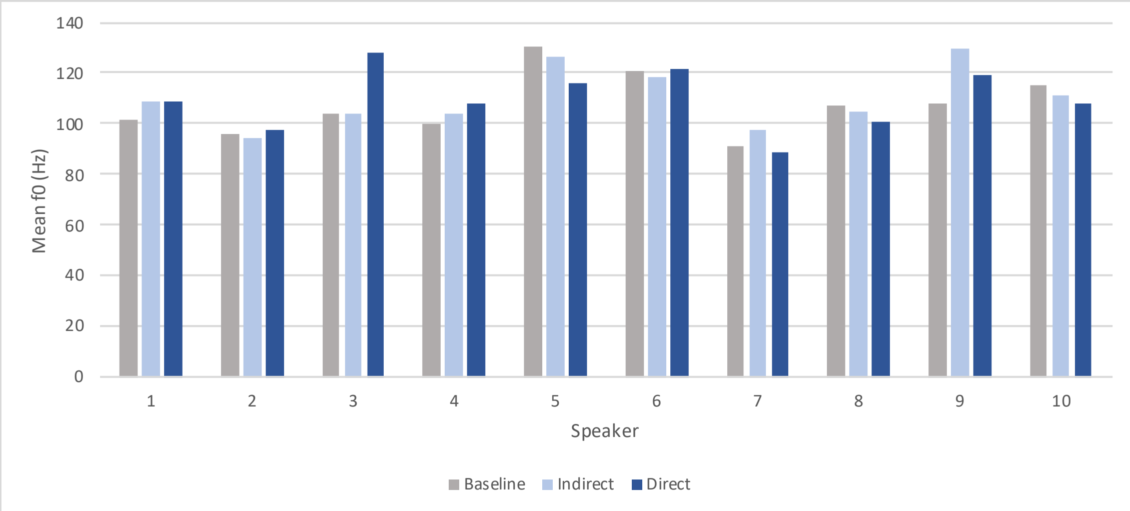

Figure 25: Change in mean f0 for direct and indirect threat conditions for male non-actors from their baseline speech

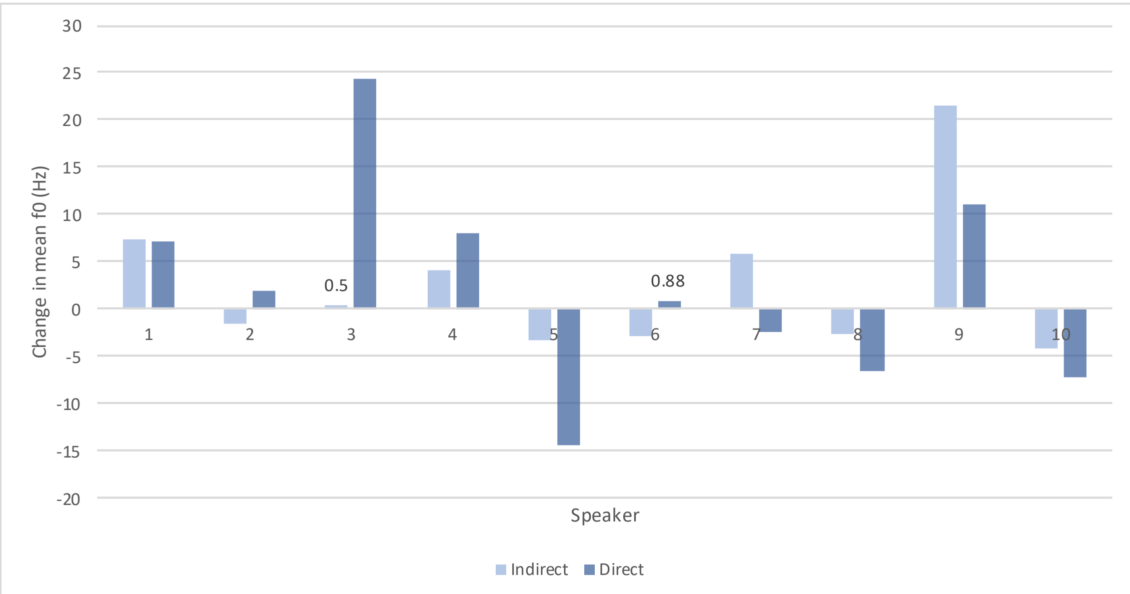

Based on these figures, there appears to be no overall tendency for male non-actors to produce threats with an increase or decrease in mean f0 (relative to their baseline speech). Speakers 1, 3, 4 and 9 deliver threats with an increase in f0 to varying degrees. Of these speakers, 3 and 4 produce direct threats with a greater mean f0 increase than indirect threats. Speaker 1 produces indirect and indirect threats with highly similar f0 means to one another. Both of these productions also show a small increase in mean f0 from Speaker 1's baseline speech. Speakers 5, 8 and 10 produce their threats with a decrease in mean f0, relative to their baseline speech. For these speakers, all produce direct threats with a lower mean f0 than indirect threats.

The remaining speakers (2, 6 and 7), produced their indirect and direct threats inconsistently with respect to whether these threats were delivered with an increase or decrease in mean f0. Speakers 2 and 6 show a slight increase in mean f0 in their direct threats and a slight decrease in mean f0 in their indirect threats. Speaker 7 produced his indirect threat with an increased mean f0 relative to his baseline speech and direct threat. Overall, there appears to be a tendency for direct threats to be produced with a greater increase or decrease to mean f0 than indirect threats. As described previously, statistical testing between the male non-actors' baseline mean f0 and indirect threats did not show a significant difference. In addition, there was no statistical difference between the male non-actors' baseline mean f0 and direct threats. These results indicate that there is no obvious pattern for male non-actors to produce threatening speech (of either type) with a difference in mean f0 from their baseline speech.

Figures 26 and 27 show the differences in mean f0 across the three experimental conditions as produced by male actors (Speakers 21-30).

Figure 26: Mean f0 for each condition produced by male actors

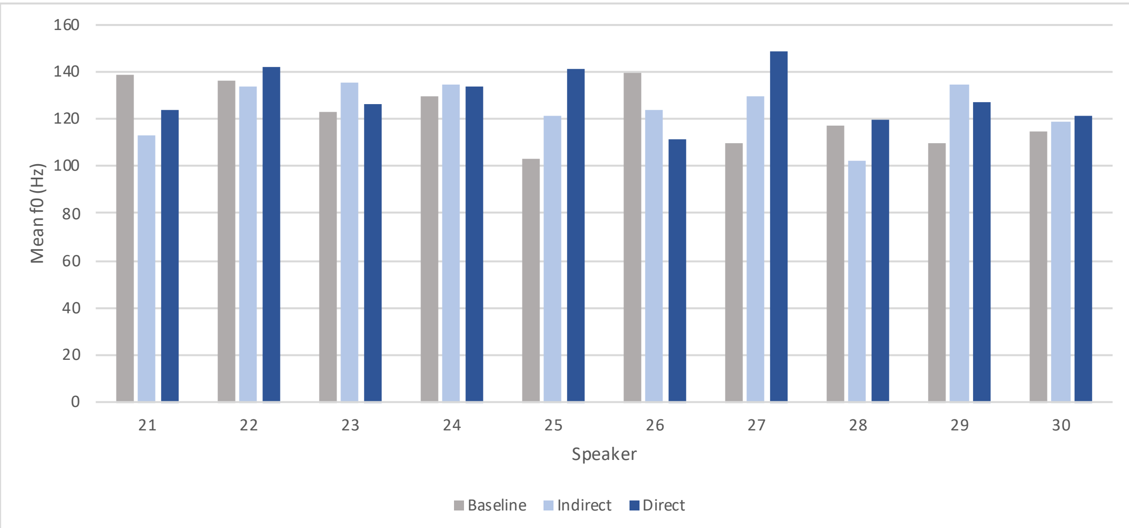

Figure 27: Change in mean f0 for direct and indirect threat conditions for male actors from their baseline speech

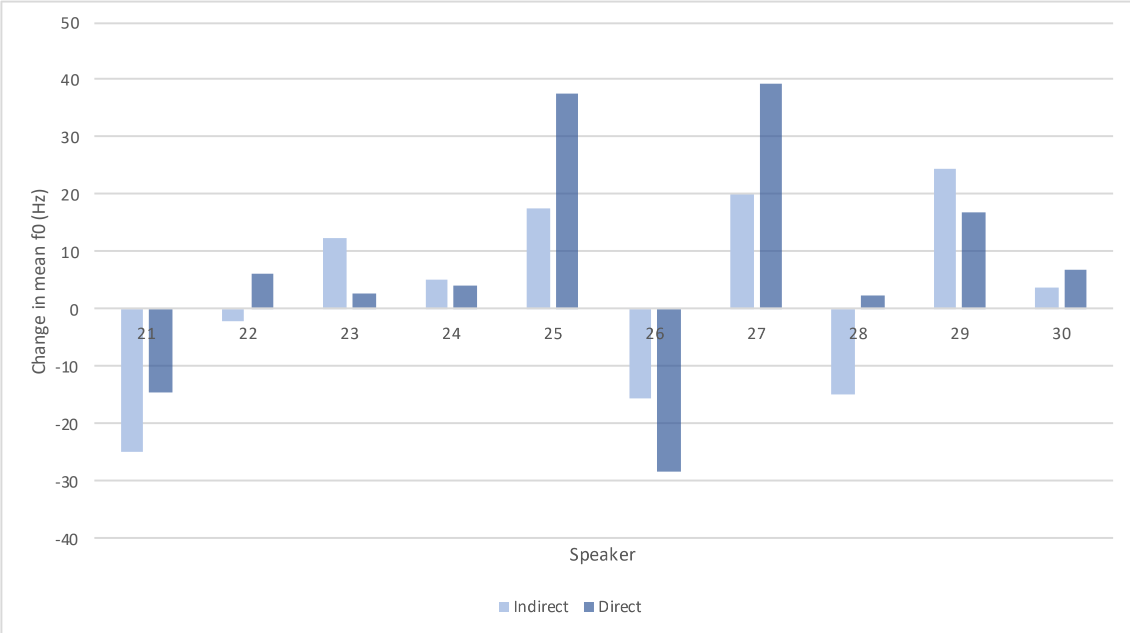

As with the male non-actors, the actor group also shows variable productions of mean f0 for each condition. Overall, (based on a visual inspection of Figure 27), the speakers in this group mostly appear to increase their raw mean f0 when producing a threat (of either type) relative to their baseline speech. Speakers 23, 24, 25, 27, 29 and 30 all follow this trend. Of these speakers, half (Speakers 25, 27 and 30) produce direct threats with a higher raw mean f0 than their corresponding indirect threats. In contrast, Speakers 23, 24 and 29 produce their indirect threats with a higher mean f0 than their corresponding direct threats.

There are also male actors who produced their threats with a decreased mean f0, relative to their baseline speech. This was the case for Speakers 21 and 26. Finally, Speakers 22 and 28 both produced their indirect threats with a decrease in mean f0 (from their baseline), and their direct threats with an increase in mean f0. In both cases, the extent of this increase or decrease varied. Statistical testing between the male actor's baseline mean f0 and indirect threats did not show a significant difference ($Z = -0.56$, $p = 0.58$). In addition, there was no statistically significant difference between the male actors' baseline mean f0 and direct threats ($Z = -1.38$, $p = 0.17$).

Therefore, the results of this analysis do not provide clear evidence to support any generalisation about how the production of threatening speech might differ from baseline speech, in terms of mean f0.

Also of interest here is whether non-actors and actors consistently produce differences in mean f0 across the three experimental conditions. A series of Mann-Whitney U tests were performed based on the data provided for each of the three conditions by male non-actors and male actors. There are significant differences between the mean f0 values for male non-actors and actors regarding their baseline speech (U=19, p = 0.019), indirect threats (U=18, 0.016), and direct threats (U=11, p = 0.003). Further examination of Figures 24 and 26 suggests that male actors typically produced all of their speech with a higher mean f0 than male non-actors, irrespective of the experimental condition.

A series of Friedman tests were run to check whether there were within-group differences in the mean f0 of male non-actors and actors in speech produced in each of the three conditions. However, these did not show any significant results (reported as $\chi^2 = 0.2$, p = 0.90 for the male non-actors, and as $\chi^2 = 3.2$, p = 0.20 for the male actors). As such, male speakers (within either acting group) did not appear to produce speech in each of the three experimental conditions with significant differences in mean f0.

No predictions from previous literature have been made about the f0 variability of threatening speech. The following figures (Figures 28 and 29), visualise the standard deviation of mean f0 across the three conditions for the male non-actors and actors. These figures represent the degree of mean f0 variation between the speakers of these groups for each experimental condition. The numbers which accompany each bar represent the degree of increase or decrease in f0 standard deviation as a percentage. For example, Speaker 1's indirect threat was produced with a decrease in f0 variation of 14.74%, compared to his baseline speech.

Figure 28: Change in standard deviation of mean f0 between baseline speech and threats produced by male non-actors

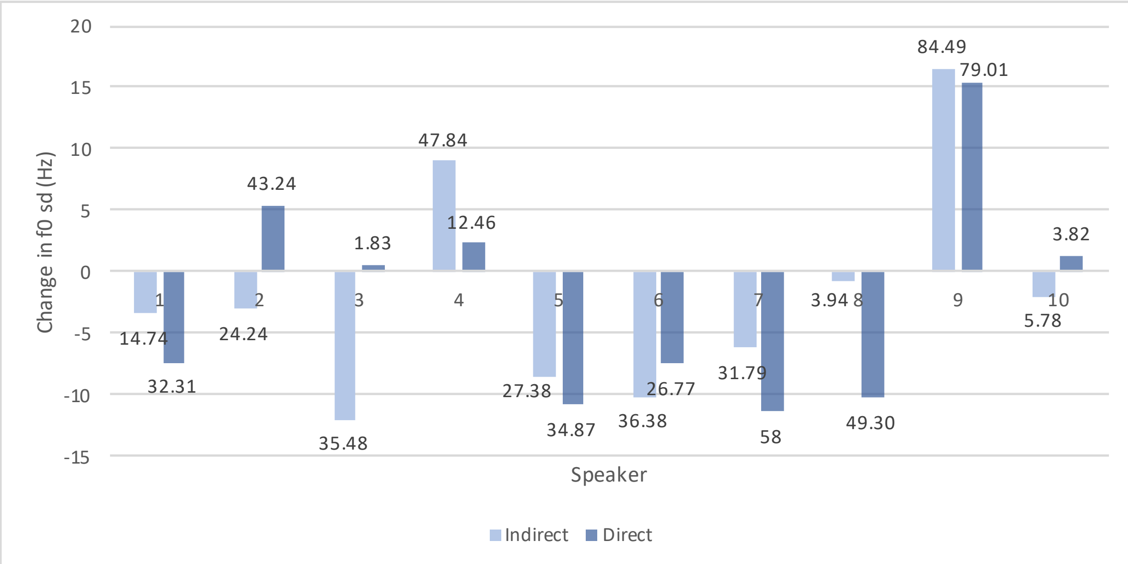

Figure 29: Change in standard deviation of mean f0 between baseline speech and threats produced by male actors

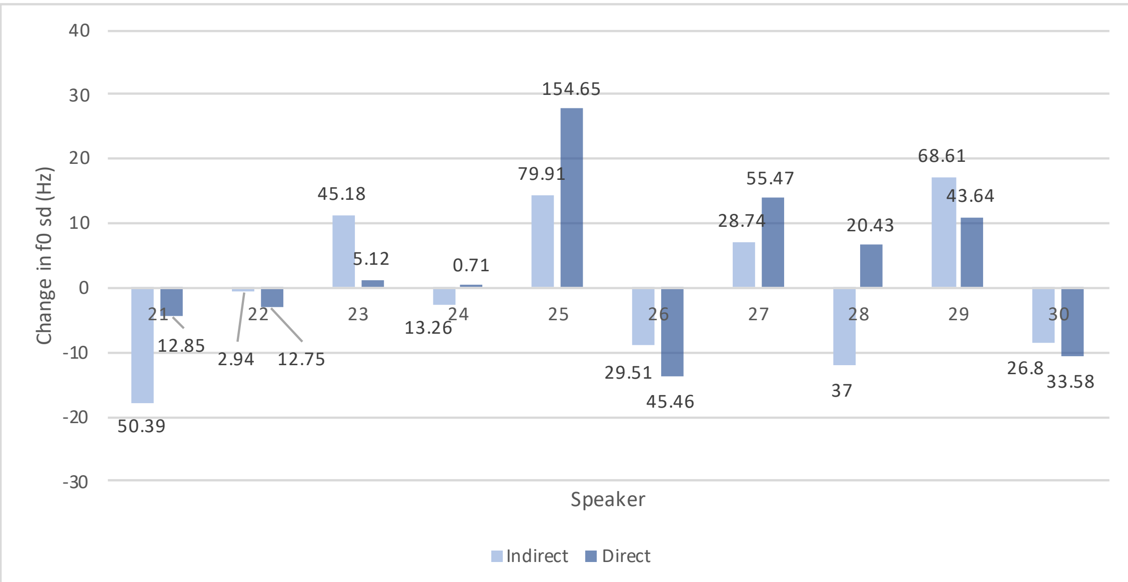

These figures suggest that overall, male speakers decreased f0 variation when producing threatening speech. That said, for Speakers 22 and 24 their extent of their decrease was considerably lower relative to other male speakers. It is less apparent whether speakers who decreased their f0 variation when making threats produced either indirect or direct threats with greater levels of decrease than the other. For Speakers 1, 5, 7, 8, 26 and 30, direct threats showed the greatest decrease in f0 variability.

Other speakers produced their threats with an increased change in f0 variation. This level of increase was also inconsistent across both non-actor and actor groups. In particular, Speakers 9 and 25 showed notably higher levels of f0 variation when making threats. Speakers 2 and 10 both decreased their f0 variation for indirect threats, and increased this variation when making direct threats.

A series of Spearman's Rank-Order correlation tests show a largely-positive relationship for male speakers between change in mean f0, and change in standard deviation. For indirect threats, this correlation was reported at $r_s = .740$, $p=<.05$. For direct threats, this correlation was reported at $r_s = .669$, $p=<.05$. In other words, increases in fundamental frequency are associated with increases in f0 standard deviation.

Summary of male speaker fundamental frequency analysis

The results of the fundamental frequency analysis of the male speech collected for this thesis can be summarised as follows:

- There is no statistical difference between the mean f0 of male participants producing speech for each condition.
- There is also no statistical difference in the mean f0 of speech taken from either male actors or male non-actors producing threats of either type (relative to their baseline speech).
- Generally, the male actors consistently produce their speech (across all conditions) at a higher mean f0 than male non-actors.
- There is evidence to suggest that male speakers overall, produce threatening speech with a decrease in f0 variation. In addition, there appears to be a positive relationship between change in f0 mean and standard deviation.

Female speakers

As detailed in Chapter 2, there are no published predictions or studies which account for the f0 of threatening speech produced by female speakers. Therefore, these data will be later discussed in relation to the predictions made about male threatening speech, and to the theories outlined in Chapter 2 which claim that threats are either a learned or an inherent behaviour.

The figure below (30) shows the median f0 and f0 range for all female speakers (Speakers 11-20 and Speakers 31-41) for each of the three experimental conditions (baseline speech, indirect threats and direct threats).

Figure 30: Mean f0 for each condition produced by all female speakers (n=21)

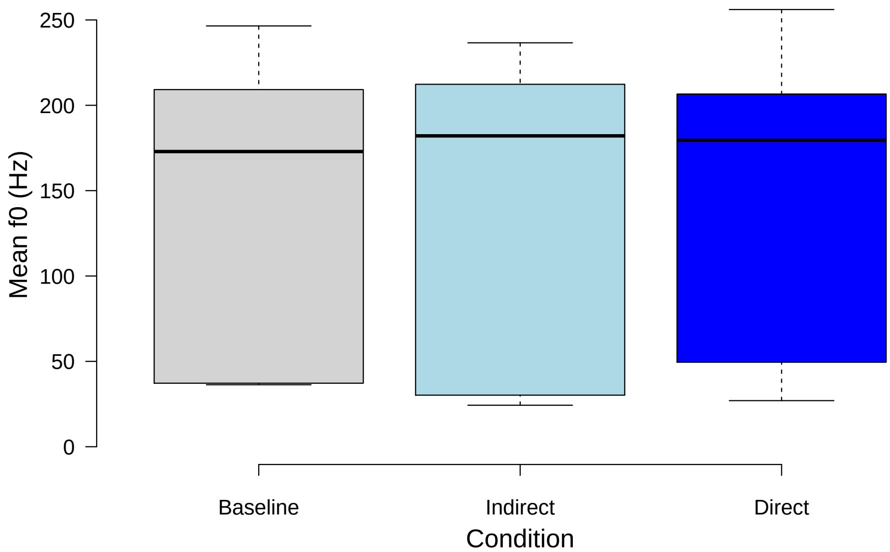

Similar to the male speakers, this image suggests a slight increase in mean f0 for the threat conditions (relative to baseline speech) as produced by female speakers. However, there appears to be little change in f0 for female speakers between the three conditions. This interpretation is supported by statistical testing (Wilcoxon signed-rank test) which yielded no significant difference between the mean f0 of female speakers producing indirect (Z =

-0.74, p = 0.45) or direct threats (Z = -0.30, p = 0.76), and their baseline speech. The following images (Figures 31 and 32) show the differences in mean f0 produced by speakers in the female non-actor group (Speakers 11-20).

Figure 31: Mean f0 for each condition produced by female non-actors

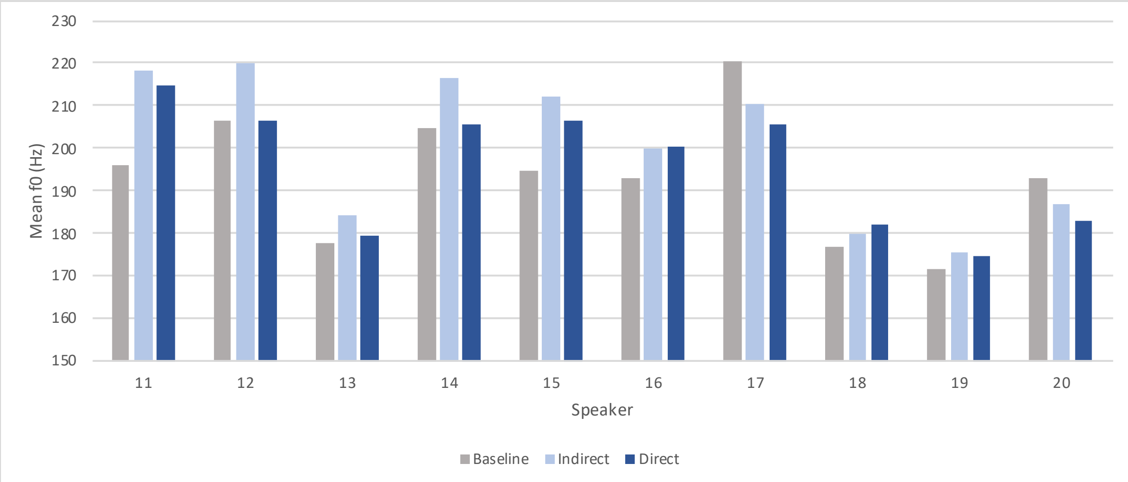

Figure 32: Change in mean f0 for direct and indirect threat conditions for female non-actors (from baseline speech)

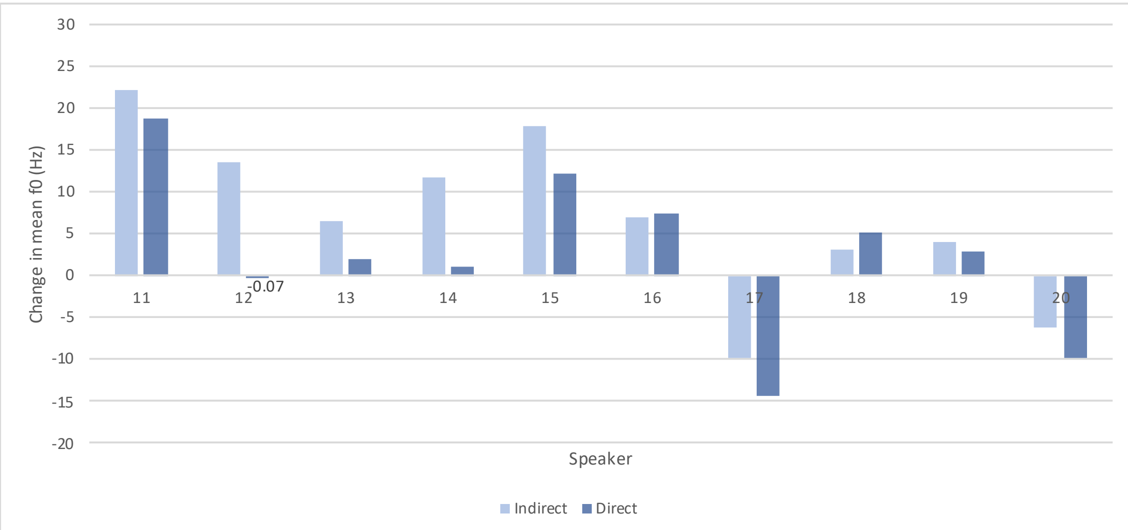

For Speakers 11, 12, 13, 14, 15, 16, 18 and 19, the raw value for mean f0 during threatening speech is higher than their corresponding baseline speech. Of these speakers, all but Speaker 18 produce their indirect threats at a higher raw mean f0 than their direct threats. Interestingly, Speaker 12 produces her direct threat at a similar mean f0 to her baseline speech. For the remaining speakers (Speakers 17 and 20), threatening speech appears to be produced with a decrease in mean f0. In addition, these speakers appear to produce their direct threats at a slightly lower f0 than indirect threats. However, there is no statistically significant difference between the female non-actors' baseline mean f0s and the f0s their corresponding direct threats ($Z = -1.07$, $p = 0.28$). Further statistical testing also did not show a significant difference between the female non-actors' mean f0 for baseline speech and indirect threats ($Z = -1.89$, $p = 0.79$).

These results suggest that while some female non-actors appear to raise their mean f0 when producing a threat, this is by no means consistent across the group. In addition, the differences between baseline speech and threatening speech (of either type) are not statistically significant. In other words, there is little evidence to suggest that female non-actors (overall) produced indirect and direct threats with changes from their baseline mean f0.

The following images (Figures 33 and 34) show the differences in mean f0 produced by female actors (Speakers 31-41).

Figure 33: Mean f0 for each condition produced by female actors

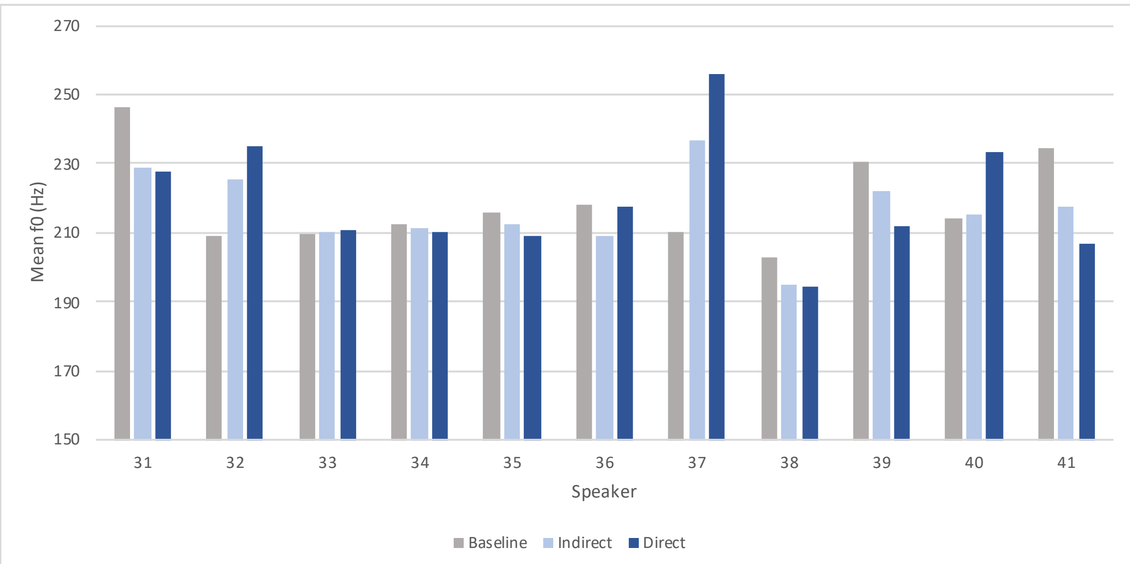

Figure 34: Change in mean f0 for direct and indirect threat conditions for female actors relative to their baseline speech

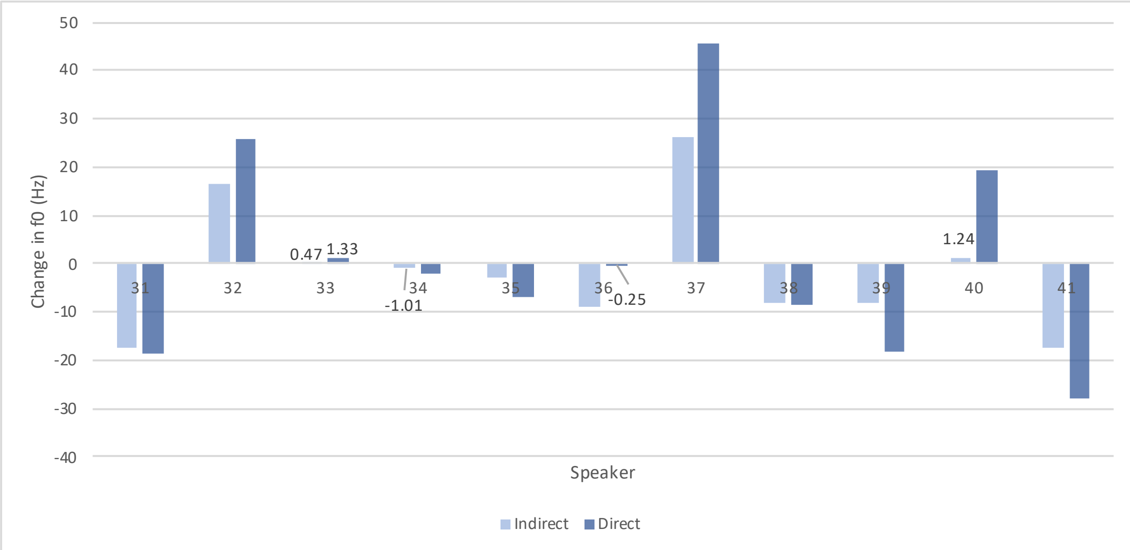

According to these figures, some female actors appear to lower their mean f0 relative to their baseline speech when uttering threats. However, this decrease is often small, as shown in Figure 5 for Speakers 33, 34, 35, and 36. Among these speakers, there does not seem to be a tendency for direct threats to be lower in mean f0 than indirect threats (or vice versa). Speakers 32, 33, 37 and 40 appear to show increases in mean f0 when producing their threatening speech. However, Speaker 33 produces a mean f0 that is only marginally increased from her baseline speech. The remaining speakers (32, 37 and 40) all produced direct threats at a higher raw mean f0 than their indirect threats. Statistical testing between the female actors' baseline mean f0 and direct threats mean f0 did not reveal a significant difference (Z = -0.27, p = 0.79). Nor did statistical testing show a significant difference between the female actors' mean f0 between baseline speech and indirect threats (Z = -0.89, p =0.37).

A closer examination of Figures 31 and 33 suggests that female actors consistently produced higher mean f0 than female non-actors across the three conditions. This is mostly supported by a series of Mann-Whitney U tests comparing the mean f0s of speech data provided by female actors and non-actors. Baseline speech was reported as significantly different (U = 10, p = 0.002) between the female non-actors and actors. Direct threats between the same non-actors and actors were also significantly different (U = 12, p = 0.002). Additionally, a significant difference was observed for indirect threats (U = 65, p = 0.057). These results indicate that the female actor group (as a whole) consistently produced their speech with a difference in mean f0 compared to the female non-actors. These differences seem to reflect a tendency for female actors to produce a higher mean f0 than female non-actors irrespective of the experiment condition.

A series of Friedman tests showed no statistical evidence to suggest that female speakers within the actor group produced a significantly different mean f0 when producing speech in any of the three conditions (χ^2 = 6.2, p = .045). However, there was a significant difference in mean f0 within the female non-actor group when producing speech in any of the three conditions (χ^2 = 1.27, p = 0.53).

The following figures (Figures 35 and 36), visualise the standard deviation of mean f0 across the three conditions. These figures show the degree of mean f0 variation between the speech of female non-actors and actors for each condition.

Figure 35: Change in standard deviation of mean f0 between baseline speech and threats produced by female non-actors

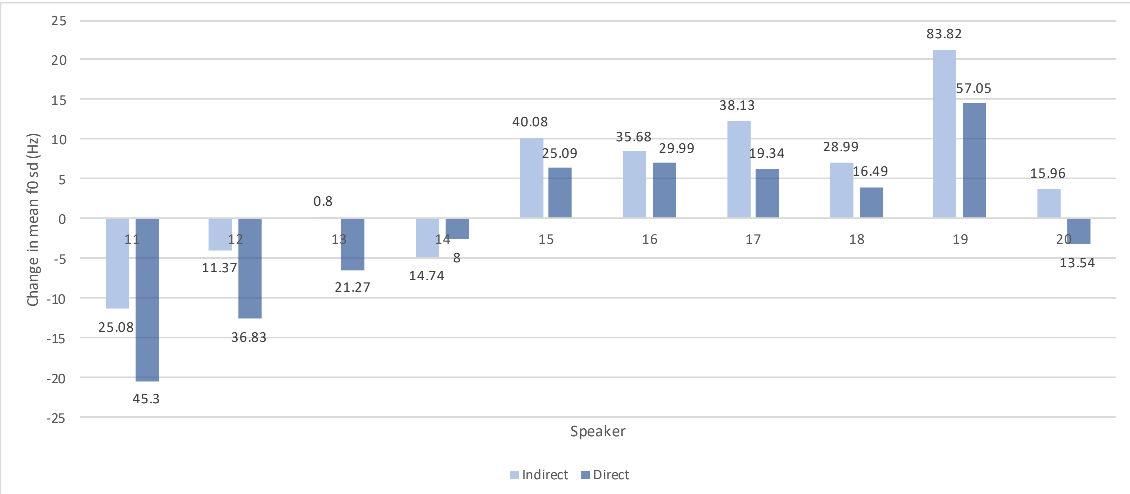

Figure 36: Change in standard deviation of mean f0 between baseline speech and threats produced by female actors

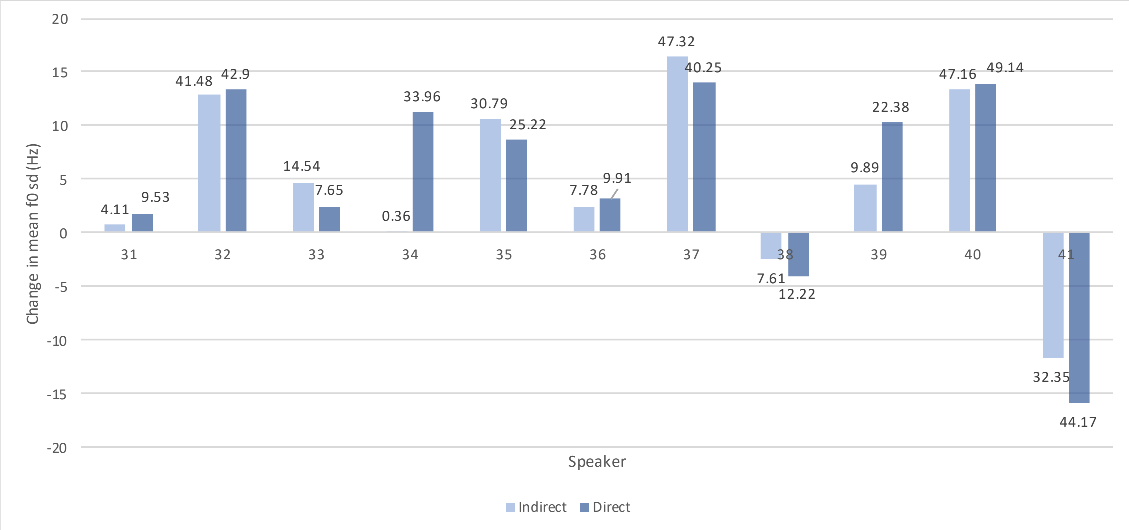

As with the corresponding figures for the male data shown in the previous subsection, the numbers displayed next to each bar represent the percentage of increase or decrease relative to that speaker's baseline speech. According to these figures (Figures 35 and 36), some of the speakers in both groups appear to produce threats with an increase in f0 variation. It is less clear whether indirect or direct threats show the greatest amount of variation. For the female non-actors who increase their f0 variation (relative to their baseline), indirect threats show more f0 variation than indirect threats. This is seen for Speakers 15, 16, 17, 18, and 19.

For the female actors, there seems to be an opposite pattern. Instead, these speakers appear to produce direct threats with a greater f0 variation than is the case for their indirect threats. This is seen for Speakers 31, 32, 34, 36, 39, and 40. Speakers 33, 35, and 37 increase f0 variation when making indirect threats.

The remaining female non-actors and actors appeared to decrease their level of f0 variation when producing a threat. For these speakers, direct threats tend to exhibit a greater decrease in variation (from baseline speech) than indirect threats. This is the case for Speakers 11, 12, 13, 38, and 41. Speakers 13 and 20 are exceptions to this pattern. However, for female non-actors, the f0 standard deviation of indirect threats ($Z = -1.48$, $p = 0.14$) and direct threats ($Z = 0$, $p = 1.0$) relative to baseline speech was not statistically significant. For female actors, the f0 standard deviation of indirect threats ($Z = -1.9$, 0.06) and direct threats (Z = -1.6, p = 0.11) was also not significantly different relative to baseline speech.

Unlike the male speakers, the female speakers did not show a largely-positive correlation between change in mean f0 and change in standard deviation. For indirect threats, this relationship was statistically non-significant. For direct threats, there was a moderately strong positive correlation (r_s=.399, p =.037).

Summary of female speaker fundamental frequency analysis

The results of the fundamental frequency analysis of the female speakers can be summarised as follows:

- There is no statistical difference in mean f0 for female participants producing speech in each condition.
- There is also no statistical difference in the mean f0 of speech taken from either female actors or female non-actors producing threats of either type (relative to their baseline speech).
- The female actors (overall) consistently produce their speech (across all conditions) at a higher mean f0 than female non-actors.
- The variability of mean f0 (as measured by standard deviation in mean f0) appears to be inconsistent both within and across the actor/non-actor groups. In addition, there was little evidence to indicate a strong correlation between change in mean f0 and standard deviation.

5.2.2 Summary and discussion of fundamental frequency analysis

This subsection has provided the results of an analysis comparing the mean f0 and f0 standard deviation from baseline, non-threatening speech and with those of threatening speech. In response to **research question 1**, there is little evidence based on these data to support the idea that threats are produced with a significant change to mean f0 relative to baseline speech. In addition, the direction of these changes when producing threatening speech also varies across these speakers. In contrast to previous predictions discussed in Chapter 2, some speakers in this experiment increased his/her f0 relative to baseline speech, even if these increases were not (collectively) statistically significant. These findings are at odds with the previous theories and research put forth by Ohala (1984) and

Mileva et al. (2017). In these studies, threats were associated with decreases to mean f0. It is possible that the threats where f0 was increased were also produced with an increased intensity (Lehiste and Lass, 1976). The intensity of these threatening recordings will be described in the following subsection.

The results of this analysis would suggest that threat production is not characterised by a singular pattern of changes to f0 (from baseline speech) as could be extrapolated from the research discussed in §2.3.2. Rather, that there are possibly different strategies used by speakers to modify their f0 for the communication of threatening speech.

Research question 1.a considers whether the phonetic features observed in threatening speech are also applicable to previous descriptions of emotional speech. Increases in f0 have been associated with the expression of happiness and anger, while decreases in f0 have been associated with the expression of boredom and sadness (Johnstone and Scherer, 2000). For the data presented in this study, some speakers increased his/her f0 when delivering a threat. This result suggests that in terms of f0, threats do not necessarily resemble the phonetic features of high arousal emotional states (such as happiness and anger) more so than low arousal emotional states (such as sadness). This will be explored in greater detail in Chapter 5.3 where the production of threatening speech will be compared to that of angry speech.

Research question 1.b addresses whether there are phonetic features which differ between indirect and direct threats. In addition, whether indirect threats are produced with notable phonetic differences which might serve to distinguish them from direct threats. The analysis presented in this subsection found no consistent trend in the production of indirect and direct threats with respect to f0. There also does not seem to be evidence to suggest that indirect threats are produced with notable differences in f0, which might serve to communicate an intention to harm in the absence of words which explicitly refer to violence (described in this research as a 'compensatory effect'). As such, it seems that indirect and direct threats generally cannot be teased apart purely on the basis of f0.

This lack of consistency in the production of mean f0 also relates to the concept of the 'frequency code' discussed in Chapter 2. In particular, to the idea that the production of lower frequencies is associated with increased body size, or physical and social dominance.

In this study, the expression of threat was not associated with consistent decreases in mean f0, as might be anticipated. Instead, it would appear that threat production is a more complex process than described in previous theoretical approaches.

However, it is possible that there are finer-grained prosodic features which speakers might use to communicate an intent to harm when the words used might suggest otherwise. For example, speakers might produce a particular word or phrase with a rising or falling tone. This, in turn, might enable a threatener to allude to an alternative underlying interpretation of the utterance. More detailed prosodic analysis of threatening utterances would shed more light on this possibility.

In response to **research question 1.d**, it seemed that actors (of either sex) produced their speech with a higher mean f0 than did non-actors. Furthermore, female non-actors generally increased their mean f0 when producing a threat. For female actors, many of the speakers decreased their mean f0 when delivering threatening speech. These findings suggest that perhaps caution should be advised when selecting stimuli for perception experiments centred on threatening language. This is because listeners might base their impressions of 'threat' on differences in pitch (Mileva et al., 2017). If listeners were asked to in some way assess the threats recorded for this study, the actor data might be perceived differently to non-actor data based on the f0 differences measured in these data. This is an area that will be addressed in a future experiment.

In addition, there is no statistical difference in mean f0 across the three experimental conditions for either male or female speakers. As such, it appears that overall, speakers (irrespective of their biological sex) do not seem to utilise f0 to differentiate an 'intent to harm', from non-threatening speech. Also, no clear pattern could be established for the degree of f0 variation among the speakers. These findings relate to **research question 1.e**, which addresses whether males and females produce phonetically different threats to one another.

The following subsection presents the results of intensity analysis performed on these data.

5.2.3 Intensity analysis

Previous literature relating to threatening or aggressive speech provides mixed predictions relating to the male production of 'volume'. There is variation across studies regarding whether speaker volume increases or decreases when these emotions or behaviours are aroused. The data presented in this subsection detail the acoustic correlate of volume (i.e. intensity). This variable was possible to measure acoustically as participants wore a headset with a microphone attached. This microphone was kept at a fixed position from the speakers' mouths throughout their recording sessions. However, there will inevitably be some differences across participants concerning their proximity to the microphone (or the microphone's proximity to the participants). Therefore, comparisons across speakers might be affected by slight variations in proximity to the recording equipment, rather than similarities or differences in vocal effort. This subsection will discuss the results of intensity analysis from the following experimental tasks:

- Task 1 (baseline read speech)
- Task 2 (indirect threat scripted by the speaker)
- Task 3 (direct threat scripted by the speaker)

Male speakers

The figure below (37) shows the mean intensity (measured in decibels - 'dB') for all male non-actors (Speakers 1-10). These data are presented as changes in the production of indirect and direct threats relative to their corresponding baseline speech.

Figure 37: Change in mean intensity for indirect and direct threat conditions for male non-actors relative to their baseline speech

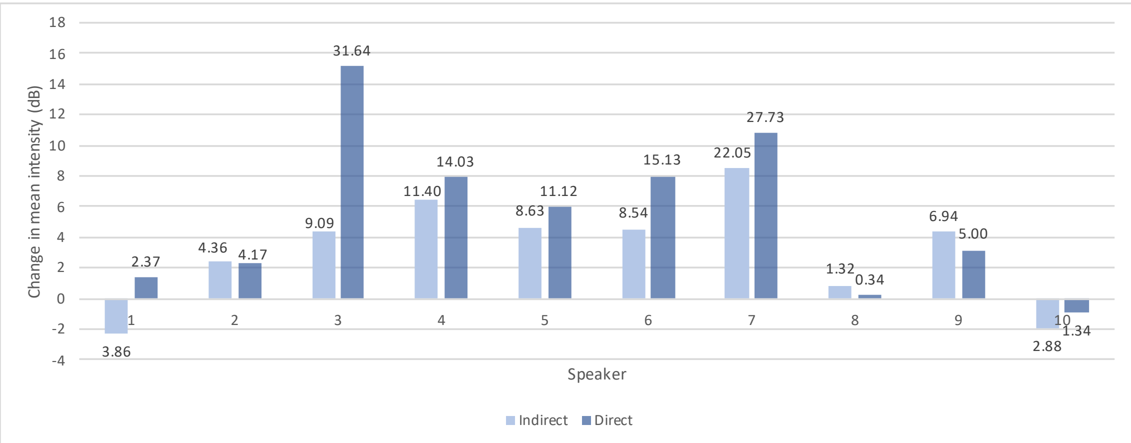

In this figure, the numbers accompanying each bar represent a percentage change from the speaker's baseline speech. For example, Speaker 1 produced his indirect threat with a 3.86% intensity decrease from his baseline speech. His direct threat, however, was produced with a 2.37% increase from his baseline speech. Overall, the male non-actors appear to produce both threats with an increase in intensity. For Speakers 1, 3, 4, 5, 6 and 7, direct threats are produced with an increase in intensity over indirect threats. For Speaker 3, this increase is notably more pronounced. The opposite is seen for Speaker 10, who decreases the intensity of both his indirect and direct threat from his baseline speech. However, his direct threat is produced with a higher intensity than his indirect threat. Speakers 2 and 8 produce their threats with similar intensity increases relative to their baseline speech.

As illustrated in the next figure (38), the standard deviation of intensity appears to be similar across the male non-actors.

Figure 38: Standard deviation of intensity across all experimental conditions for male non-actors

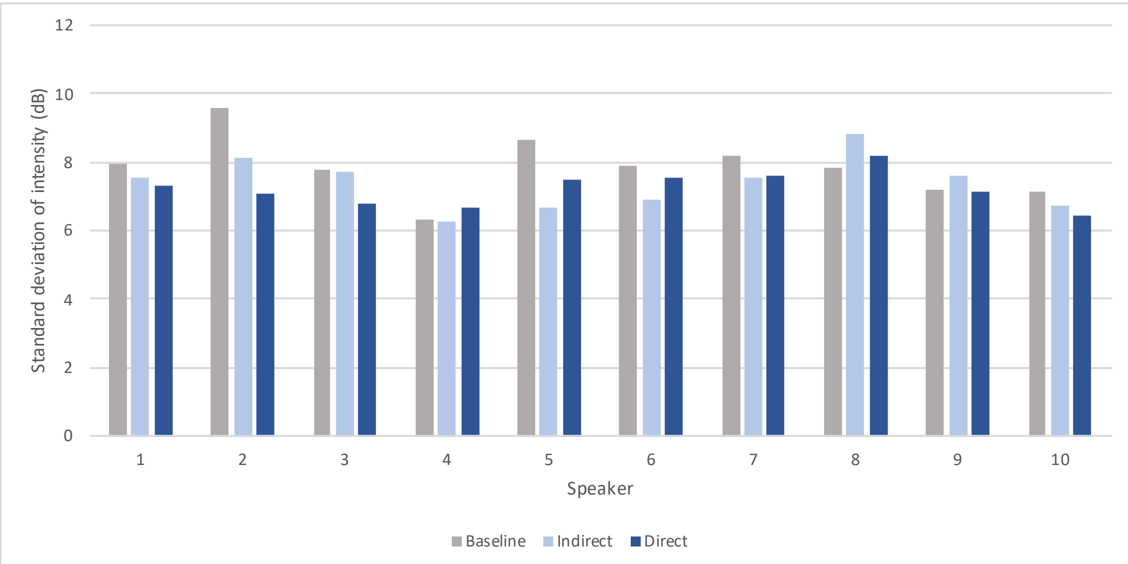

The following figure (39) shows the differences in mean intensity produced by male actors (Speakers 21-30). As before, the numbers positioned above or below each bar represent a percentage increase or decrease in intensity from the speaker's baseline speech.

Compared to male non-actors, the male actors appeared to be less consistent regarding changes in intensity between threatening and baseline speech. Speakers 22, 23, 27 and

Figure 39: Change in mean intensity for indirect and direct threat conditions for male actors from their baseline speech

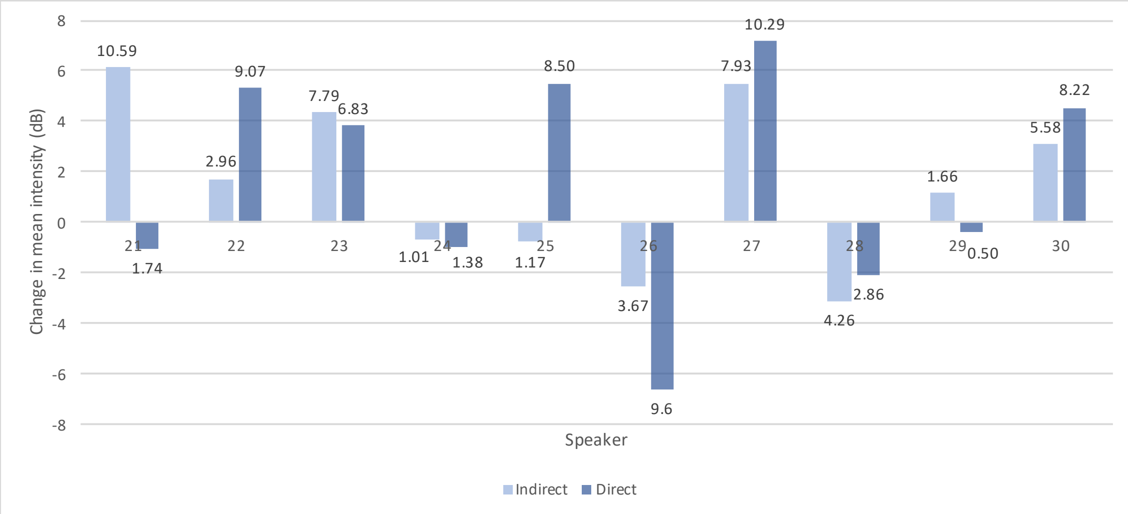

30 each produced their threats with increased intensity relative to their baseline speech. Except for Speaker 23, direct threats were produced with a greater intensity than indirect threats. Speakers 24, 26 and 28 decreased their intensity when producing threatening speech, relative to their baseline speech. Speaker 21 produced his indirect threat with a notable increase in intensity, compared to his direct threat.

The figure below (40) shows the standard deviation of intensity of the male actors. Overall, there appears to be little variation concerning standard deviation. However, in relation to other speakers, Speaker 21 appears to vary his intensity more during his direct threat production.

Figure 40: Standard deviation of intensity across all experimental conditions for male actors

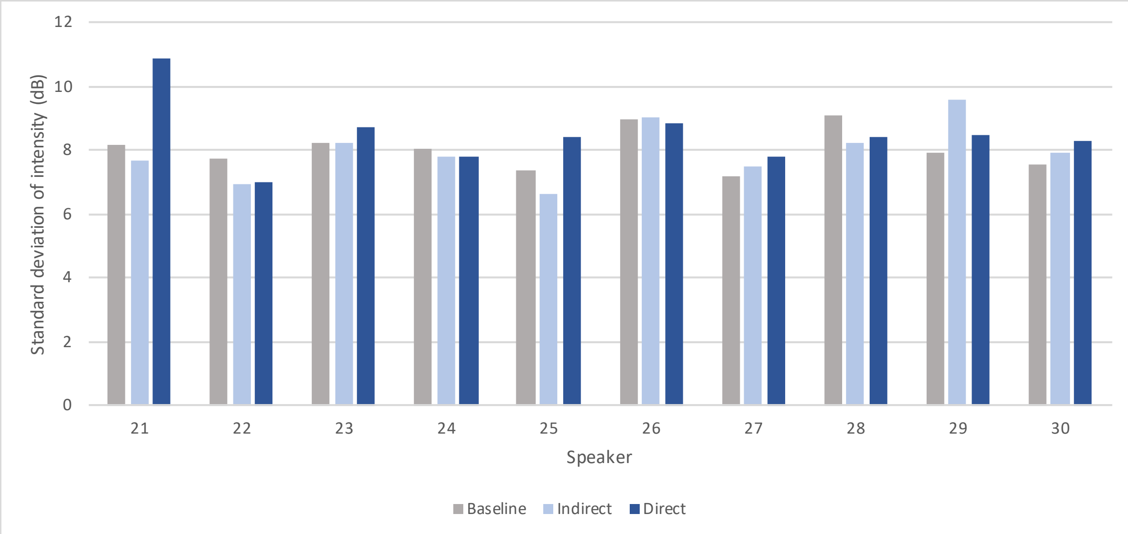

Summary of male speaker intensity analysis

The results of the intensity analysis for the male non-actors and actors can be summarised as follows:

- Overall, male non-actors produce their threats with an increase in intensity.

Female speakers

As with the male speakers, there is no clear prediction regarding the intensity of threatening, angry or aggressive speech. The figure below (41) illustrates the mean intensity for all of the female non-actors recorded for this study (Speakers 11-20). As before, percentage changes from the speaker's baseline speech are presented alongside each bar.

Figure 41: Change in mean intensity for indirect and direct threat conditions for female non-actors from their baseline speech

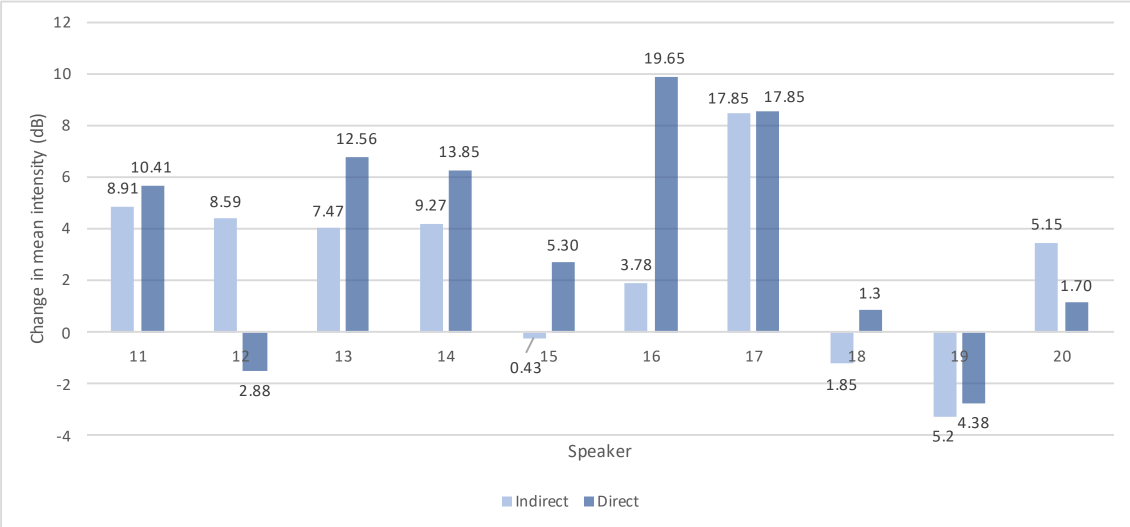

Overall, there appears to be a tendency for female speakers to produce threatening speech with an increase in intensity. This is the case for Speakers 11, 13, 14, 16, 17, and 20. For these speakers, there is variation regarding the extent of the increase between indirect and direct threats. For Speakers 11, 13, 14 and 16, direct threats were produced with the greatest increase in intensity. Speaker 17 produced both of her threats with highly similar increases in intensity, relative to her baseline speech. Speaker 20 produced indirect threats with a greater intensity increase than her corresponding direct threat.

The remaining female non-actors appeared to not follow this overall tendency to increase intensity during threat production. Speaker 12 produced her indirect threat with a notably higher intensity than her direct threat. In addition, her direct threat decreased in intensity relative to her baseline speech. Speaker 15 produced her indirect threat with a slight decrease in intensity, and her direct threat with an increased intensity, compared to her baseline speech. Finally, Speaker 18 produced her indirect threat with a decrease in intensity, and her direct threat with a slight increase in intensity in relation to her baseline speech.

As shown in Figure 42 below, the standard deviation of intensity appears to be similar overall across the female non-actors. Speaker 12 can be seen to markedly increase her intensity more than for her direct threat.

Figure 42: Standard deviation of intensity across all experimental conditions for female non-actors

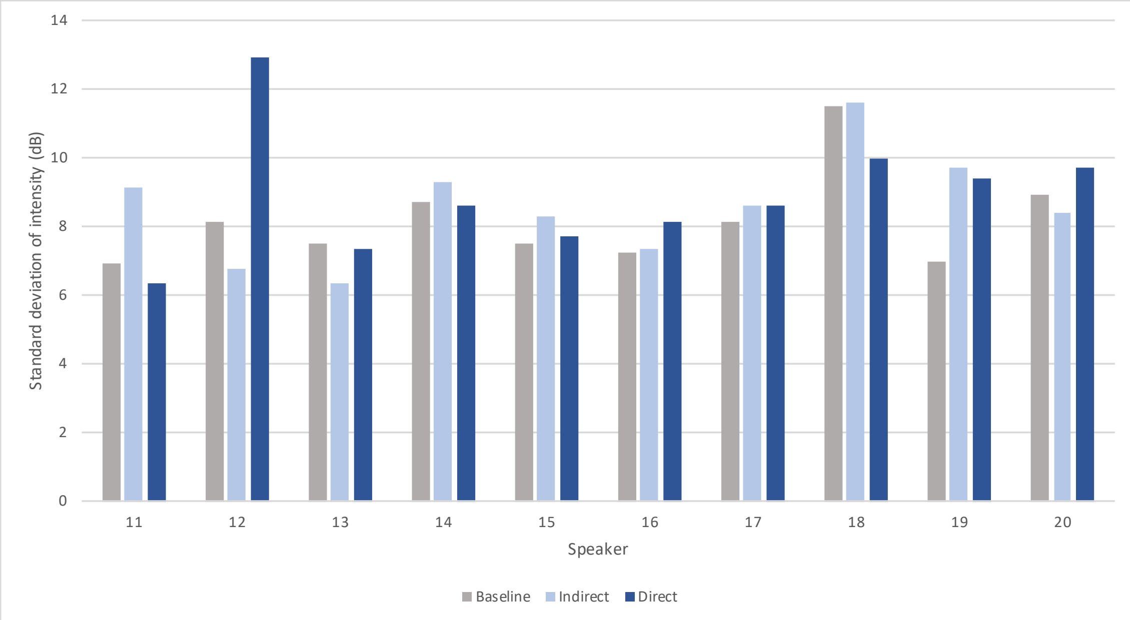

Figure 43 (below) shows the differences in mean intensity for the female actors (Speakers 31-41) recorded for this study.

Figure 43: Change in mean intensity for indirect and direct threat conditions for female actors from their baseline speech

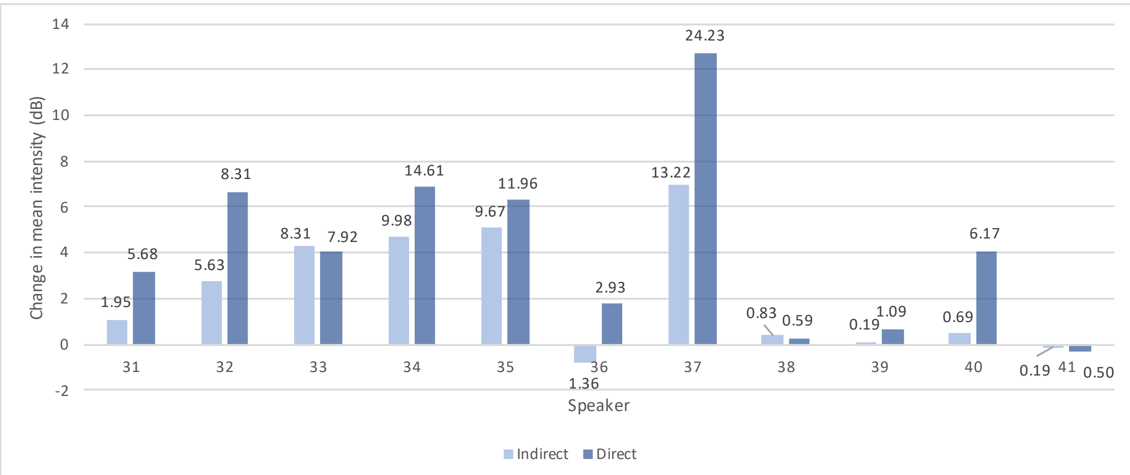

With the exception of Speakers 36 and 41, Figure 44 provides evidence of an apparent tendency for the female actors to produce their threats with an increase in intensity. However, there is variation concerning the extent of this increase. For the majority of these speakers (31, 32, 34, 35, 37, 39 and 40), direct threats were produced with greater intensity than their corresponding indirect threats. In particular, Speaker 37 produced her threatening speech with a markedly heightened intensity relative to her baseline speech. Speaker 36 produced her indirect threat with a decrease in intensity from her baseline speech, while her direct threat was produced with an increase in intensity. Speaker 40 produced both of her threats with a slight decrease in intensity relative to her baseline speech.

Figure 44: Standard deviation of intensity across all experimental conditions for female actors

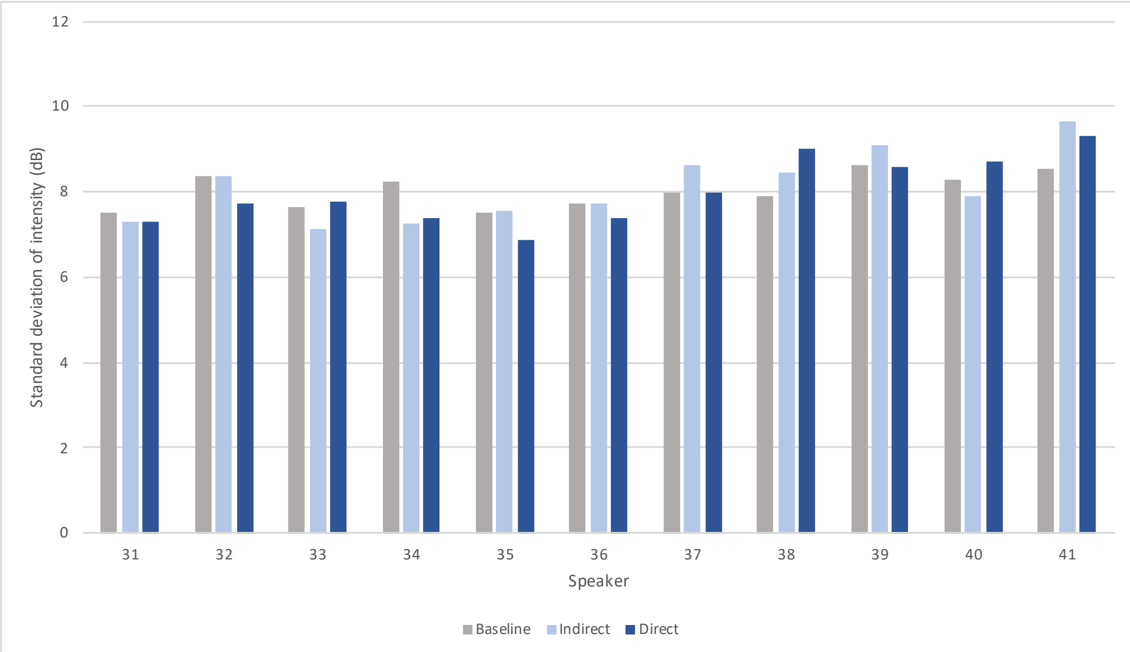

Summary of female speaker intensity analysis

The results of intensity analysis for the female speakers recorded for this study can be summarised as follows:

- Overall, both female non-actors and female actors produced their threatening speech with increased intensity.

5.2.4 Correlation between f0 and intensity

It is possible that there is a relationship between these speakers regarding changes in mean f0 and mean intensity. Previous research has shown that increasing f0 can correspond with increases in intensity (Lehiste and Lass, 1976). The table below (23) shows the output of a series of Pearson correlation coefficient tests between changes in mean f0 and changes in mean intensity between the speakers' baseline speech and indirect and direct threats.

Table 23: Pearson output - correlation between changes in mean f0 and mean intensity comparing baseline vs threat speech

Speaker group	Indirect threat	Direct threat
Male non-actors	0.20 (s.d. = 5.75)	0.51 (s.d. = 8.26)
Male actors	0.15 (s.d. = 11.45)	0.79 (s.d. = 14.43)
Female non-actors	-0.18 (s.d. = 7.36)	0.002 (s.d. = 7.07)
Female actors	0.70 (s.d. = 9.30)	0.81 (s.d. = 14.96)

The following images (45-48) show these correlations between mean f0 and intensity for each of the speaker groups recorded for this study.

Figure 45: Correlation between mean f0 and mean intensity across the male non-actor group

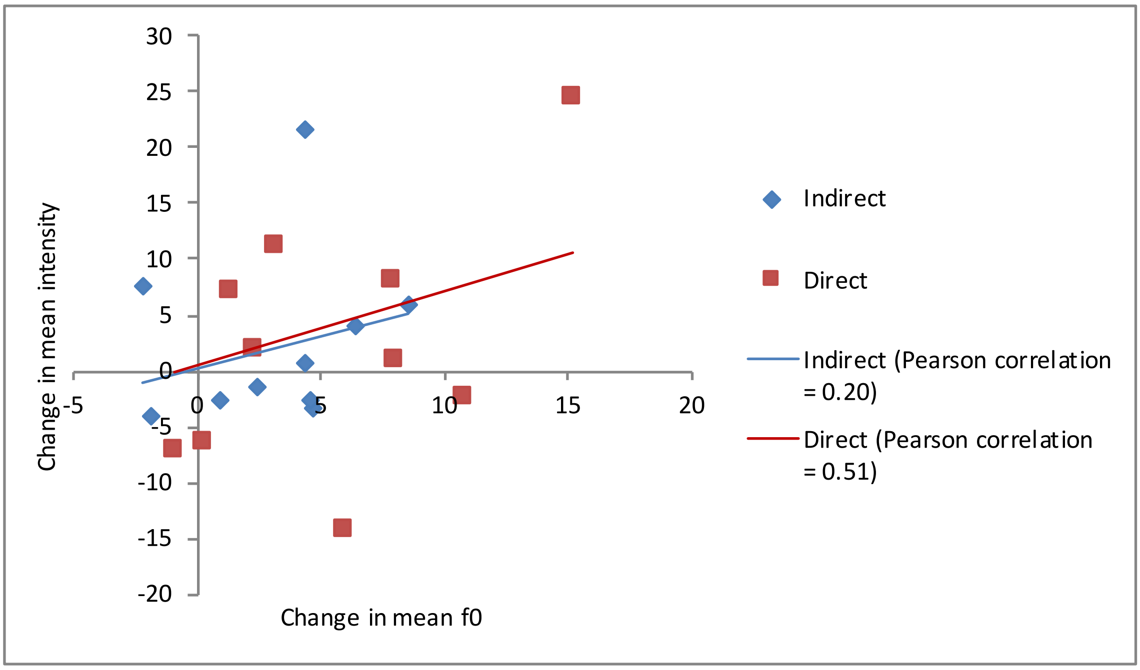

Figure 46: Correlation between mean f0 and mean intensity across the male actor group

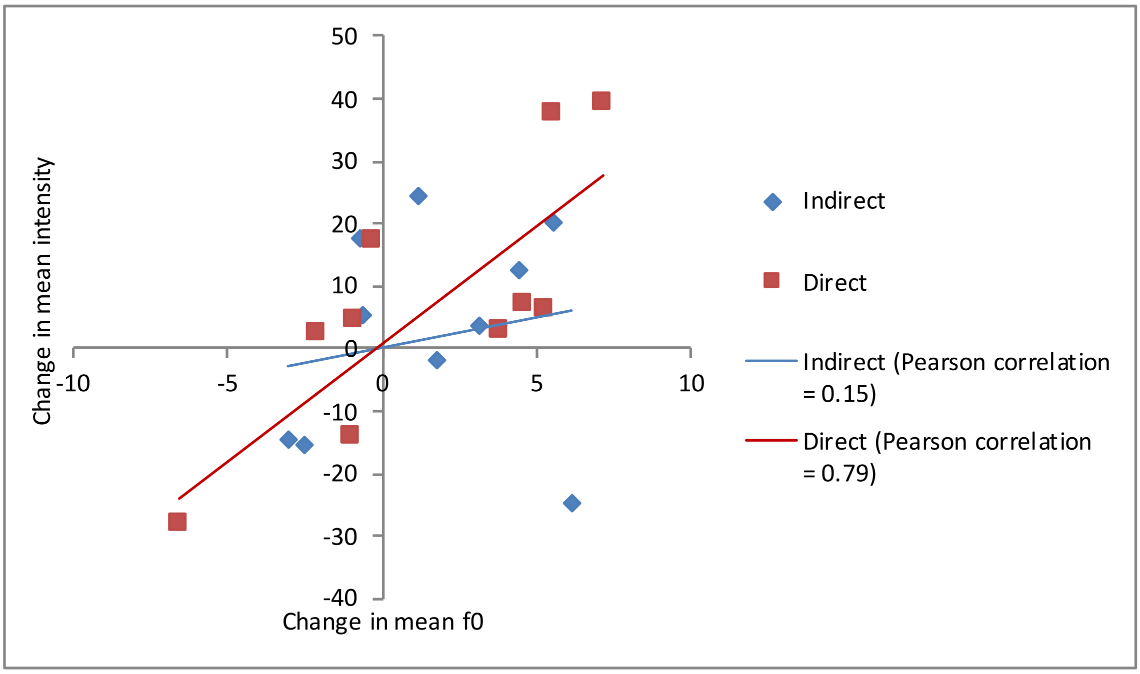

Figure 47: Correlation between mean f0 and mean intensity across the female non-actor group

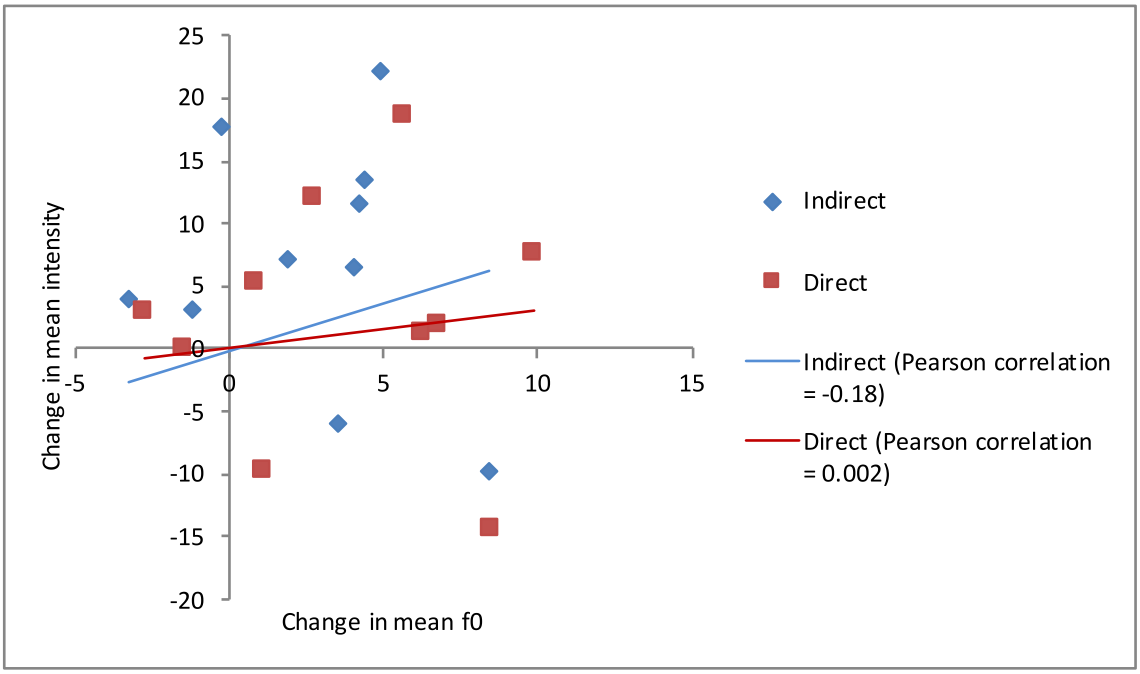

Figure 48: Correlation between mean f0 and mean intensity across the female actor group

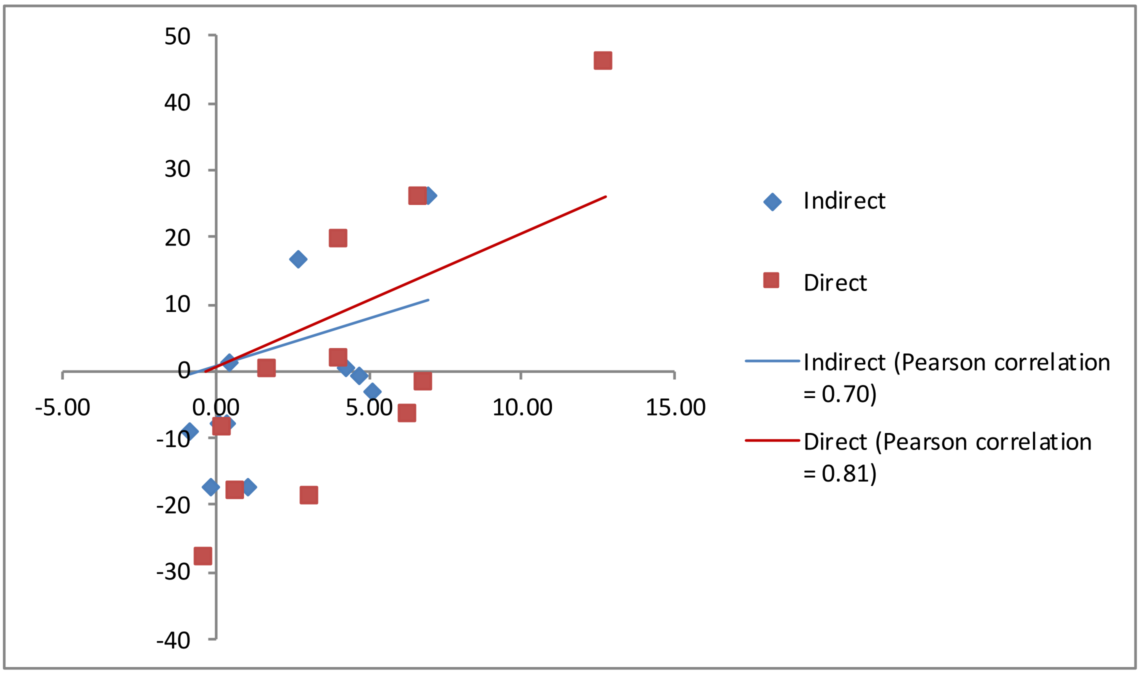

These results mostly show a positive linear relationship between changes in mean f0 and changes in mean intensity. In particular, there is a strong positive relationship between changes in mean f0 and changes in mean intensity for male and female actors producing direct threats. For female non-actors, there is a negative correlation between the same variables; for their direct threats, there is a near non-linear relationship. These results suggest that there was not necessarily a robust linear relationship between changes to mean f0 and changes to mean intensity in these data. In other words, speakers might produce speech with an increased f0, but without a marked increase in intensity, and vice versa.

5.2.5 Summary and discussion of intensity analysis

This subsection has presented the results of an analysis of the intensity of the speech of each speaker recorded for this study producing threatening and non-threatening speech. **Research question 1** addresses whether there are phonetic features which can be associated with threatening speech. These speakers (across the board) seem to produce threatening speech with a greater intensity relative to non-threatening, baseline speech. The extent of this increase varies across the relevant speakers. As such, this research provides some evidence that threatening speech can be characterised by an increase in intensity.

This apparent tendency to increase intensity appears to be consistent with previous literature about the phonetic features of angry speech. This relates to **research question 1.a** which addresses the possible link between the phonetic production of threatening speech and emotional speech. Chapter 2.0.2 further explores the possible overlap between the phonetic production of threatening speech and angry speech.

Research question 1.b considers whether indirect and direct threats differ phonetically. In these data, there appears to be a trend for direct threats to be produced with higher intensity than indirect threats. The reason for this trend is unclear. It is possible that speakers reading aloud texts involving themes of violence or aggression might increase their intensity in a way which is analogous to increasing intensity when conveying anger.

Research question 1.d considers the potential for differences in vocal production between participants who report acting experience and those who report no formal acting experience. Overall, male non-actors and male actors appeared to not produce notable differences between threatening and non-threatening speech. The female actor and non-actor groups, however, did seems to produce threatening speech with some differences in intensity. Therefore, regarding intensity, there is some evidence to indicate that actors produce threatening speech data which differs from non-actors. As discussed in the previous subsection, this has implications for the use of actor/non-actor speech data as stimuli in perception experiments. There is also the potential for research which examines the production of threats, to base its findings on the speech of only one of these groups. In other words, the findings of intensity (or f0) analysis of actors might not apply to real-life scenarios which presumably involve people without formal acting experience.

Finally, **research question 1.e** addresses the possibility that speakers of different biological sexes produce spoken threats differently. These data suggest that there seems to be little difference overall between male and female speakers in terms of intensity. As such, this research suggests that males and females do not appear to deliver threatening speech with notable differences in intensity relative to baseline, non-threatening speech.

The following subsection presents the results of articulation rate analysis performed on these data.

5.2.6 Articulation rate analysis

Previous literature does not provide a clear picture of the tempo or pace of threatening or angry speech. As described in §2.3.4, angry speech is generally described as being faster than emotionally-neutral speech. This subsection presents the results of articulation rate analyses for all of the speakers recorded for this research. This analysis was selected as it provides a quantitative measure of the speaker's tempo or pace during fluent speech. This subsection will discuss the articulation rates of the following experimental tasks:

- Task 1 (baseline read speech)
- Task 2 (indirect threat scripted by the speaker)
- Task 3 (direct threat scripted by the speaker)

As discussed earlier in Chapter 3.4.2, the output of an articulation rate analysis is a numerical figure which describes the number of syllables produced per second. An articulation rate output between 4-6 syllables per second is typical for English-language speakers (Takefuta, Jancosek & Brunt, 1971). Outputs which are lower than 4 would indicate slow-paced speech; outputs higher than 6 would suggest fast-paced speech.

A repeated measures ANOVA was performed with articulation rate (AR) values for speech recorded across the different experimental conditions (i.e. baseline, indirect and direct threat speech) as dependent variable, and speaker sex and level of acting experience as fixed factors, together with the threeway interaction between experimental condition, speaker sex and level of acting experience.

The ANOVA results showed no significant main effects but did find a significant interaction between experimental condition and level of acting experience ($F(1.80, 66.63) = 5.41$, $p = .008$)).

To interpret this result, the following plot was generated in SPSS software (49).

Figure 49: Mean articulation rates of baseline, indirect and direct threat speech

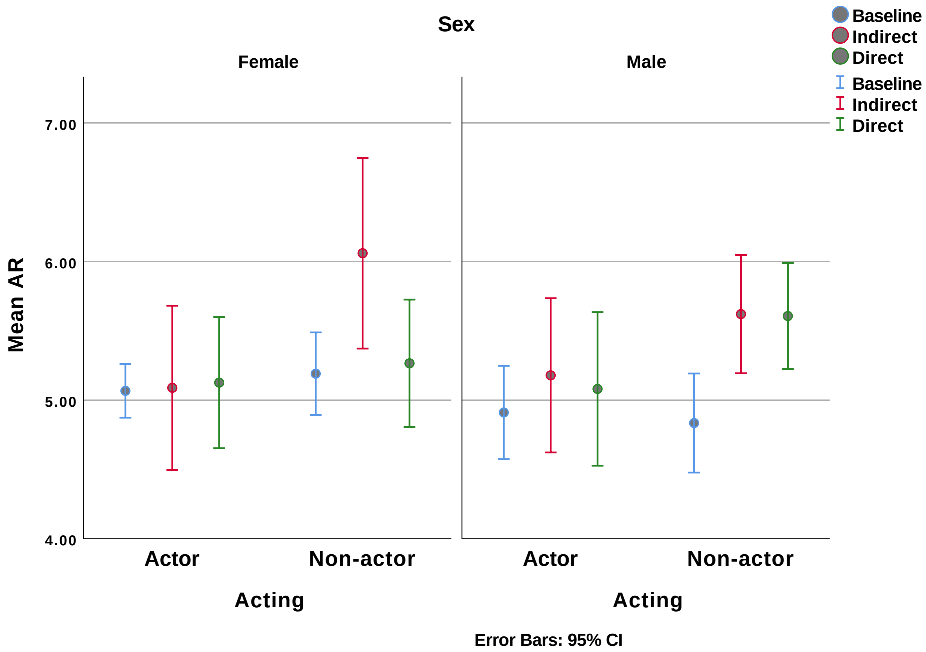

This plot suggests that the apparent variation in AR values was not solely due to experiment condition (that is, there is indeed no main effect of condition). Rather, that non-actors (particularly the female non-actors) generally seem to produce higher AR values than their actor counterparts.

The following section details a further analysis into each of these groups, beginning with the AR results pertaining to the male participants. This is to explore whether there are any within-group tendencies for the production of AR across different the experimental conditions. This subsection will present the results for each speaker group (male non-actors, male actors, female non-actors and female actors) in turn. Following the results of each speaker group, there will be a discussion of how these findings relate to previous relevant literature.

Male speakers

Figure 50 shows the articulation rates of each of the male non-actors (Speakers 1-10) producing baseline speech and indirect and direct threats.

Figure 50: Articulation rates for each condition produced by male non-actors

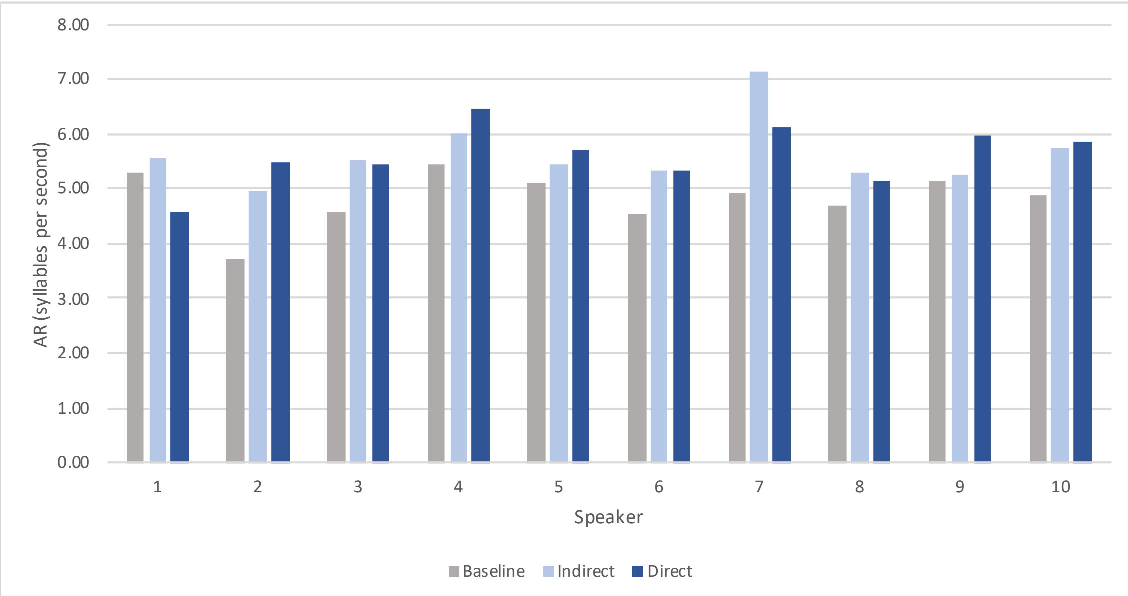

This figure shows that the majority of these speakers produced AR values which fell within the expected range for English-language speakers. Only Speaker 2 (in his baseline speech), Speaker 4 (in his direct threat) and Speaker 7 (in both of his threats) fall outside of this typical AR range. In general, the male non-actors produced threatening speech with raw ARs which were greater than their corresponding baseline speech. This was the case for all speakers except Speaker 1, who delivered his direct threat with a lower AR value than his baseline speech.

A repeated one-way ANOVA test showed a statistically significant result for the difference in AR values for the male non-actors across the three experimental conditions ($F = 11.105$, $p = 0.001$). The output of a Wilcoxon signed-rank test between the AR values of baseline speech and indirect threats was statistically significant (Z = -2.80, p = 0.005). There was also a statistically significant result between the baseline speech and the direct threats produced by these male non-actors (Z = -2.50, p = 0.012). As such, there is a tendency for these speakers to deliver threatening speech with a difference in AR in relation to their corresponding baseline speech. This tended to be an increase in AR during threatening speech relative to baseline speech.

The following figure (51) shows the articulation rate values for the male actors (Speakers 21-30) across the same three experimental conditions.

Figure 51: Articulation rates for each condition produced by male actors

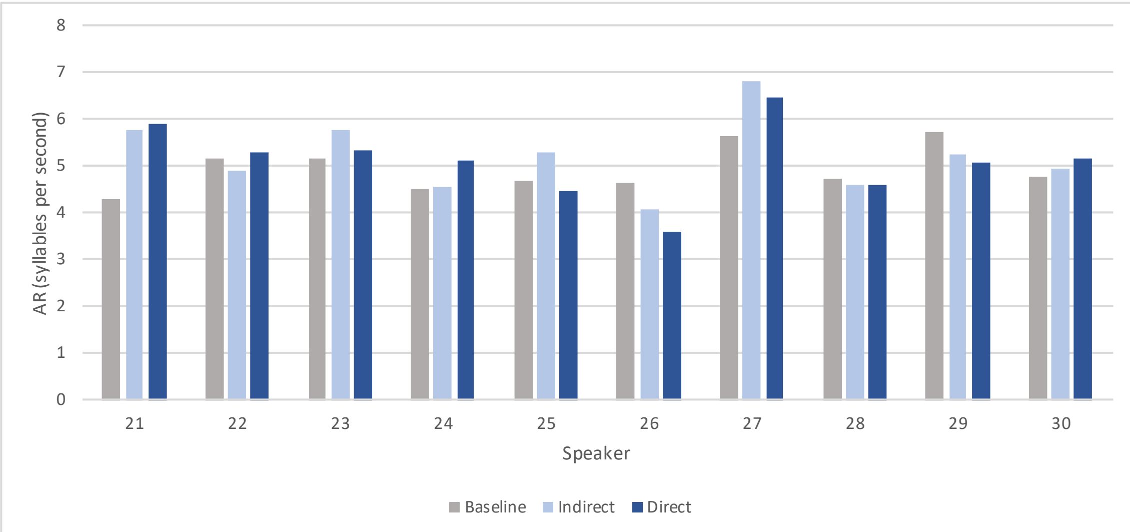

The figure shows as before, that the majority of the male actor recordings fell within the expected AR values for English-language speakers. Only Speaker 26's direct threat AR value was lower than this typical AR range. Also, only Speaker 27's threats were produced with an AR which was greater than this range.

For the male actors, there is a less clear picture concerning the articulation rate of threatening speech in relation to baseline speech. As with the male non-actors, the following male actors produced their threatening speech with a raw increase in tempo: Speakers 21, 23, 24, 27 and 30. The remaining male actors mostly produced threats with a raw decrease in tempo. Speaker 22 delivered his indirect threat with a slower tempo and his direct threat with a faster tempo relative to his baseline speech. In contrast, Speaker 25 delivered his indirect threat with a faster tempo and his direct threat with a slower tempo relative to his baseline speech.

A repeated one-way ANOVA test, however, showed a statistically non-significant result for the difference in AR values for the male actors across the three experimental conditions ($F = 0.90$, p 0.42). Further testing showed a statistically significant difference between the AR values for the male actor's baseline speech and their corresponding indirect threat ($Z = -2.47$, $p = 0.012$). There was no statistically significant difference between the AR values for the same speaker's baseline speech and their corresponding direct threat ($Z = -0.56$, $p = 0.58$). There was, however, a statistically significant result between the AR values of indirect and direct threats ($Z = -2.65$, $p = 0.008$).

These results suggest that the male actors tended to produce their indirect threats with a notable difference in tempo. However, the degree and direction of this change (i.e. an increase or decrease) do not seem to be consistent across these speakers. In addition, these results indicate a difference across these speakers in the tempo of indirect and direct threats. Speakers 21, 22, 24 and 30 produced their direct threats with a greater AR than their corresponding indirect threats. The remaining speakers in this group delivered indirect threats with a greater AR than their corresponding direct threats. As such, there appears to be a slight tendency for the male actors to produce their indirect threats with a faster tempo than direct threats.

Summary of male speaker articulation rate analysis

The results of AR analysis for the male speakers recorded for this research can be summarised as follows:

- Male non-actors tended to produce threatening speech with a faster tempo in relation to baseline, non-threatening speech.
- For male non-actors, both indirect and direct threats were produced with a statistically significant difference in AR from baseline speech. The AR values of these indirect and direct threats were not significantly different from one another.
- Male actors showed greater variation in terms of their AR production.
- For male actors, the AR values of indirect threats were statistically different from corresponding baseline speech. In addition, there was a statistical difference between the AR values of these speakers' indirect and direct threats.

The following section presents the results of articulation rate analysis relating to the female speakers recorded for this research.

Female speakers

Figure 52 shows the articulation rates of each of the female non-actors (Speakers 11-20) producing baseline speech and indirect and direct threats.

Figure 52: Articulation rates for each condition produced by female non-actors

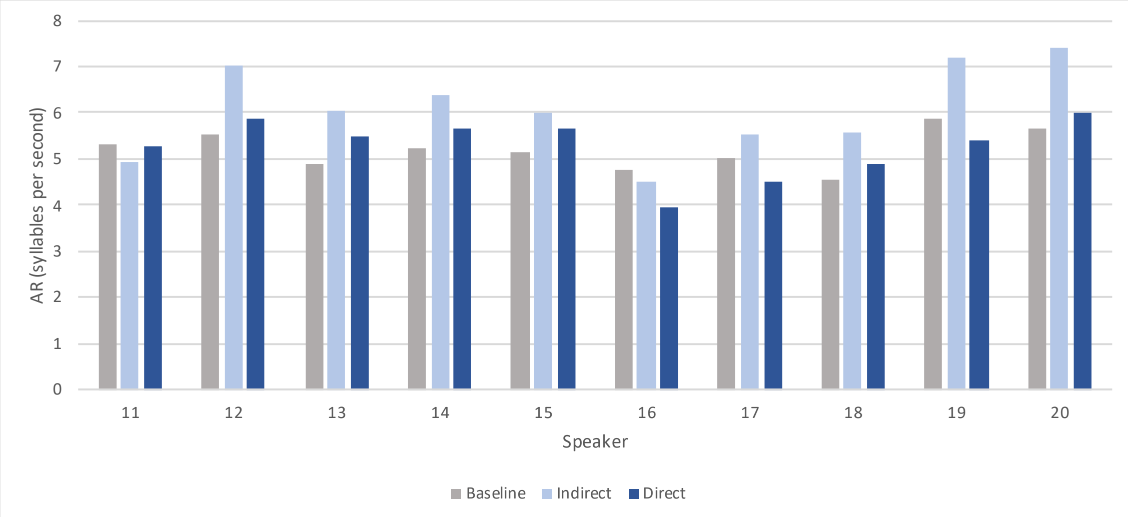

For the female non-actors, there are a number of speakers who produce speech which falls outside of the typical 4-6 syllables per second range. These speakers are Speakers 12, 14, 16, 19 and 20. Only Speaker 16, produces speech which is lower than this range during her direct threat. Speakers 12, 14, 19 and 20 produce ARs which are higher than this typical range. All of these instances occur during the production of indirect threats.

Figure 47 suggests that there may be a slight tendency for threatening speech to be produced with a faster tempo than baseline speech. This is evident for the speech delivered by Speakers 12, 13, 14, 15, 18 and 20. There are some exceptions to this pattern. Speaker 16 produces both of her threats with a decreased AR value. Speakers 11, 17 and 19 deliver indirect and direct threats with variations in the direction of AR relative to their corresponding baseline speech. For all of these female non-actors, indirect threats are produced with a higher raw AR value than direct threats.

A repeated one-way ANOVA test showed a statistically significant result for the difference in AR values for the female non-actors across the three experimental conditions ($F = 12.63$, $p = 0.0003$). A Wilcoxon signed-rank test did not show any statistically significant differences between the AR values of baseline speech and and direct threats ($Z = -0.36$, p = 0.72) as produced by female non-actors. However, there was a significant difference between the AR values of baseline speech and indirect threats for female non-actors (Z = -2.50, p = 0.01). This difference was typically manifested as an increase in AR when producing indirect threats (relative to baseline speech).

The following figure (53) shows the articulation rate of each of the female actors (Speaker 31-41) producing baseline speech and indirect and direct threats.

Figure 53: Articulation rates for each condition produced by female actors

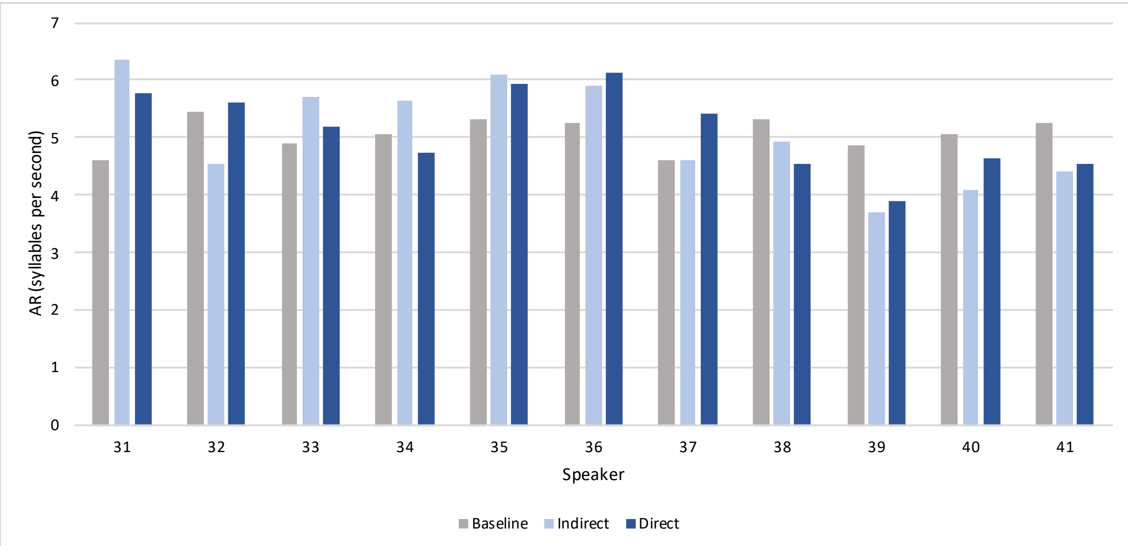

Figure 53 shows that, as with the other speakers in this study, the female actors also mostly produce AR values that are within the expected range for English speakers. Speakers 31, 35 and 36 each produce AR values that are slightly above 6 syllables per second. These higher AR values are recorded for these speakers producing threatening speech. Only Speaker 39 delivers an AR value which is lower than this typical AR range. These lower AR values are found in her threatening speech.

Overall, there seems to be a mixed picture for the tempo of threatening speech relative to non-threatening speech. Some of the speakers appear to increase their tempo from their baseline speech when producing threatening speech (indirect and direct). This includes Speakers 31, 33, 35 and 36. Other speakers in this group appear to lower their tempo relative to their baseline speech. This includes Speakers 38, 39, 40 and 41. The remaining speakers (Speakers 32, 34 and 37) vary in terms of whether the tempo of indirect or direct threats is higher or lower than their corresponding baseline speech.

These speakers also vary in terms of whether indirect threats or direct threats seem to be produced with a higher AR. The majority of the female actors (Speakers 32, 36, 37, 39, 40 and 41) produce direct threats with a higher AR than indirect threats.

A repeated one-way ANOVA test resulted in no statistically significant output for the difference in AR values for the female actors across the three experimental conditions ($F = 0.03$, $p = 0.97$). A series of Wilcoxon signed-rank tests also did not show any statistically significant differences between the AR values of baseline speech and indirect threats ($Z = -0.27$, $p = 0.79$), as well as baseline speech and direct threats (also $Z = -0.27$, $p = 0.79$) as produced by female actors. These non-significant outputs suggest that the female actors did not tend to consistently produce their threatening speech with a notable difference in tempo relative to their corresponding baseline speech. These results appear to be consistent with the female non-actor AR data previously presented.

Summary of female speaker articulation rate analysis

The results of AR analysis for the female speakers recorded for this research are summarised as follows:

- For female non-actors, indirect threats were overall produced with a statistically significant difference in AR relative to baseline speech. This difference was predominately an increase in AR during indirect threat production.
- For female actors, both indirect and direct threats were produced with no statistically significant difference in AR from baseline speech. The AR values of these indirect and direct threats were also not significantly different from one another.
- Female actors showed greater variation in terms of their AR production. Several speakers increased their AR when producing a threat. The same number of speakers decreased their AR when producing a threat.
- For female actors, there was no statistically significant difference between the AR values recorded for indirect and direct threats and baseline speech.

The following section presents the results of statistical comparisons made between all of the non-actors and actors (irrespective of their sex).

Comparison of non-actor and actor articulation rate analysis

The results of Mann-Whitney U tests showed no statistically significant difference between the AR values produced by non-actors and actors during baseline speech or direct threats ($U = 188.5$, $p = 0.43$). There was, however, a statistically significant result for the difference in AR production between non-actors and actors during indirect threats ($U = 118$, $p = 0.017$).

These results suggest that, overall, the non-actors and actors produced similar AR values during these experimental tasks. For indirect threats, there appears to be a difference in AR production between these acting groups. Based on Figures 50-53, there seems to be a tendency for the non-actors to produce indirect threats with an increased tempo compared to the actors.

The following section provides an overview of these articulation rate findings, and discusses their relevance to previous literature.

5.2.7 Summary and discussion of articulation rate analysis

This subsection has presented the results of a quantitative analysis into the tempo of threatening speech relative to baseline, non-threatening speech. Firstly, in response to research question 1, these data suggest that these speakers tend to increase their tempo when producing threatening speech. This was, however, by no means consistent across all of the speakers. There were a number of speakers who decreased their tempo when producing threatening speech. The degree of these increases or decreases also varied across these speakers. Crucially, these increases or decreases mostly did not fall outside of the typical or expected AR range for speakers of English. At present, there are no publicly-available predictions or research about the tempo of threatening speech. Therefore, this research suggests that speakers may opt for a slightly faster tempo during threatening speech.

Research question 1.a addresses whether the phonetic features of threatening speech resemble those observed for emotional states. The results of this AR analysis are also consistent with previous phonetic descriptions of anger. In particular, previous descriptions of 'hot' anger or rage (Burkhardt and Sendlmeier, 2000). The relationship between these speakers' production of threatening speech and angry speech will be presented in more detail in section 5.3 .

Research question 1.b considers whether indirect and direct threats are produced with phonetic differences to one another. In these data, there seems to be some evidence that speakers produce indirect threats with a higher AR than direct threats. This lends some support to the idea that indirect threats are phonetically distinct from more directly-phrased threats.

This said articulation rates provide a broad picture of the tempo of speech. It is possible that speakers use other temporal features which might distinguish threatening speech from non-threatening speech. For example, speakers might include more pauses or pauses that are longer in duration when making a threat. These features are not included in the calculation of articulation rates. As previously described in the discussion of f0 analysis, examining threatening speech using a finer-grained prosodic approach would provide additional insight into the tempo or pace of spoken threats.

Research question 1.d focuses on whether the phonetic productions of threats differs between non-actors and actors. In this study, non-actors and actors seemed to produce mostly similar ARs for baseline speech and direct threats. Indirect threats appeared to be produced more quickly by the non-actors relative to the actors. However, the actors tended to produce more varied ARs to one another when producing their threats.

These findings suggest that there are phonetic differences between these two groups. This could have potential implications for the use of non-actor or actor threatening speech data as stimuli in perception research. In particular, if listeners are found to assess threatening language on the basis of tempo. Currently, more research is needed to explore the link between the perception of tempo and the assessment of threats.

Finally, **research question 1.e** considers the possibility of male and female differences during spoken threat production. Overall, it appears that both male and female speakers seem to produce indirect threats with a higher AR than baseline speech or direct threats. There is little evidence to suggest that there were notable differences between the sexes in terms of the tempo produced during threatening speech. There is no publicly-available research (or hypotheses) about the potential differences between the tempo of male and female threatening speech. As such, this research suggests that there are, overall, no notable differences between male and female speakers for the tempo of threatening speech.

The next subsection shows the results of a series of vocal profile analyses performed on these data.

5.2.8 Vocal profile analysis

Previous phonetic literature has only briefly touched on the possible changes speaker make to their vocal tracts to sound threatening. This literature is primarily centred on the theories of sound symbolism and the 'frequency code' (Ohala, 1984). These theories attempt to explain why certain sounds (or types of sounds) are associated with seemingly abstract concepts. Of relevance to this research, threatening vocalisations have been linked with lower pitched, loud noises.

Related to these theories, there have also been attempts to account for the reasons for changes to the vocal anatomy of male speakers around puberty. As described earlier in this research, it has been proposed that bodily changes during puberty result in male vocal tracts being disproportionally large in relation to their bodies or to the vocal tracts of prepubescent children. This growth and elongation of the vocal tract are said to enable males to vocalise noises that are associated with sounding 'threatening'. As such, males are said to be able to more effectively communicate to others that they are a threat.

Based on these theories, threatening speech could be expected to be characterised by increases in intensity and decreases in fundamental frequency. As discussed in the previous results subsections, the speakers recorded for this study appeared not to produce threats which consistently met these expectations. It is still possible, however, that speakers modify their 'typical' or baseline vocal tract features, or use different phonation types, when producing threatening speech. In other words, these speakers might (subconsciously or otherwise) shape their vocal tracts in such a way that it is longer than it is during baseline speech. This could be achieved by, for example, rounding the lips while speaking or lowering the height of their larynx.

To assess these sorts of vocal tract changes a vocal profile analysis was performed on the following recordings for each speaker:

- Task 1 (baseline read speech)
- Task 2 (indirect threat scripted by the speaker)
- Task 3 (direct threat scripted by the speaker)

As described in Chapter 4 (methodology), this analysis is more subjective and impressionistic than the previous phonetic methods presented in this research so far. The analysis involved the author carefully listening to each recording, and cataloguing the presence of (and extent of) non-neutral vocal tract features for each recording using the modified VPA protocol.

For genuine forensic phonetic casework currently performed in the United Kingdom, however, this analysis would likely be performed by multiple people. This allows for multiple phoneticians to calibrate their own individual impressions to form a more definitive analysis. Owing to the quantity of data collected for this research, the vocal profile analysis of the current data has only been performed by the author. As such, this subsection serves to only catalogue the initial impressions of the voice quality and vocal tract features pertaining to threatening speech. Subsequent research in this area could then more closely scrutinise any pertinent vocal tract tract features using acoustic methods.

All of the vocal tract features and vocal settings outlined on the modified VPA protocol were listened for in this analysis. Doing so allows this research to identify features which might be of interest to future studies about threatening speech. There is also phonetic literature which has described the vocal tract features of angry speech. Data which addresses the possible similarities between the voice quality and vocal tract features of threatening speech and angry speech will be discussed in more detail in Chapter 5.3.

In the next section, overall impressions relating to these speakers vocal profiles will be presented and discussed.

5.2.9 Results of vocal profile analysis

Across the data collected for this research, there is little impressionistic evidence to support the consistent use of particular vocal features or voice qualities when delivering threatening speech. In particular, there are very few participants who seemed to lower their larynx during threatening speech production based on an impressionistic analysis. There are also no speakers who appear to protrude or round their lips during threatening speech.

What seemed to be more apparent for these speakers, is for changes to jaw movement during threatening speech (relative to their corresponding baseline speech). This usually consisted of minimising jaw movement (sometimes as though they were speaking through gritted teeth). This appears to be the case for Speakers 1, 5, 23, 27, 29 and 32. Other speakers seem to increase the extent of their jaw movement during threatening speech. This was noted for the threatening speech produced by Speakers 3, 4, 9 and 37. These jaw movement findings will be discussed later in this research in relation to angry speech.

There are also a few speakers who produce their threats with spread lips (as though smiling). This was noted for Speakers 22, 32 and 40 during their threatening speech. Speaker 32 also appeared to accompany her lip-spreading with some laughter during her threatening speech.

There are a number of speakers who appear to produce threats with very little change with respect to vocal features from their baseline speech. This appears to be the case across the non-actor and actor data. However, there are a notable number of the female non-actor speakers who - on an impressionistic level - appear to produce strikingly similar vocal tract features between their baseline and threatening speech. This appears to be the case for Speakers 11, 12, 14, 15, 16, 19 and 20.

The results of vocal profile analysis performed on all of the baseline and threatening speech collected for this study are summarised as follows:

- There appeared to be no clear tendency for threatening speech (of either type) to be associated with particular vocal tract features or vocal settings.

5.2.10 Summary of vocal profile analysis

This subsection has presented the findings of a series of vocal profile analyses performed by the author on the baseline and threatening speech collected for this research. The findings of this analysis intended to offer a first impression of the vocal features of threatening speech. These findings can be explored more thoroughly in future research. As discussed earlier in this section, analysis which is as subjective as vocal profile analysis should ideally be subjected to calibration by multiple investigators.

In response to **research question 1**, there appears to be no evidence of particular vocal tract features or vocal settings which appear to be associated with threatening speech. This finding differs from the predictions put forward in research which is based on theories stemming from biological or evolutionary perspectives (Ohala, 1984; Xu, Kelly and Smillie, 2013). Namely, this research goes against the expectation that speakers would deliver threats using vocal tract features which elongate the vocal tract.

In addition, a substantial number of these speakers appeared to not make notable changes to vocal tract features or vocal setting from baseline speech to threatening speech. Some speakers, however, produced threatening speech with changes to jaw position. This finding will be discussed later in this study in relation to angry speech.

Research question 1.b centred on the possible phonetic differences between indirect and direct threats. In terms of vocal tract features or vocal setting, there appeared to be little evidence of notable changes between indirect and direct threats. This finding indicates that speakers to not distinguish these types of threats thought changes to vocal tract features and vocal settings.

Research question 1.d considered whether there is evidence to suggest that non-actors and actors produced threatening speech differently to one another. With respect to vocal tract features, this appears to be unclear. There was a considerable amount of variation across these speaker groups in terms of the vocal tract features used when delivering threatening speech. There seemed to be a more notable difference between the vocal tract features of the male and female participants. This relates to **research question 1.e** which focuses on whether there are differences in threat production based on speaker sex. In terms of vocal tract features, this appeared to be the case. This was largely driven by the impression that female speakers consistently did not make notable changes to their vocal tract features between their baseline and threatening speech.

5.2.11 Summary of phonetic comparison of baseline speech and threatening speech

Overall, there was no consistent pattern of phonetic features associated with a threatening vocal production across the speakers recorded. This finding might suggest that, rather than a single strategy for signalling a threatening intent, there may be instead multiple methods employed by speakers. For example, Speakers 9, 15, and 27 all increased both mean fundamental frequency and articulation rate when making threats of either type. On the other hand, Speakers 26, 38, 39 and 41 all appear to decrease their mean f0 and AR when producing a either an indirect or direct threat. The remaining speakers appear to produce threats with variable increases and decreases to mean fundamental frequency, intensity and AR.

For the former group of speakers (those who increased f0, AR and intensity), their productions appear to more closely resemble the previous literature of anger, particularly so-called 'hot anger' or 'rage' (Burkhardt and Sendlmeier, 2000). The next section provides a more detailed comparison between the threatening speech and angry speech produced by these speakers.

5.3 Comparison of the phonetic features of baseline speech, threatening speech and 'anger' speech

This subsection aims to shed light on the possible phonetic similarities between threatening speech and angry speech. This intends to address **research question 1.a**, which focuses on whether the phonetic features of threatening speech also resemble other emotional speech. As seen in Chapter 5 (Results), there were a number of phonetic features which were common across the threatening speech data which have also been linked with productions of anger in previous research. These included increases in f0, intensity and tempo.

The data in this subsection presents the results of the analysis of f0, intensity, articulation rate and vocal profile. This analysis was performed on the following data collected for this research:

- Task 1 (baseline read speech)
- Task 2 (indirect threat scripted by the speaker)

- Task 3 (direct threat scripted by the speaker)
- Task 4 (emotionally-neutral sentence read aloud by the speaker in an 'angry' way)

This subsection will present the results for each speaker group (male non-actors, male actors, female non-actors and female actors) in turn. The methods used to calculate these data are the same as described for the simulated data in Chapter 4 (Methodology). Following these results, there will be a discussion of how these findings relate to previous literature relating to anger speech.

5.3.1 Fundamental frequency

Male speakers

Figure 54 shows the change in mean f0 for the male non-actors (Speakers 1-10) producing indirect and direct threats, and angry speech relative to their baseline speech.

Figure 54: Change in mean f0 for threatening speech and angry speech produced by male non-actors from baseline speech

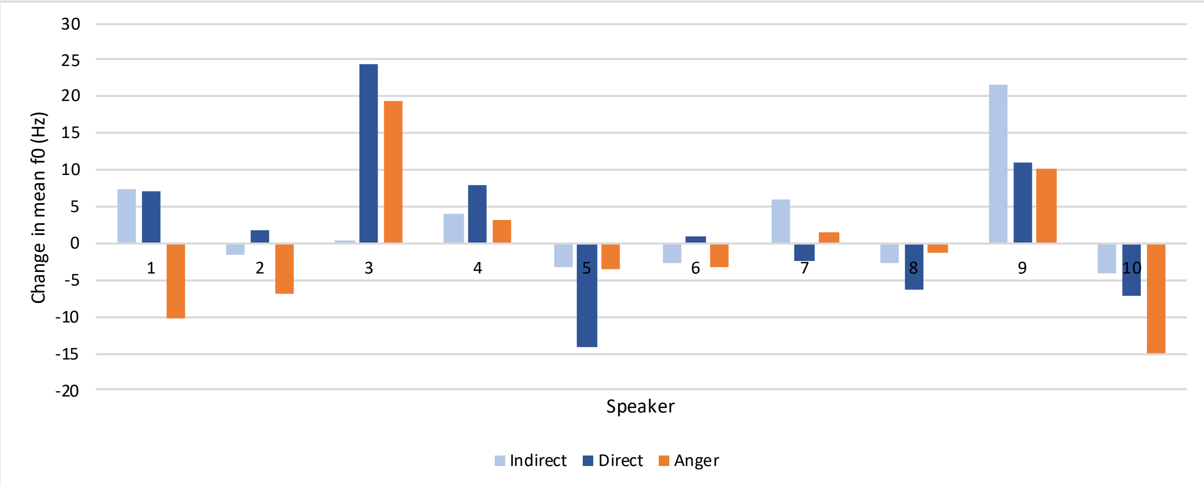

From this figure, there does not seem to be a clear relationship between the mean f0 of threatening speech and angry speech. This interpretation is supported by a statistically non-significant result of a one-way ANOVA test between these data ($F = 0.56$, $p = 0.64$). Figure 49 suggests some speakers deliver angry speech with a decrease in raw mean f0 relative to their baseline speech. This can be seen in the data for Speakers 1, 2, 5, 6, 8 and 10. The remaining male non-actors increase their mean f0 when producing angry speech relative to their baseline speech. These findings indicate that there is no tendency for the male non-actors to produce angry speech with a similar mean f0 to one another.

The following figure (55) shows the mean f0 for the male actors (Speakers 21-30) producing baseline speech, indirect and direct threats, and angry speech.

Figure 55: Change in mean f0 for threatening speech and angry speech produced by male actors from baseline speech

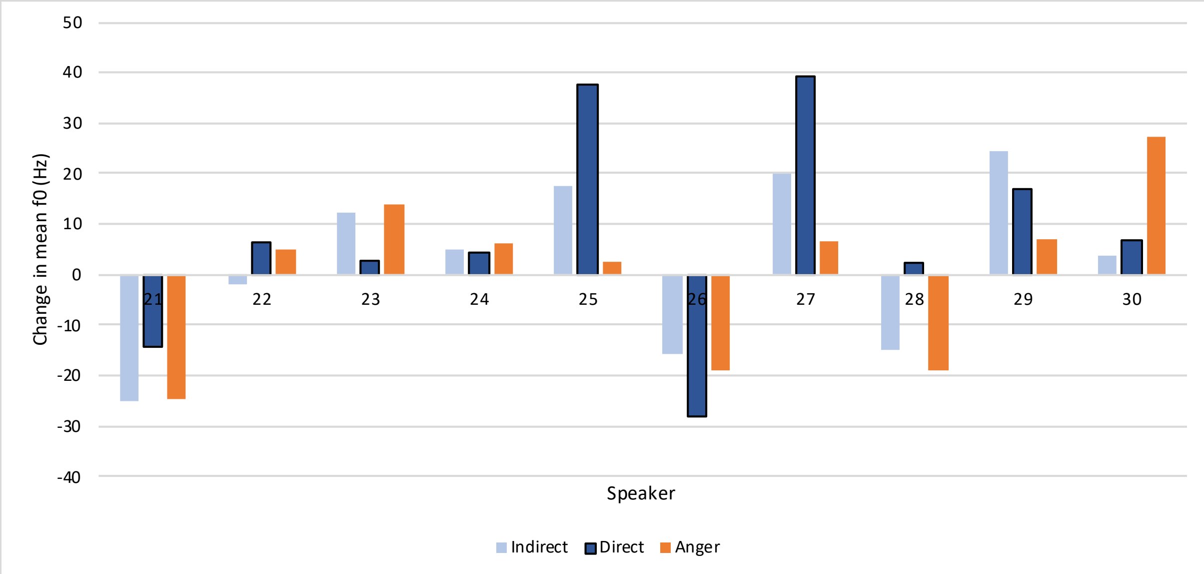

Figure 55 suggests that there is variation in the mean f0 of angry speech. Some the male actors appear to produce angry speech with a raw increase in mean f0 relative to their baseline speech. This can be seen in the data for Speakers 22, 23, 24, 25, 27, 29 and 30. The degree of this increase varies slightly between these speakers. The remaining male actors appear to produce angry speech with a decrease in mean f0 relative to their baseline speech. A repeated measures ANOVA test did not produce a statistically significant result for the mean f0 of the male actors producing speech across these experimental conditions ($F = 0.83$, $p = 0.49$).

These results suggest that male actors seem to produce comparable mean f0s across these four experimental conditions. As such, while threatening speech and angry speech are alike in terms of f0, angry speech also has similar mean f0s to baseline, emotionally-neutral speech.

Summary of male speaker fundamental frequency analysis: angry speech

The results of the fundamental frequency analysis of angry speech produced by the male speakers can be summarised as follows:

- There is no statistical difference between the mean f0 of angry speech (as well as the other experimental conditions) produced by the male non-actors or male actors.

The following section presents the results of fundamental frequency analysis relating to the angry speech delivered by the female speakers recorded for this research.

Female speakers

Figure 56 shows the change in mean f0 for the female non-actors (Speaker 11-20) delivering indirect and direct threats, and angry speech relative to their baseline speech.

Figure 56: Change in mean f0 for threatening speech and angry speech produced by female non-actors from baseline speech

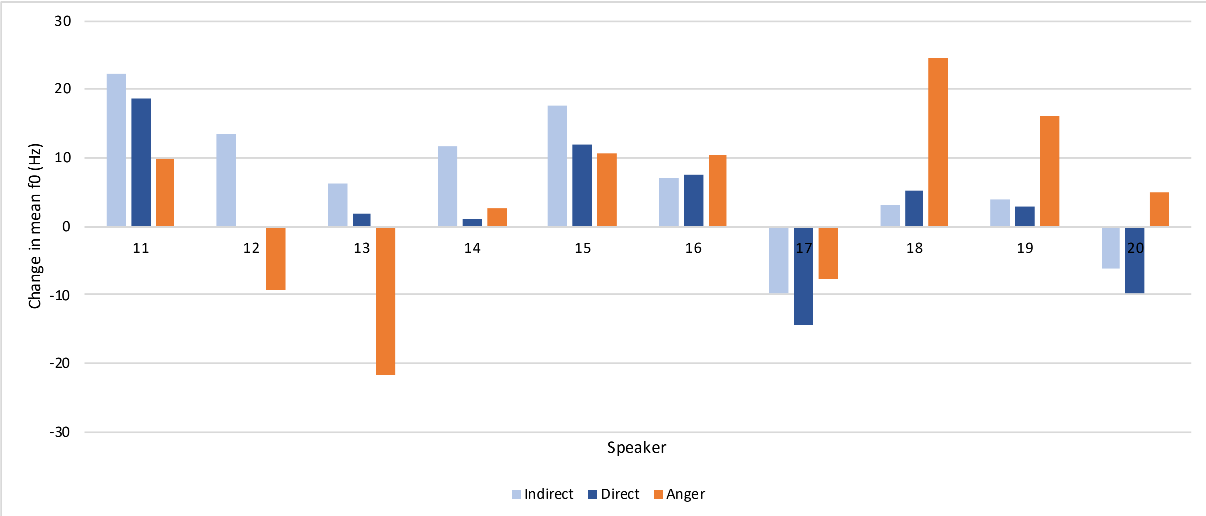

Figure 56 suggests that some of the female non-actors produce angry speech with a raw increase in mean f0 relative to baseline speech. This is the case for Speakers 11, 14, 15, 16, 18, 19 and 20. The degree of this increase is variable between these speakers. The remaining speakers (Speakers 12, 13 and 17) deliver angry speech with a lower mean f0 than their baseline speech.

There appears to be variation as well, in terms of the mean f0 of angry speech relates to threatening speech. In some cases, angry speech is produced with a higher mean f0 than threatening speech. This increase is particularly pronounced for Speakers 18 and 19. Speakers 16, 17 and 20 also increase their raw mean f0 when delivering angry speech relative to threatening speech.

For Speakers 11, 12, 13, and 15, angry speech has a lower raw mean f0 value than threatening speech (of either type). Finally, Speaker 14 produces angry speech with a lower mean f0 than her indirect threat only.

A subsequent one-way ANOVA test did not produce a statistically significant result between the mean f0 values for the female non-actors across these experimental conditions (F = 1.24, p = 0.31).

The following figure (57) shows the change in mean f0 for the female actors (Speaker 31-41) delivering indirect and direct threats, and angry speech relative to their baseline speech.

Figure 57: Change in mean f0 for threatening speech and angry speech produced by female actors from baseline speech

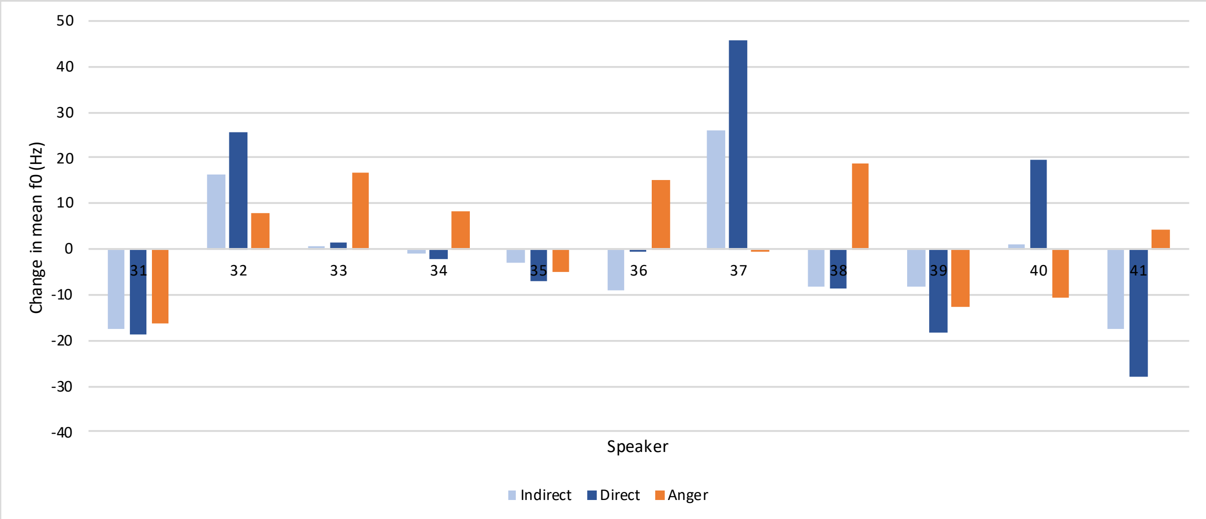

As before, some of the female actors produced their angry speech with an increase in mean f0 relative to their baseline speech. This can be seen in the data for Speakers 32, 33, 34, 36, 38 and 40. The remaining speakers decrease their mean f0 when delivering angry speech compared to baseline speech.

There is also a mixed picture when it comes to the production of angry speech relative to threatening speech. Speakers 31, 33, 34, 36, 38 and 40 all increase their f0 when delivering angry speech in relation to threatening speech. Similarly to the female non-actor data, a one-way ANOVA test also did not produce a statistically significant result between the mean f0 values for the female actors across these experimental conditions (F = 0.23, p = 0.88).

Summary of female speaker fundamental frequency analysis: angry speech

The results of the fundamental frequency analysis of angry speech produced by the female speakers can be summarised as follows:

- There is no statistical difference between the mean f0 of angry speech (as well as the other experimental conditions) produced by the female non-actors or male actors.

Summary of fundamental frequency analysis: angry speech

Overall, there was little evidence in these data to suggest that these speakers consistently produced angry speech with a notable difference in mean fundamental frequency relative to baseline and threatening speech. This finding was irrespective of both speaker sex or the level of acting experience reported.

The next section presents the results of intensity analysis relating to the angry speech delivered by these same speakers. The methods used to measure the intensity of the angry speech recorded for Task 4 are the same as those described in Chapter 4 (Methodology). As before, no statistical comparisons between made across these data. This is because intensity measurements are strongly affected by the speaker's proximity to the recording source.

In the next section, the data for the male participants will be presented first. This will be followed by the intensity data provided by female participants.

5.3.2 Intensity

Male speakers

Figure 58 shows the change in mean intensity for threatening speech and angry speech relative to baseline speech for the male non-actors (Speakers 1-10).

Figure 58: Change in mean intensity for threatening speech and angry speech produced by male non-actors from baseline speech

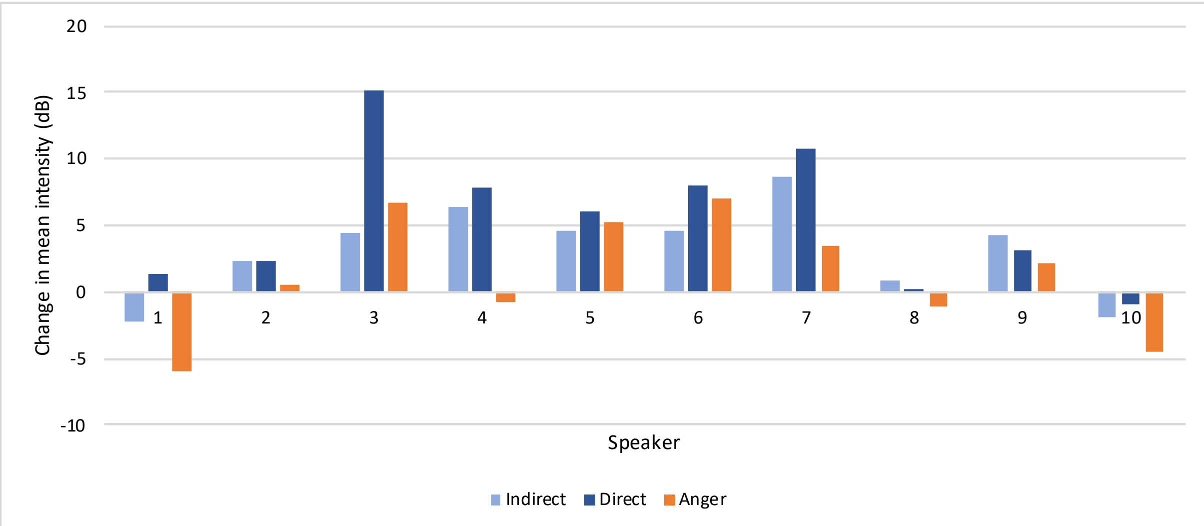

For many of these speakers, the mean intensity of angry speech appears to be increased relative to their corresponding baseline speech. The exceptions to this pattern are Speakers 1 and 10 who decrease their mean f0 during angry speech. A number of the male non-actors produce angry speech with a lower mean intensity than threatening speech. This can be seen in the data for all of the participants in this group except for Speaker 3, whose angry speech is delivered with a higher mean intensity than his indirect threat.

Figure 59 shows the change in mean intensity of angry and threatening speech relative to baseline speech for the male actors (Speakers 21-30).

Figure 59: Change in mean intensity for threatening speech and angry speech produced by male actors from baseline speech

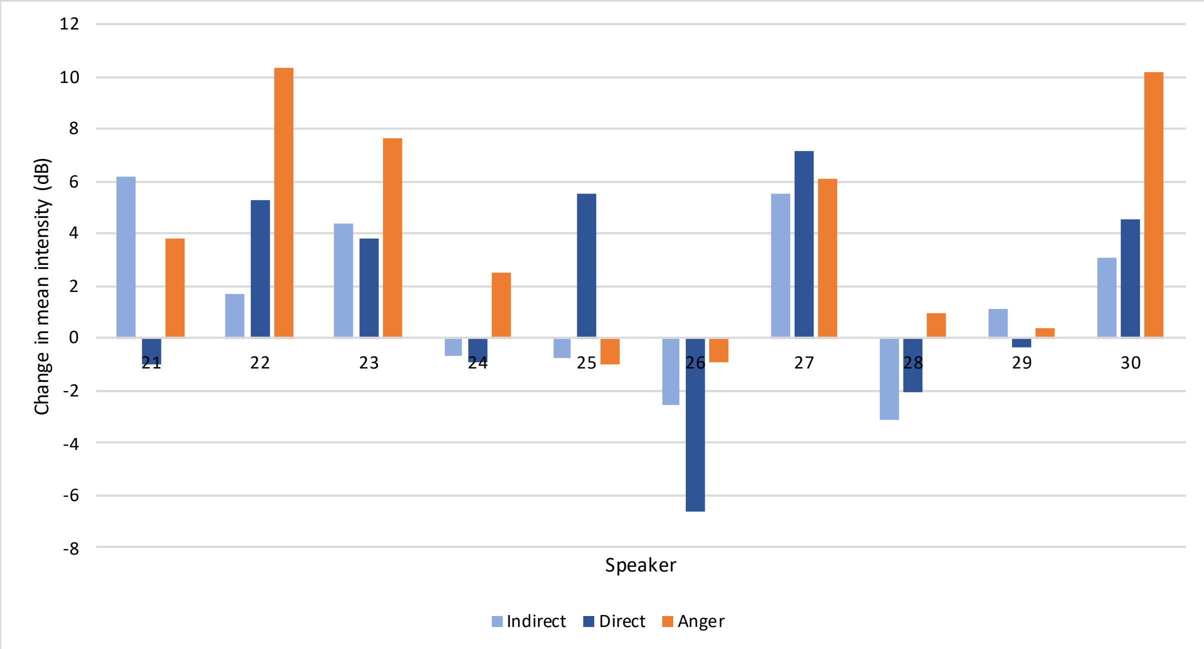

For the male actors, there is a slight trend for angry speech to be delivered with an increase in mean intensity relative to baseline speech. This can be seen in the data for Speakers 21, 22, 23, 24, 27, 28, 29 and 30. Speakers 25 and 26 show a decrease in the mean intensity produced during baseline speech when delivering angry speech.

In addition, there are a number of speakers who produce angry speech when an increase in mean intensity relative to threatening speech (both indirect and direct). This applies to Speakers 22, 23, 24, 27, 28 and 30. Speakers 21 and 29 produces angry speech with an increase relative to his direct threat and a decrease compared to his indirect threat. Finally, Speaker 25 produces angry speech with a decreased mean intensity relative to both his indirect and direct threats.

Summary of male intensity analysis: angry speech

The results of the intensity analysis of angry speech produced by the male speakers can be summarised as follows:

- Male non-actors and male-actors tended to increase their mean intensity when producing angry speech relative to baseline speech.
- For the male non-actors, angry speech tended to be produced with a lower mean

intensity than threatening speech.

- For the male actors, angry speech tended to be produced with a greater mean intensity than threatening speech.

The following section presents the findings of intensity analysis performed on the angry speech of the female speakers recorded for this study.

Female speakers

Figure 60 shows the change in mean intensity of angry and threatening speech relative to baseline speech for the female non-actors (Speakers 11-20).

Figure 60: Change in mean intensity for threatening speech and angry speech produced by female non-actors from baseline speech

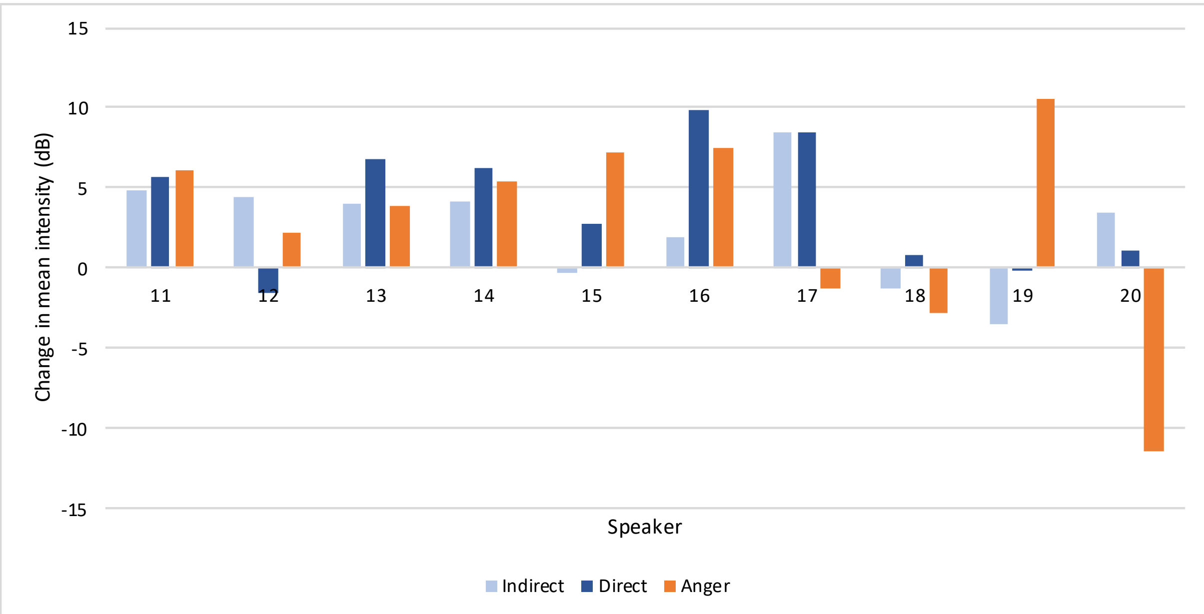

This figure suggests a tendency for the female non-actors to produce angry speech with a greater mean f0 relative to their baseline speech. Speakers 17, 18 and, in particular, Speaker 20 are exceptions to this pattern.

There does not seem to be as clear of a trend for the mean intensity of angry speech relative to threatening speech. Speakers 11, 15 and 19 produce angry speech with an increase in mean intensity relative to threatening speech. Speakers 13, 17, 18 and 20 produce angry speech with a decrease in mean intensity relative to threatening speech. Speakers 12 and 14 produce angry speech with a mean intensity that is higher than their direct threats but lower than their indirect threats. In contrast, Speaker 16 produces her angry speech with a lower mean intensity than her direct threat, and a greater mean intensity than her indirect threat.

Figure 61 shows the change in mean intensity of angry and threatening speech relative to baseline speech for the female actors (Speakers 31-41).

Figure 61: Change in mean intensity for threatening speech and angry speech produced by female actors from baseline speech

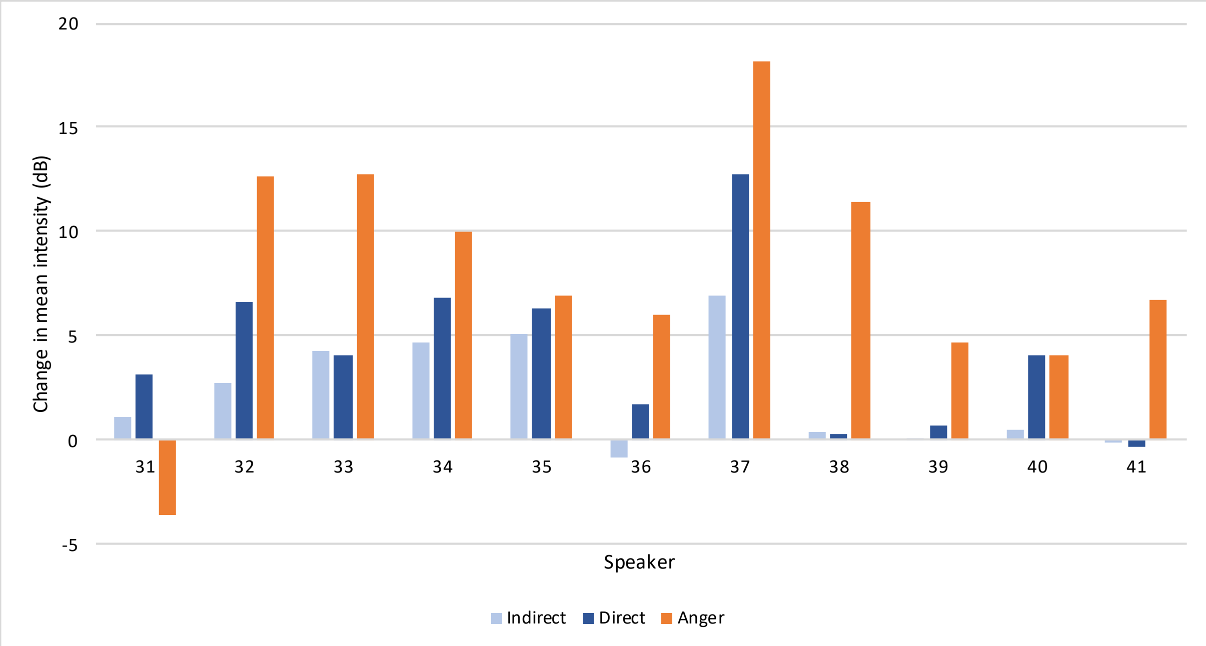

The female actors show a clear tendency to increase their mean intensity relative to their baseline speech. The exception to this pattern is Speaker 31. In addition, these speakers also have a tendency to increase their mean intensity relative to their threatening speech. Speaker 40 produces her direct threat with an ever so slight increase in mean intensity than her angry speech.

Summary of female intensity analysis: angry speech

The results of the intensity analysis of angry speech produced by the female speakers can be summarised as follows:

- Female non-actors tended to increase their mean intensity when producing angry speech relative to baseline speech. There was also a slight tendency for these speakers to decrease the mean intensity of angry speech relative to threatening speech. However, this pattern was by no means consistent across the female non-actors.

- Female-actors mostly increased their mean intensity when producing angry speech relative to baseline speech and threatening speech.

Summary of intensity analysis: angry speech

There was an overall tendency among these speakers to produce angry speech with an increased intensity relative to their corresponding baseline speech. A more complicated relationship arose between the mean intensity of angry speech compared to threatening speech. For male and female non-actors, angry speech was produced with a decreased intensity relative to threatening speech. For male and female actors, there was instead a tendency to deliver angry speech with an increased intensity relative to their threatening speech.

The next section presents the results of an articulation rate analysis relating to the angry speech delivered by these same speakers. The methods used to measure the ARs of the angry speech collected in Task 4 are the same as those described in Chapter 3. It should be emphasised that the duration of these angry speech recordings are short (between 1-3 seconds). As explained earlier in this study, AR measurements are more reliable when taken from longer samples of fluent speech. Therefore, the AR measurements calculated from these angry speech recordings should be treated with caution.

5.3.3 Articulation Rate

Male speakers

Figure 62 shows the articulation rates for baseline, threatening and angry speech produced by the male non-actors (Speakers 1-10).

Figure 62: Articulation rates for baseline, threatening and angry speech produced by male non-actors

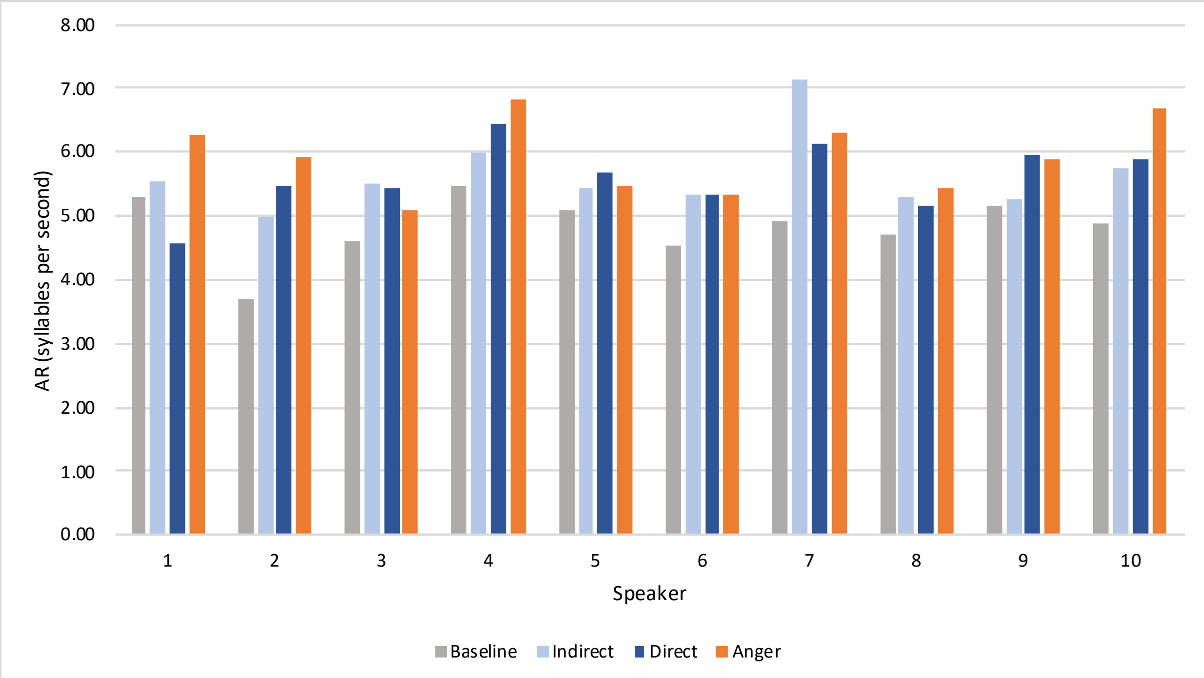

All of the male non-actors produce angry speech with a higher raw AR than their corresponding baseline speech. For half of these speakers, angry speech was produced with a greater AR than threatening speech. This can be seen in the data for Speakers 1, 2, 4, 8 and 10. The degree of these tempo increases varies between these speakers. For Speakers 1, 4, 8 and 10, the articulation rate exceeds the expected 4-6 syllable per second range for English speakers. None of the male non-actors fell below this AR range when delivering angry speech. The results of a one-way ANOVA test show a statistically significant difference between the ARs produced by the male non-actors across these four experimental conditions ($F = 11.83$, $p = 0.00004$).

Figure 63 shows the articulation rates for baseline, threatening and angry speech produced by the male actors (Speakers 21-30).

Figure 63: Articulation rates for baseline, threatening and angry speech produced by male actors

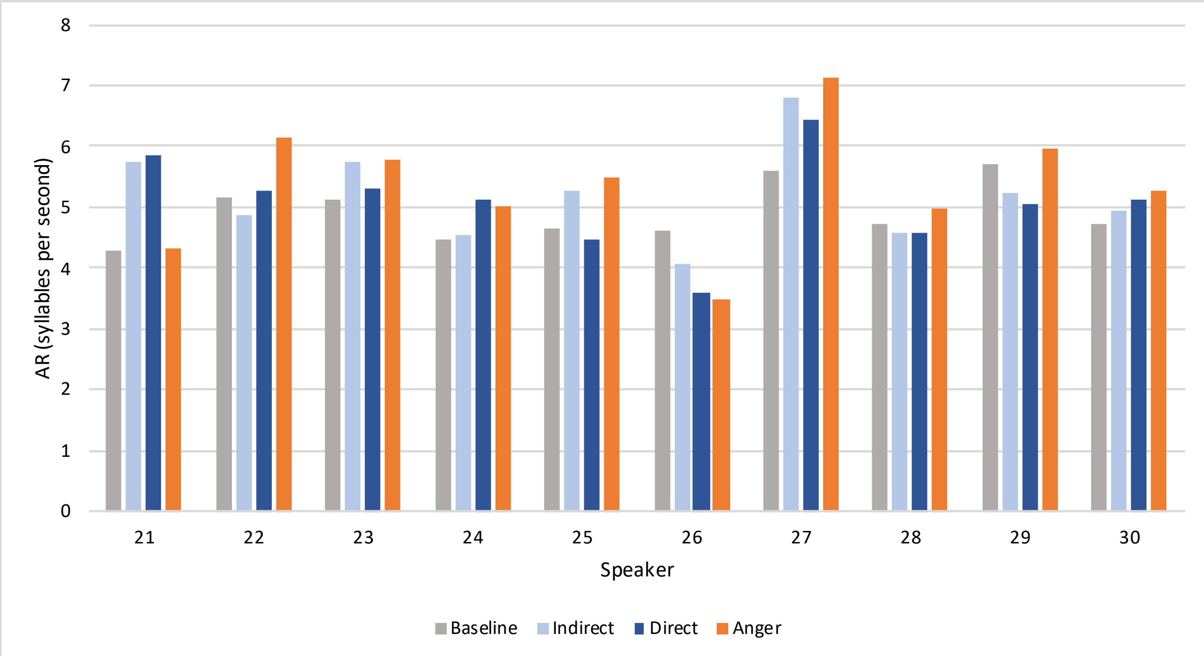

This figure suggests some of the male actors deliver angry speech with an increased AR relative to their baseline speech. Speakers 22 and 27 deliver their angry speech with an AR that exceeds the expected 4-6 syllable per second range for English speakers. Speaker 26 is the only speaker in this group to produce angry speech with a slower tempo than baseline speech. Speaker 26's AR value is also lower than the expected 4-6 syllables per second range. However, the results of a one-way ANOVA test show no statistical significance between the ARs produced by the male actors across these four experimental conditions ($F = 1.41$, $p = 0.26$).

There is also a tendency for the male actors to produce angry speech with greater AR values than threatening speech. This can be seen in the data for Speakers 22, 23, 25, 27, 28, 29 and 30. Speaker 24 delivers his direct threat with an ever so slight increase in AR relative to his angry speech.

The results of the articulation rate analysis of angry speech produced by the male speakers can be summarised as follows:

- Male non-actors (collectively) did not produce angry speech with AR values with a significant difference between baseline speech and threatening speech.

- Male actors (in general) produced their angry speech with an increase in AR relative to the other experimental conditions.

Female speakers

Figure 64 shows the articulation rates for baseline, threatening and angry speech produced by the female non-actors (Speakers 11-20).

Figure 64: Articulation rates for baseline, threatening and angry speech produced by female non-actors

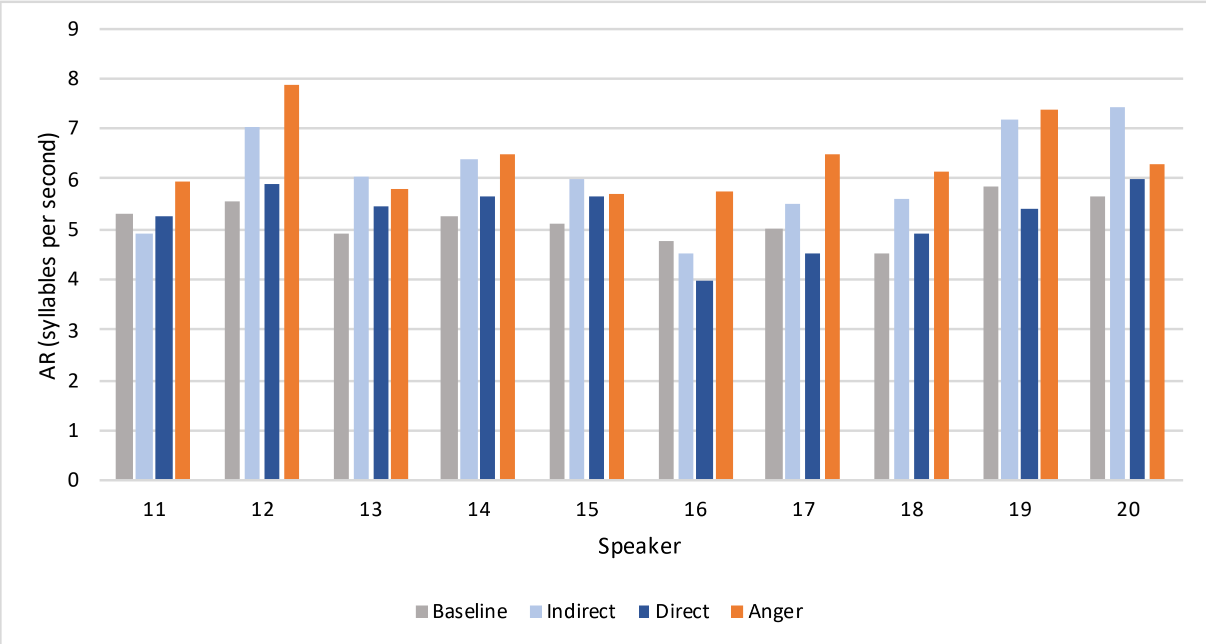

For the female non-actors, all of the anger speech recordings were produced with an increase in AR relative to the corresponding baseline speech. Speakers 12, 14, 17, 18, 19 and 20 all produced ARs which exceed the 4-6 syllables per second range expected for English speakers. There is also a tendency for the female non-actors to deliver angry speech with a greater AR than threatening speech. This is the case for Speakers 11, 12, 14, 16, 17, 18 and 19. Speakers 13, 15 and 20 each produce their angry speech recording with a higher AR than their direct threat, but which is lower than their indirect threat.

A one-way ANOVA test showed a statistically significant result between the ARs produced by the female non-actors across these experimental conditions (F = 16.13, p = <.0001). Further statistical testing showed a significant difference between the AR values measured for baseline speech and anger speech, and baseline speech and direct threats (both Z = -2.80, p = 0.005).

Figure 65 shows the articulation rates for baseline, threatening and angry speech produced by the female actors (Speakers 31-41).

Figure 65: Articulation rates for baseline, threatening and angry speech produced by female actors

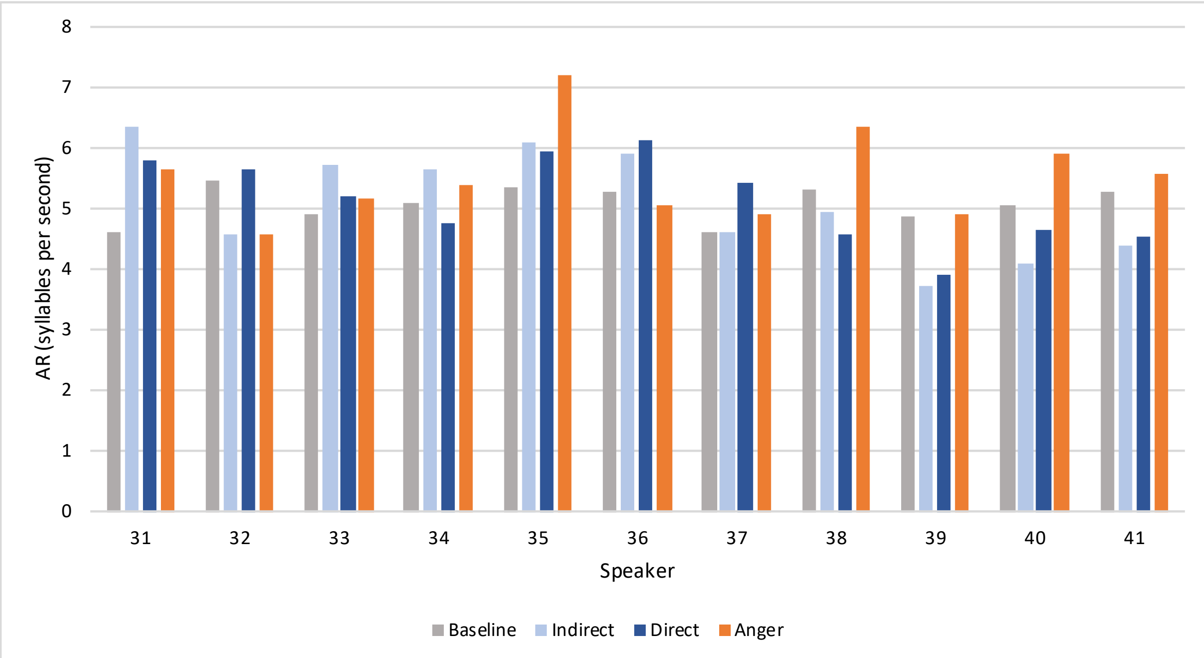

Most of the female actors deliver angry speech with a faster tempo than their baseline speech. Speakers 32 and 36 are exceptions to this trend. This group varies considerably in terms of the relationship between the tempo of angry and threatening speech. Speakers 31, 33 and 36 produce angry speech which is slower than threatening speech of either type. Speakers 32 and 37 produces their direct threats with a faster tempo than their angry or indirect threat speech.

In contrast, Speaker 34 delivers her indirect threat more quickly than her angry or direct threat speech. The remaining speakers (Speakers 35, 38, 39, 40 and 41) produce their angry speech with an increased AR relative to threatening speech. As such, there were more female actors who showed an increase in AR when producing anger (relative to either type of threatening speech), than speakers who showed more varied AR results across these experimental conditions.

A one-way ANOVA test did not show a statistically significant result between the ARs produced by the female actors across these experimental conditions ($F = 0.23$, $p = 0.88$).

The results of the articulation rate analysis of angry speech produced by the female speakers can be summarised as follows:

- Female non-actors (collectively) produce angry speech with AR values which are significantly different to baseline speech and direct threat speech. This was observed to largely manifest as an increase in AR during threatening speech relative to baseline speech.
- Female actors (collectively) did not produce angry speech with AR values with a significant difference to baseline speech and threatening speech.

Summary of articulation rate analysis: Anger

To summarise, there were mostly no statistical differences between the AR values of angry, baseline and threatening speech. As such, there was no clear pattern in the ARs across these experimental conditions and speaker groups. It should be stressed that these findings should be interpreted with caution due to the very short sample duration of the angry speech recordings. Longer speech samples would have enabled more robust AR calculations.

The following section presents the findings of a vocal profile analysis performed on the angry speech recordings collected in Task 4 of the experiment. These voice descriptions will later be discussed in relation to the speaker's corresponding baseline speech and threatening speech.

5.3.4 Vocal profile analysis

As discussed in Chapter 2, angry speech has been associated with jaw and tongue movement. In particular, angry speech has been linked with jaw opening and tongue movement which utilises a greater amount of the space in the speaker's oral cavity (Kienast and Sendlmeier, 2000).

Earlier in this research, it was noted that some speakers appeared to adjust their jaw position when producing threatening speech. However, there did not seem to be any notable tendency for threatening speech to be characterised by changes made to jaw (or tongue) position. This also appeared to be the case across the Task 4 'angry speech' recordings.

These auditory impressions could be supported by measuring the formants of vowels that are of a similar quality (for example, measuring all instances of [i] vowels). The output of these measurements enable the first and second formants (f1 and f2 respectively) of the vowels present in baseline and threatening speech to be plotted in a visual vowel space. This plot would provide some insight into whether the speaker spoke with minimised or extensive jaw and tongue movement. In particular, f1 measurements would relate to impressions of jaw opening. This is because high f1 measurements would suggest the tongue is positioned lower in the speaker's mouth. By opening the jaw, the tongue would be able to move to a lower position than it would otherwise.

Such an analysis was not presented here, due to the lower quantities of speech found in much of these threatening speech data. As such, there many threatening speech recordings which feature very few vowels. This creates difficulties when attempting to make meaningful comparisons between the vowels produced in both baseline and threatening speech. Exploring the extent of jaw movement during threatening and angry speech could form the basis of future research in this area.

5.3.5 Summary and discussion of the comparison of the phonetic features of baseline speech, threatening speech and 'anger' speech

This subsection has presented data which sought to address **research question 1.a** which focused on whether the phonetic features of threatening speech were similar to other emotional states. The motivation for this research question primarily stemmed from previous literature which likened threatening behaviour to angry behaviour. Comparative to threatening speech, angry speech has been the subject of considerably more phonetic research. Therefore, it was of interest to this study to assess the degree to which these behaviours overlapped during vocal production.

It was found that overall, these speakers production of 'anger' appeared to not be notably (or significantly) different from their threatening (or baseline) speech. However, angry speech appeared to be characterised by these speakers by a higher mean f0, intensity and speech tempo relative to baseline speech. This finding is comparable to much of the previous literature on the phonetic production of anger. How angry speech relates to threatening speech appears to be more considerably more complicated. It is possible that by examining longer samples of anger speech (than what has been presented in this sub-section), a clearer picture might emerge regarding how threatening and angry speech production relate to one another.

What these results appear to suggest is that the relationship between productions of threatening language and emotional state is not clear-cut, or could thought of as independent of one another. For example, it is entirely feasible that a speaker producing a threat feels nervous, angry, authoritative (or any other combination of emotional states). In addition, there might also be variation across speakers in terms of how speakers interpret terms such as 'threat' or 'anger'. In other words, the lack of significance noted in this section could be due to speakers interpreting these behavioural or emotional states differently. This would likely result in a lack of consistent results for the broad phonetic properties of speech examined for this research.

5.4 Comparison of authentic and simulated threat data

This study has so far presented the results of a phonetic analysis of both authentic and simulated threats. The data collected from these sources were recorded in very different circumstances and consist of entirely different linguistic content. In response to **research question 1.c**, this subsection aims to focus on comparing the phonetic features of authentic and simulated threats. This subsection is motivated, in part, by the difficulties of accessing relevant, authentic threats which are suitable for detailed phonetic analysis.

From a research perspective, it would be ideal if simulated threats were consistent with what can be observed from authentic threats. In particular, there are ethical and practical difficulties of using authentic materials in production and perception experiments. For example, audio data taken from police records might be provided to researchers on the basis that no one other than the researcher(s) listens to them. Using simulated threats instead could circumvent these restrictions. If the speech data recorded for this research drastically differed from authentic materials, then there would be an argument to avoid making claims about threats based on the analysis of simulated threats.

To this end, the following sections provide a comparison of the authentic and simulated threats collected for this research. More specifically, the f0 and ARs calculated from the authentic threat data will be compared to simulated versions of these threats. As previously described, all of the authentic speakers are male. In addition, some of these recordings were subsequently bandpass-filtered (if they were not originally a telephone call recording). To allow for more robust comparisons to be made between these authentic threats, only Task 5 data produced by the male non-actors (Speakers 1-10) will be presented here. These simulated data were selected for the following reasons:

- **Male speakers only.** This is because all of the authentic speakers are male. As discussed earlier in this research, it is possible that some aspects of threatening speech production differ between the sexes. This subsection intends to focus purely on differences between authentic and simulated materials. As such, including female speakers in this comparison would add a possible confounding factor.
- **Non-actors only.** This is because it is not known whether the people who produced these authentic threats had (or considered themselves to have) acting experience. Based on the findings of this study, there seem to be differences between non-actors and actors particularly in terms of the f0 and intensity of threatening speech.
- **Task 5 data only**. These data consisted of speech that was based on orthographic transcriptions of the authentic threats analysed for this research. Using this data allows for comparisons that are based on the phonetic production of the same (or highly similar) linguistic content. As previously noted in this research, the speakers recorded for this study did not hear the original recordings of these threats.

These simulated data were bandpass-filtered before phonetic analysis.

The following section provides a comparison of the mean f0s recorded for the authentic and simulated data described above.

5.4.1 Fundamental frequency

Figure 66 shows the mean f0 measured for each of the authentic threats collected for this study (red boxplots). In addition, this figure also shows the average mean f0 corresponding to each of these threats as produced by Speakers' 1-10 (blue boxplots).

Figure 66: Comparison of median f0 between each authentic threat and corresponding simulated threats produced by all of the male non-actors

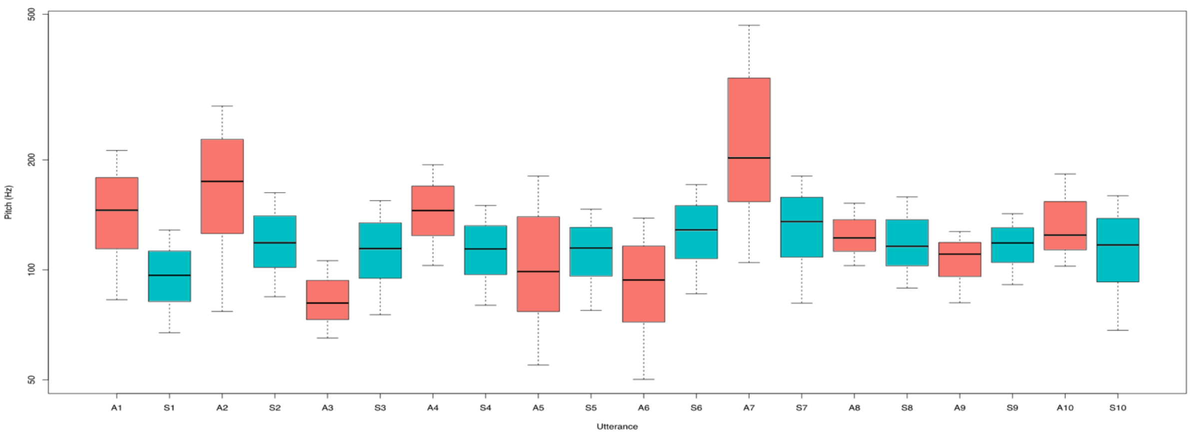

This figure suggests that the simulated data (collectively), is more consistent in terms of f0 than the authentic data. In other words, speakers of real-life threats seem to vary more in their f0 production from one another. On the other hand, these authentic transcripts appear to have been read aloud by Speakers 1-10 with a similar mean f0 to one another.

Another observation is the tendency for the authentic data to be produced with a greater f0 range than the simulated data. This is especially true of speakers A1, A2, A5, A6 and A7. In contrast, the simulated data appears to show less f0 variation for the production of each transcript.

The following section presents the results of articulation rate analysis on these same data.

5.4.2 Articulation Rate

Figure 67 shows the AR calculated for each of the authentic threats collected for this study (red bars). This figure also shows the averaged AR values corresponding to each of these threats as produced by Speakers' 1-10 (blue bars).

Figure 67: Comparison of articulation rate between authentic threats and corresponding simulated threats

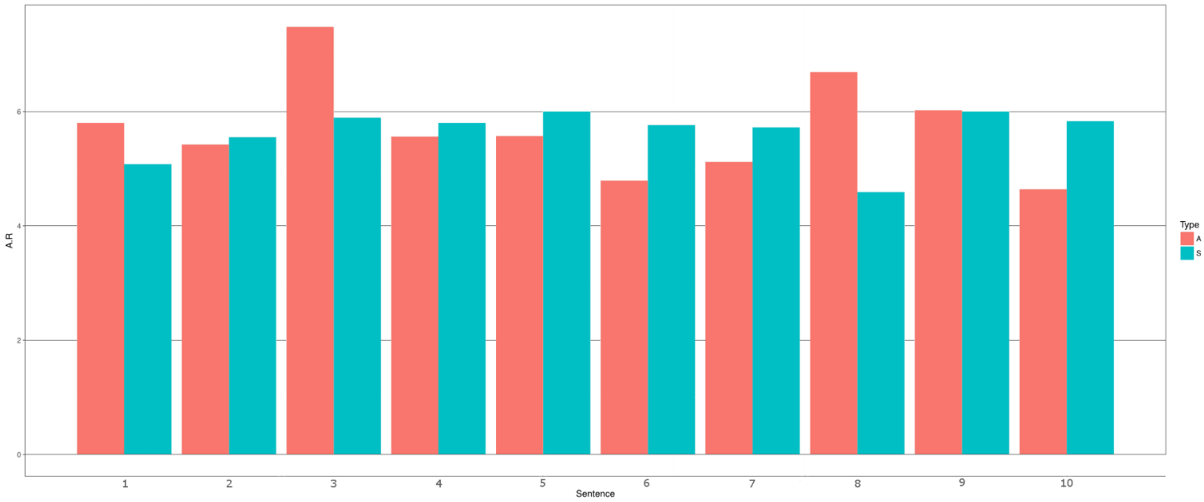

For both groups, these AR results mostly fall within the typical English-language range of 4-6 syllables per second. The exceptions to this are found in the authentic threats produced by A3 and A8. The average ARs produced by speaker in experimental conditions vary in terms of whether they are higher or lower than the original authentic recordings.

The speakers recorded for this study read aloud the following transcripts with an increased AR in relation to the original authentic threat: A2, A4, A5, A6, A7 and A10. The remaining transcripts were delivered at a lower tempo than the original authentic recording.

5.4.3 Summary and discussion of the comparison of authentic and simulated threat data

This subsection aimed to identify the phonetic similarities or differences between authentic threat data and simulated threat data. The results of a phonetic comparison of authentic and simulated threats can be summarised as follows:

- For both the mean f0 and AR results, the simulated data appears to be more consistent to one another.
- These authentic data vary more to one another in terms of mean f0 and AR.

In response to **research question 1.c**, there is evidence to suggest that authentic materials can differ from simulated threats in terms of salient phonetic features such as f0 and tempo. As such, it is plausible that in perception studies, threats sourced from these groups could be distinguished from one another by listeners. However, this has not been explored in the current study.

Therefore, care should be taken when selecting threatening speech data for research purposes. Caution should also be advised when expressing findings based on simulated materials. In other words, while there are benefits to using simulated threat data, researchers should be mindful that these data might perform differently to authentic data. This in turn, should be reflected in any conclusions which were formed on the basis of simulated data. These observations are also applicable to studies which focus on other language crimes.

Nevertheless, this research was motivated primarily by a lack of publicly-available research into the phonetic features of threatening speech. In particular, to further our understanding what how lay-people understand as a 'threatening tone of voice.' This lack of research may be in part due to difficulties in accessing relevant materials. To allow for more much needed peer-reviewed research in this area, simulated threats can provide a starting point.

In light of the findings of this subsection, there are a number of possible ways which could allow simulated threats to be more phonetically similar to the authentic threat data. Using experimental reproductions of threatening language could avoid some of the practical and ethical pitfalls faced when using authentic data. For the current study, participants were not instructed on how to sound 'threatening.' This was in order to study how these speakers interpreted this instruction.

For the collection of threats for use in future studies, participants could be provided with more detailed materials to aid their production. For example, by allowing them to listen to publicly-accessible threatening speech recordings. This could include recordings taken from film and television.

Also, participants could be provided with transcriptions which offer more guidance on how the original authentic recording sounded. For example, the transcriptions could indicate words or phrases which are produced with an increased volume. This could be shown using typographical features (such as underlining text or writing in capital letters). The use of typographical features in the data collected for this research will be presented in Chapter 5.6.

Chapter 4 focused mainly on the form a threat can take, that is to say, the words that are used to convey a threatening intent. Crucially, these data do not reflect threats which are produced spontaneously. Despite this drawback, these data can instead focus on how authors construct the production of threatening language. In other words, these data can provide insight into the deliberate (and perhaps unconscious) choices made by lay-people to appear threatening. Related to the linguistic form of these threats, is the use of typographical features used by authors. These features can modify the phonetic production of a live-reading or the interpretation of the words stated in the threat. To this end, the following subsection presents phonetic data which relates to where authors have included typographical features in their written scripts.

5.5 The relationship between typographical features and subsequent phonetic production

Research question 3 poses whether typographical features present in the written threats collected for this research can be related to the subsequent phonetic production of the threat. These texts were handwritten by authors who were aware that they would later read aloud these documents in a 'threatening' way. The inclusion of typographical features in these scripts might provide some insight into the phonetic (and/or prosodic) features which the author considers to aid his/her production of a threatening message. Compared to the research questions discussed so far, this question has a broader scope than threatening messages (or even other forensically-relevant language). This is because the inclusion of research question 3 was motivated more organically, to investigate how the authors of the scripts collected for this research might use typographical features to convey different phonetic or prosodic cues which had not been the focus of the analysis shown earlier in Chapter 5.

Writers across written media frequently use variations in, for example, font size and bold typing. These alterations to the physical appearance of text can reflect, for example, how speech had originally been produced or ought to be interpreted by the reader. The intended interpretations of these typographical features can, of course, be disputed by readers or listeners. The potential for misinterpreting the pragmatic meaning or intent of scripted speech also has substantial implications in other areas of forensic phonetics. For example, the interpretation of court proceedings by a forensic linguistic researcher could be influenced by how the courtroom transcriber has chosen to capture speech in a written form.

For example, there is potential for inaccurate or unrepresentative transcriptions of earwitness testimony by police or courtroom transcribers. By converting speech into a written format, it is possible that phonetic features are not captured, or are noted inconsistently. This runs the risk of creating multiple possible interpretations of how the speech was originally produced. For example, in a legal case where there is no recording of an alleged threat, a written transcription would be the closest representation of the offence available to a court and/or jury. Reading aloud a potentially inaccurate transcription of an alleged threat in a courtroom setting would present the court and/or jury with evidence which might be misleading. This, in turn, might influence the legal decisions that are made.

In this research, typographical features are defined as changes or modifications that are made to the written font that is otherwise used throughout a text. Examples are the use of *italic font* or CAPITALISATION. The term also refers to additional features such as underlining text or using “quotation marks” to enclose a word or phrase.

Each script produced for Task 3 was manually tagged for the following typographical features:

- noticeable changes in font size throughout the text
- the use of italicised font
- noticeable changes in the boldness of font throughout the text
- capitalisation of words and/or phrases
- underlining words and/or phrases
- the use of quotation marks
- the use of ellipsis

It should be noted that participants were only instructed to write down what he/she *would* say when making a threat(s) in this hypothetical scenario. Participants were not instructed to write down *how* they would read aloud their threats. Therefore, it is perhaps unsurprising that throughout these data there are only a few instances of typographical features. Of the typographical features listed previously, there are no examples of noticeable changes to font size and font boldness or the use of italicised font in these data. There are, however, some examples of the use of quotation marks, ellipsis, underlined text, and capitalised text. Each of these occurrences will be presented below, along with additional phonetic analysis corresponding to these utterances.

The pitch and intensity scales presented for each speaker are relative to his/her minimum and maximum f0 and intensity recorded during Task 1. This is because the Task 1 recording provides a more substantial reference sample of read speech. Therefore, any notable phonetic changes which are seemingly based on typography can be discussed and presented in relation to the speaker’s ‘typical’ f0 or intensity production. In cases where the figures recorded for f0 or intensity are higher during the threatening speech than the corresponding baseline speech, the pitch and intensity scale axes on these figures have been adjusted accordingly. These cases will be flagged up in their respective subsections.

The following subsections discuss the phonetic realisation of texts written using quotation marks, capitalisation of text, underlined text, and ellipsis.

5.6 Quotation marks

The following participants included quotation marks in at least one of their written threat texts: Speaker 3, 25, 31 and 33. This subsection will explore whether the phonetic realisation of these underlined sections of script was consistent across these speakers.

5.6.1 Speaker 3: indirect threat data

There is one typographical feature in Speaker 3's script(s). Within his indirect threat, he writes the following:

'Mr West, it seems you are incapable of taking upon yourself the responsibility of closing down your diabolical "farm."'

In this sentence, the word ‘farm’ is enclosed in double quotation marks. Figure 68 (next full page) shows a waveform, intensity trace and pitch trace of Speaker 3’s production of the phrase ‘diabolical farm.’

Figure 68: Waveform, intensity trace and pitch trace for Speaker 3’s production of the phrase: ‘diabolical farm’

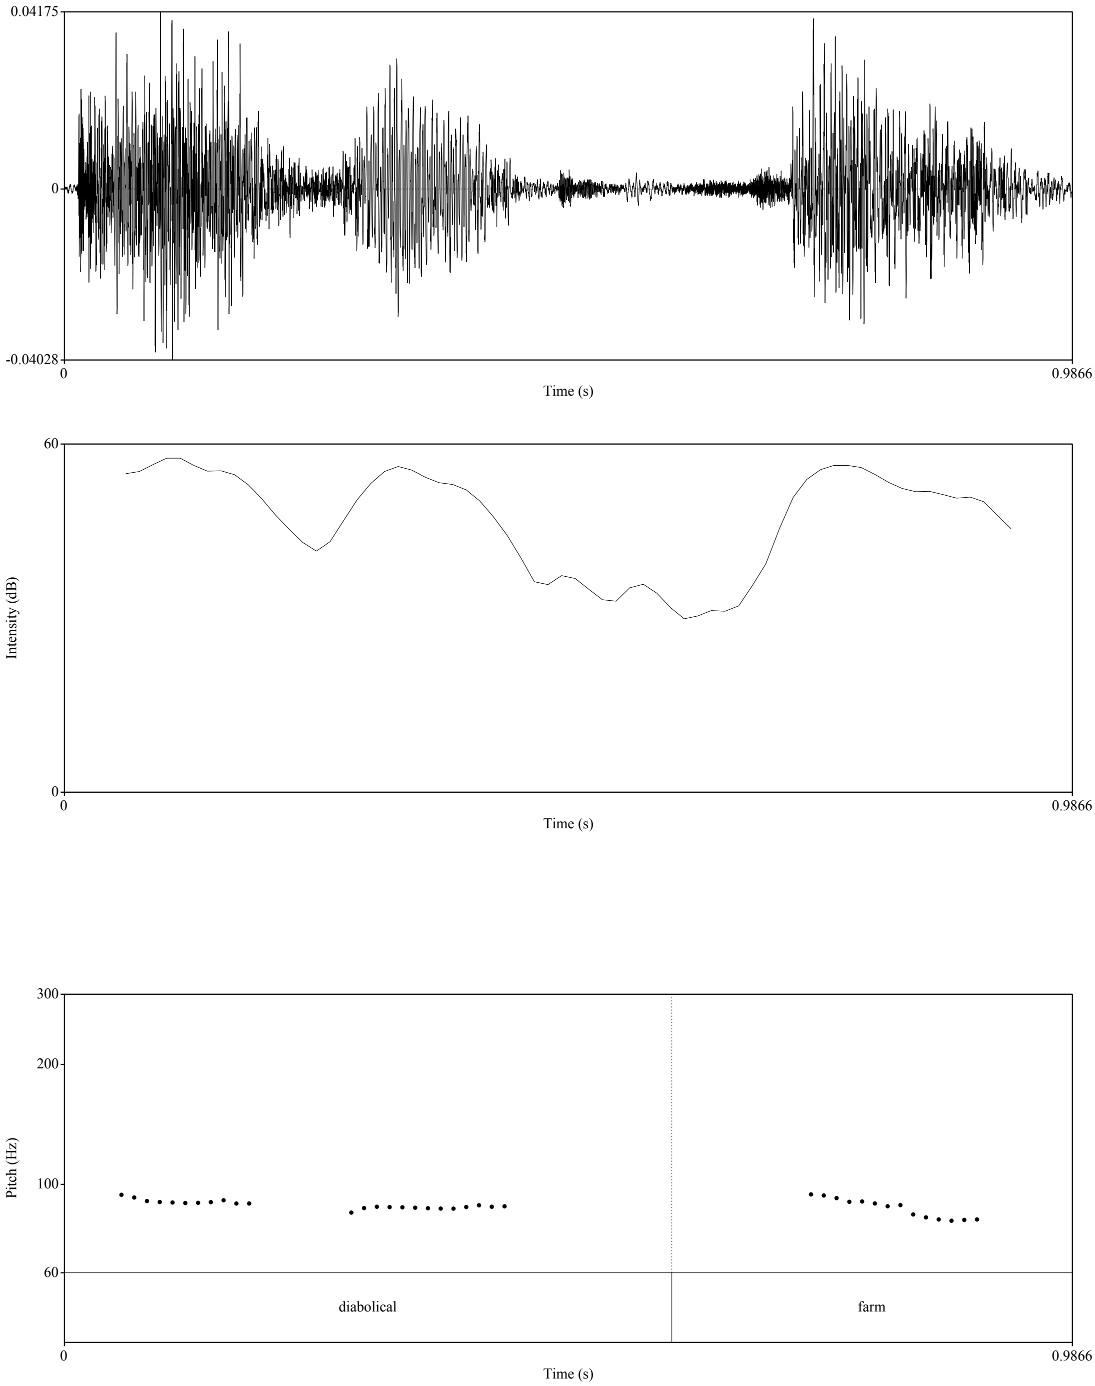

Auditorily, the word ‘farm’ appears to be not noticeably stressed in relation to the other words preceding it. This auditory impression is supported by the acoustic data presented in Figure 68. This is because the intensity and pitch level of the word ‘farm’ appears to be similar to that of the preceding word ‘diabolical.’ In the pitch trace, however, there is a falling contour on the word ‘farm.’ In addition, there is little acoustic or auditory evidence to suggest that there is a noticeable pause between the two words. The inclusion of a pause might have contributed towards an impression of emphatic stress. There is, however, a pause of approximately 0.9 seconds following the word ‘farm’ before the start of the following sentence. It is likely that the pause following the word ‘farm,’ as well as the falling pitch tone on this word, is due the position of this word at the end of a sentence.

Despite the lack of acoustic or auditory evidence to suggest any noticeable effect on production, the author nevertheless chose to include this typographical feature on the word ‘farm.’ It is possible that in doing so, the author is implying is that the farm would be better described in some other way. In other words, there might be an implication that the alleged abuse taking place at this location is such that it cannot be considered to meet the basic standards the author expects of a legitimate farm. Without feedback from the author, however, this is only conjecture.

5.6.2 Speaker 25: direct threat data

Speaker 25 uses one example of a typographical feature in his threatening script. This can be seen in the following extract of text taken from his direct threat:

‘What I do have however (sic) is a very angry brother with a shedload (sic) of rather “explosive materials.”’

In his recording, however, the additional phrase ‘shall we say’ is produced thusly:

‘What I do have however (sic) is a very angry brother with a shedload (sic) of shall we say rather explosive materials.’

Figure 69 shows the waveform, intensity and pitch traces taken from Speaker 25’s production of the phrase ‘shall we say explosive materials.’ In this figure, the intensity level axis has been adjusted from this speaker's baseline speech. This is because Speaker 25 produces his direct threat with a higher intensity than that recorded in his baseline speech.

Figure 69: Waveform, intensity trace and pitch trace for Speaker 25's production of the phrase: 'shall we say explosive materials'

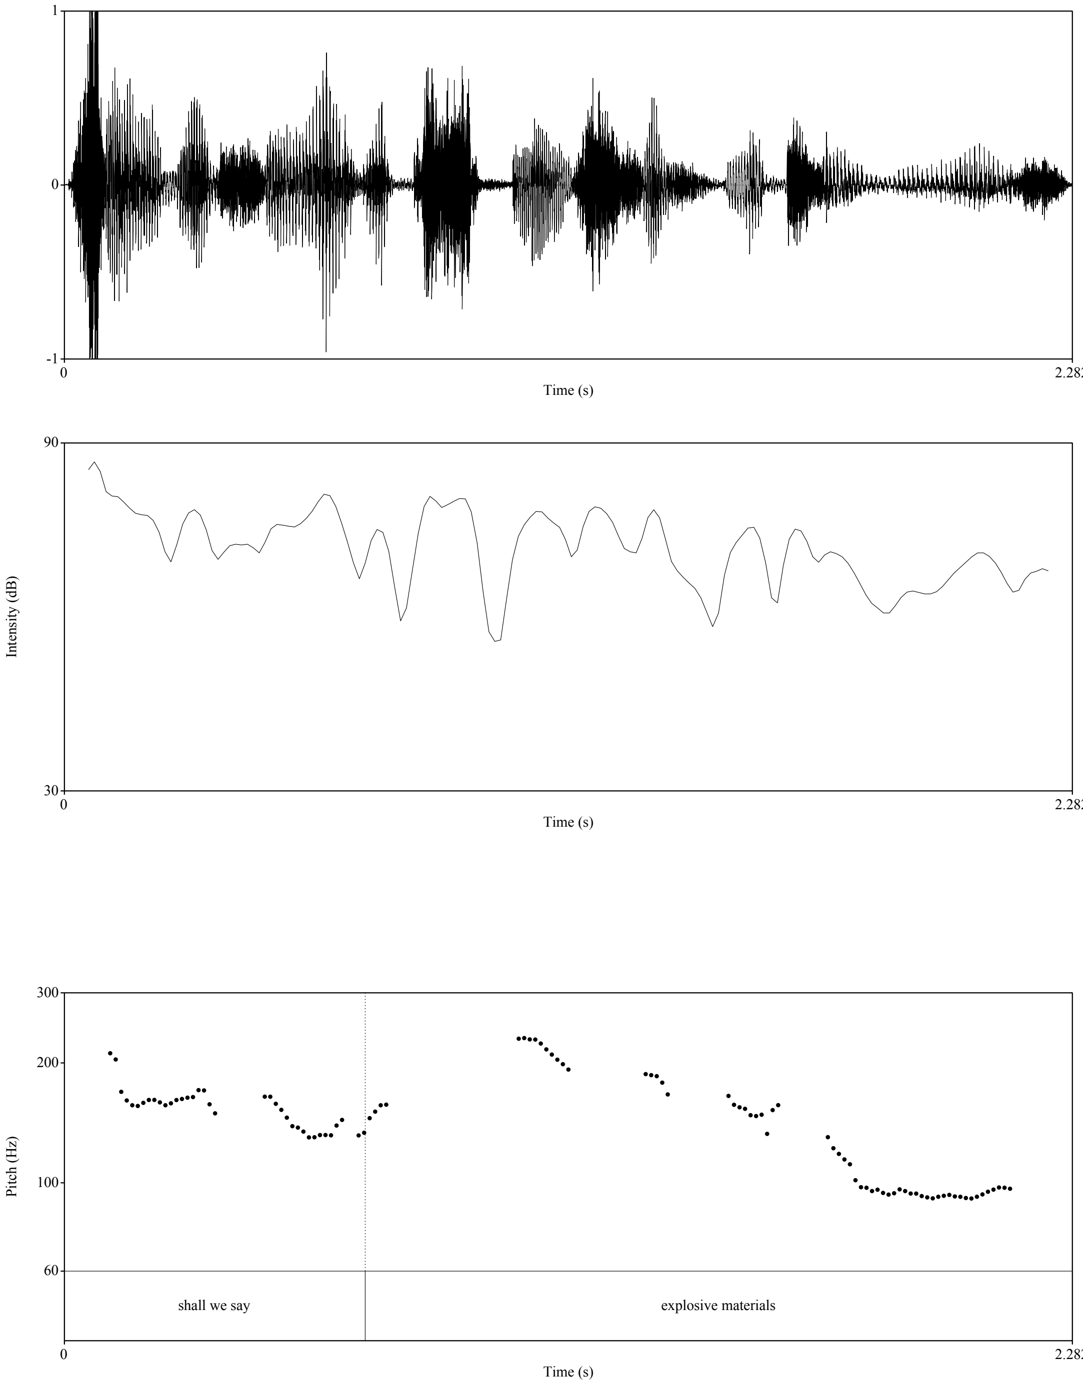

Similarly to Speaker 3, Speaker 25 produces a phrase embedded in double quotations at the end of a sentence. This might account for the falling tone around the word 'materials' as shown on the pitch trace of Figure 69. The phrase 'explosive materials' is produced with a falling tone. Also, the intensity level around of the speech surrounding this phrase is at a similar level to that of the surrounding speech that was not transcribed within quotation marks. Auditorily and acoustically, there appears to be little evidence of emphatic stress on this phrase relative to the surrounding utterance.

This said, this speaker might have attempted to convey some meaning through the use of quotation marks around 'explosive materials.' Possibly, the speaker might have intended to allude to a bomb or similar device, rather than express this information more overtly. In other words, the speaker might be attempting to convey a hint about the nature of the harmful action he 'intends' to commit. As before, the exact reason for embedding this phrase in quotation marks is only known to the author.

5.6.3 Speaker 31: direct threat data

Speaker 31 writes the following in her direct threat text:

'You can't hide behind "Mr West" now, we all know your real name...'

In this sentence, the target's title 'Mr West' is embedded in double quotation marks. Figure 70 shows the waveform, intensity trace and pitch trace taken from Speaker 31's production of the phrase 'You can't hide behind "Mr West" now...' In this figure, the axis for the intensity level has been adjusted. This is because Speaker 31 produces her direct threat with a greater intensity than what was recorded for her baseline speech.

Figure 70: Waveform, intensity trace and pitch trace for Speaker 31's production of the phrase: 'You can't hide behind "Mr West" now...'

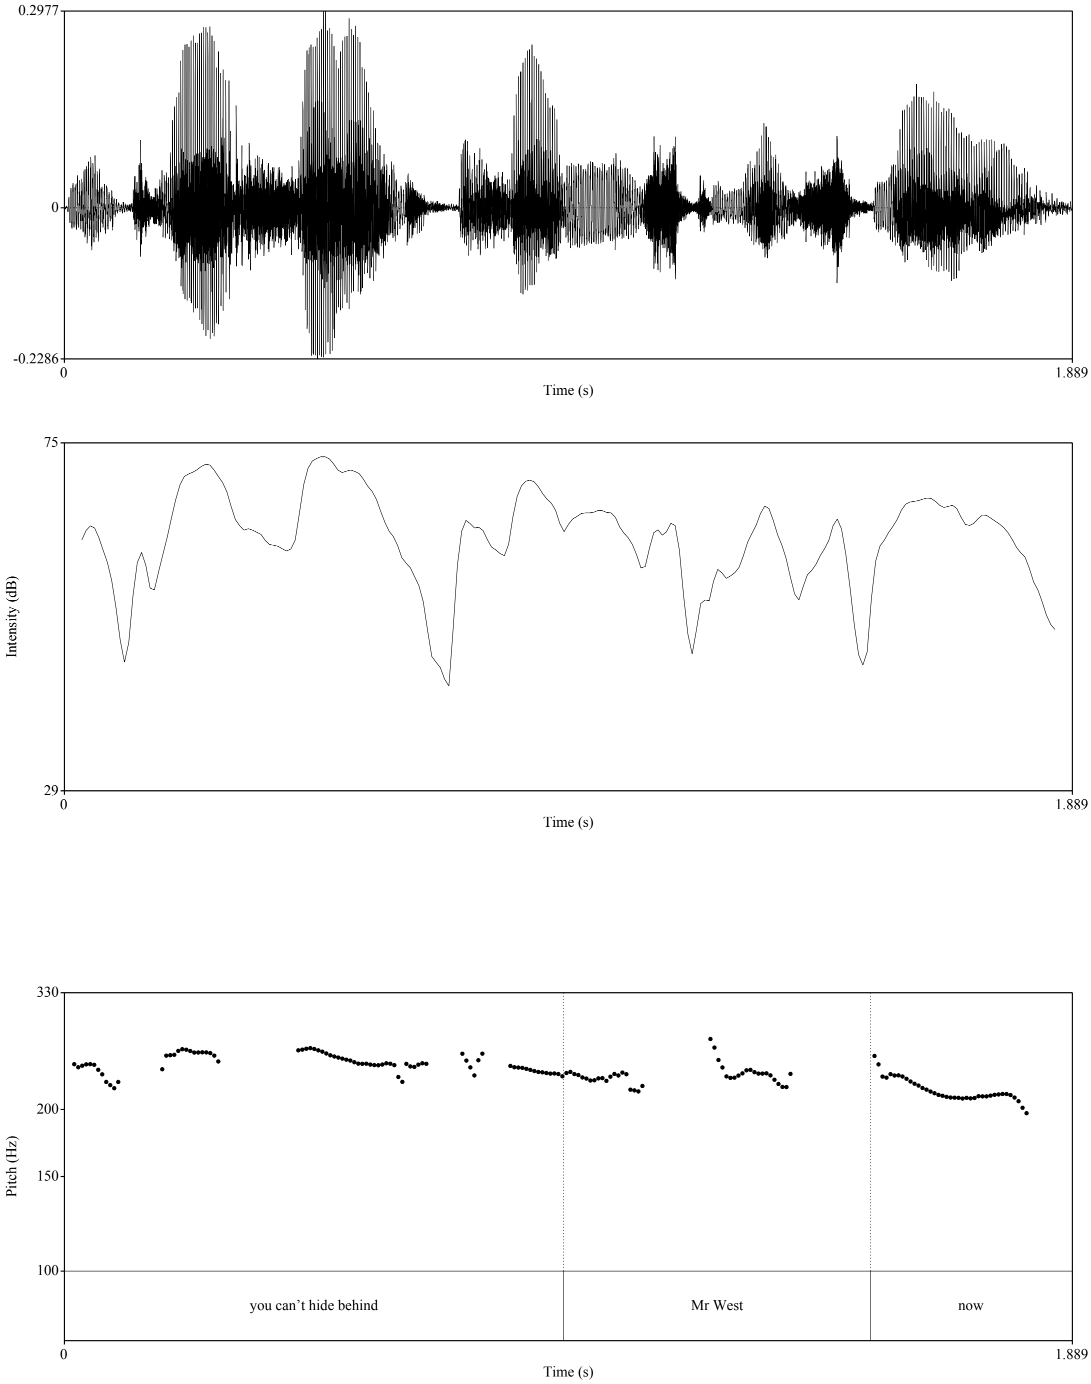

On the basis of auditory and acoustic analysis, there is little evidence to suggest that this speaker produces the phrase 'Mr West' with any notable emphatic stress. Instead, her pitch and intensity level around the production of this word appears to be mostly stable, except for a sharply-falling tone around the onset of the word 'West.' There are also no pauses before or after this phrase which might contribute to an impression of stress. Based on the content of her message, however, it can be inferred that the speaker is suggesting that the target has been concealing his true identity (or his full name). By embedding 'Mr West' in quotation marks, the speaker might be attempting to draw attention to what she considers to be a 'false' (or misleading) identity.

5.6.4 Speaker 33 data: direct threat data

Speaker 33 provides multiple examples of additional typographical features in her direct threat script. Firstly, in the phrase:

'With regards to your message I would like to inform you that you cannot "chose"(sic) whether you close down the farm or not.'

Here, the word 'choose' is enclosed in quotation marks. Figure 71 shows the waveform and spectrogram (Figure 72) for an excerpt taken of this utterance.

Figure 71: Waveform, intensity trace and pitch trace for Speaker 33's production of the phrase: 'cannot "choose" whether you close...'

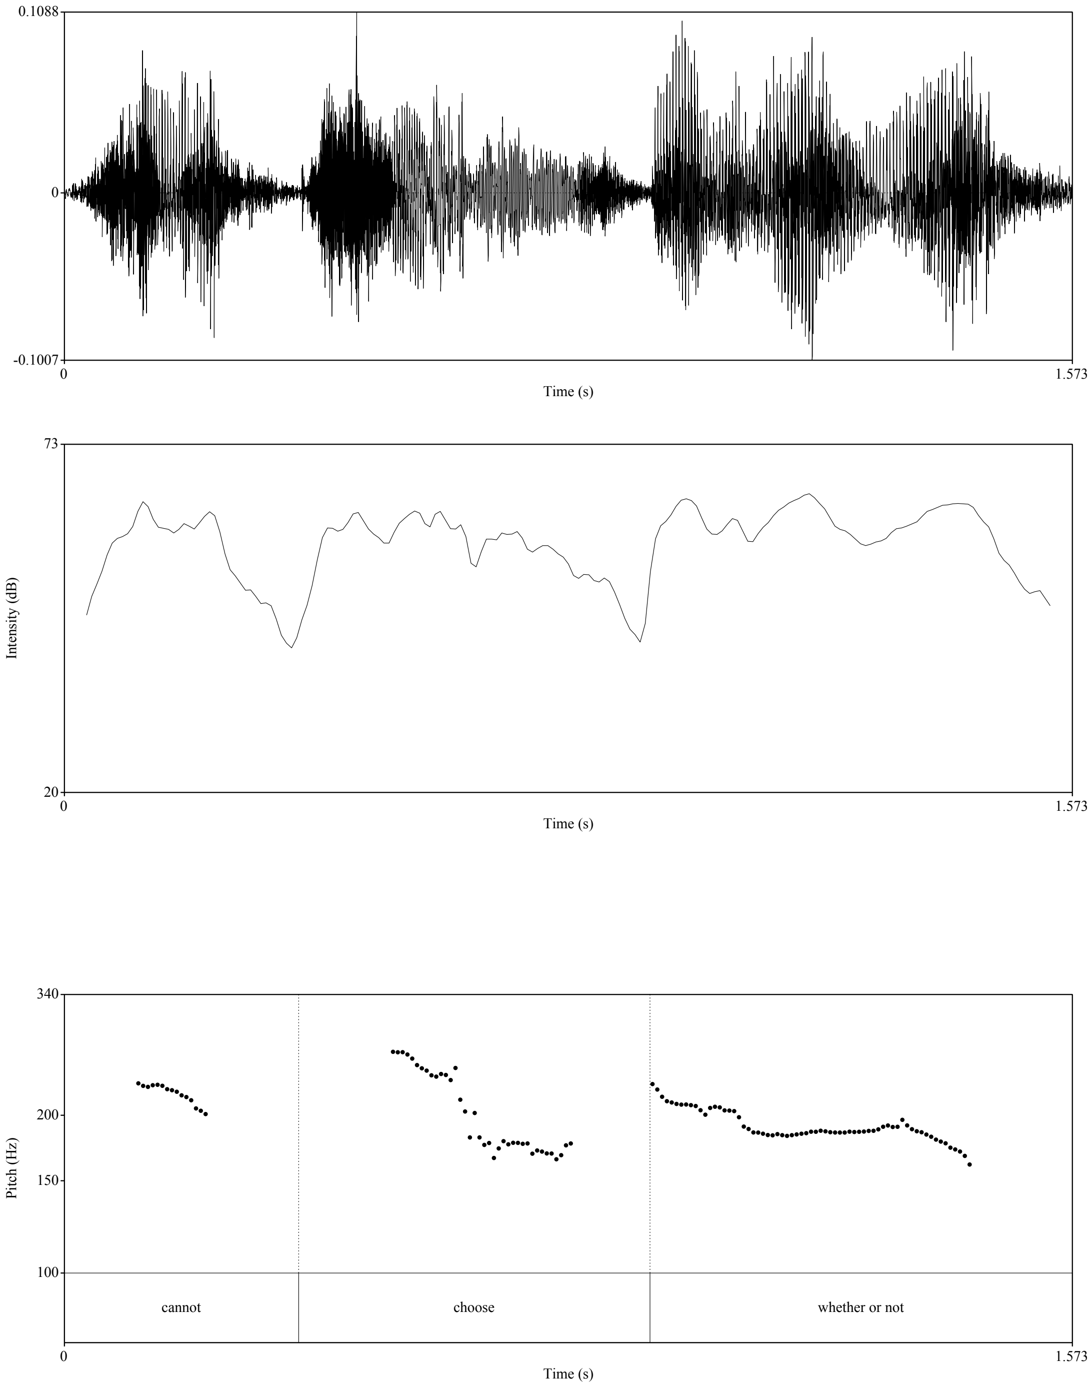

There are some prosodic features in this sample which suggest that the word ‘choose’ is prosodically prominent in this utterance. Firstly, this word is produced with a falling tone from a pitch level that is relatively high for this speaker. Secondly, that the vowel in ‘choose’ is notably elongated. This is shown in closer detail in Figure 72.

Figure 72: Spectrogram of Speaker 33’s production of the word: ‘choose’

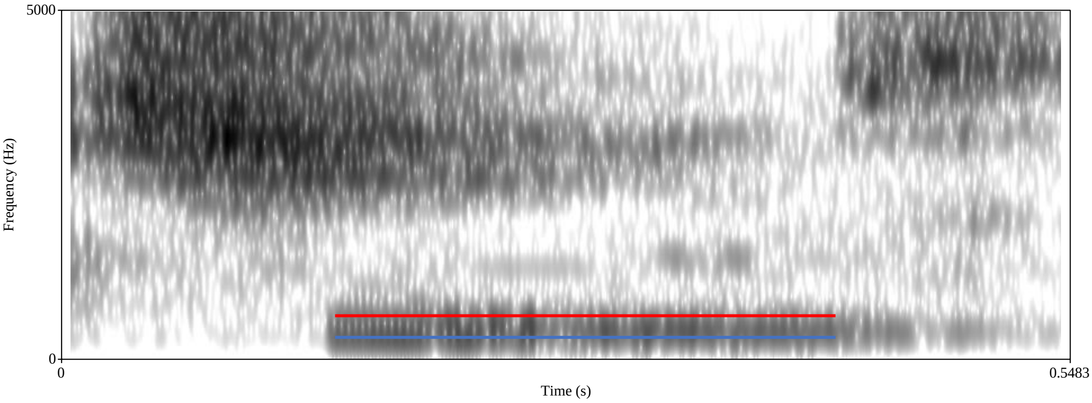

In this figure, there are two lines to mark the approximate location of the first formant (blue) and the second formant (red). The location of these formants, as well as auditory analysis, are consistent with the production of a rounded, close, back vowel [u]. This vowel is sustained for a relatively long period of around 0.28 seconds. In addition, this word is followed by a pause of around 0.3 seconds before the next utterance begins. These factors appear to contribute towards an impression of prosodic stress on the word ‘choose.’

It is possible that the speaker stresses this word to emphasise to her target that closing down the farm is no longer an option. Instead, closing down the farm is framed as a necessary action that the target ought to take or might otherwise be forcibly taken by the threatener.

The same typographical feature is also seen in this speaker’s indirect threat script:

‘...everybody else who opposes your cruelty will use our weapons in order to end what you are doing, whether you have ‘chosen’ to or not.’

This time, the word ‘chosen’ is enclosed in quotation marks using single quotes as opposed to double quotes.

Figure 73: Waveform, intensity trace and pitch trace for Speaker 33's production of the phrase '...have 'chosen' to or not.'

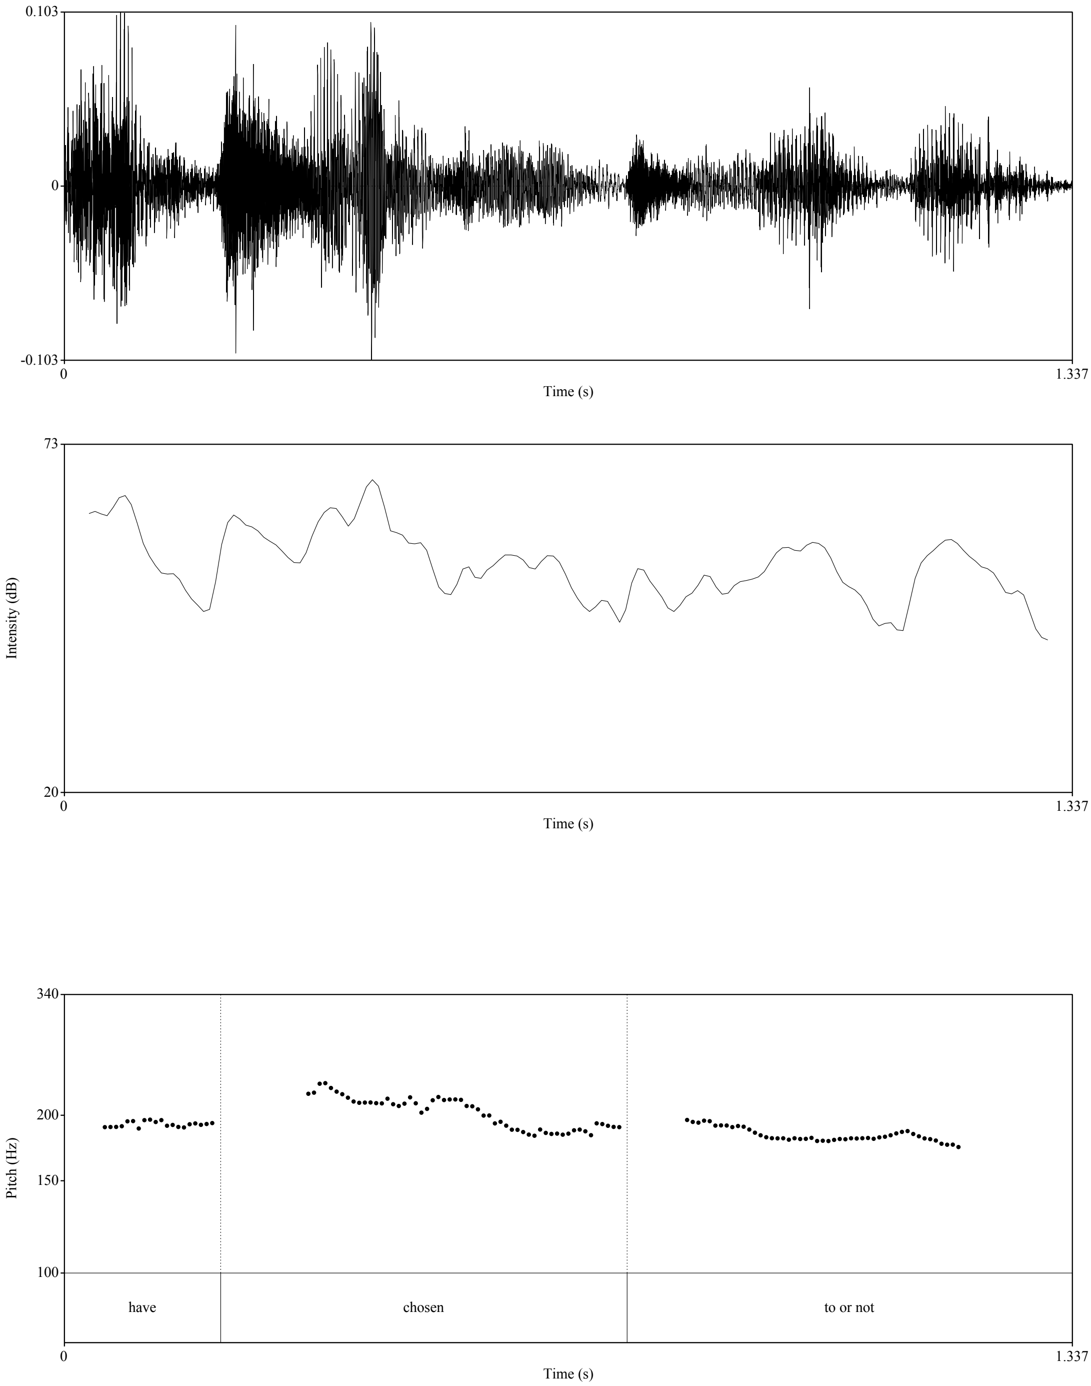

Auditorily, the word ‘chosen’ sounds the most prominent within this sentence. In particular, this word sounds louder than the surrounding utterance. This impression is supported by a peak in amplitude produced around the timing of the first vowel in this word. This peak can be seen in both the waveform and the accompanying intensity trace (Figure 73). The preceding voiceless postalveolar affricate also appears to generate a relatively high amount of turbulent noise. This is shown in the following figure (74). The leftmost dark section of this spectrogram is indicative of higher levels of energy being generated during the production of this consonant.

Figure 74: Spectrogram of Speaker 33’s production of the word: ‘chosen’

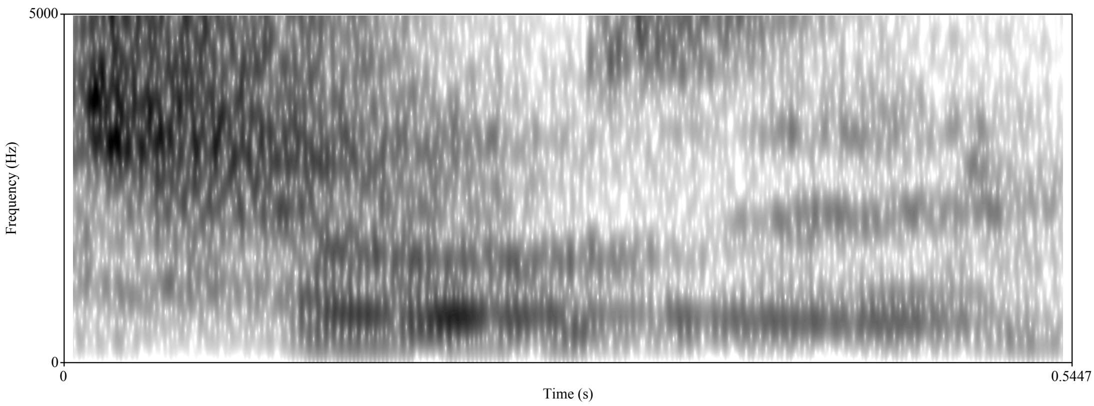

To summarise across these speakers, there appears to be no single, consistent phonetic realisation associated with the use of quotation marks. For Speakers 3, 25, and 31, there appeared to be little evidence of a specific phonetic realisation for the production of words written with quotation marks which was distinct from accompanying words and/or phrases which were not embedded in quotation marks. For Speaker 33, her production of the word ”choose” appeared to be produced with a notably falling tone, as well as elongation of the segments [ʧ] and [u].

5.7 Capitalisation of text

There was only one example of text written in a capitalised font across the written data collected for this research. In addition to the text embedded in quotation marks (described previously), Speaker 33's direct threat script also included an example of capitalisation of a word in the following sentence:

'Closing down the farm is the ONLY option."

Here, the word 'only is written using capital letters, while the rest of the text uses capital lettering only when beginning a new sentence. Figure 75 shows the waveform, intensity trace and pitch trace taken from Speaker 33's production of the phrase '...the ONLY option'

Figure 75: Waveform, intensity trace and pitch trace for Speaker 33's production of the phrase '...the ONLY option.'

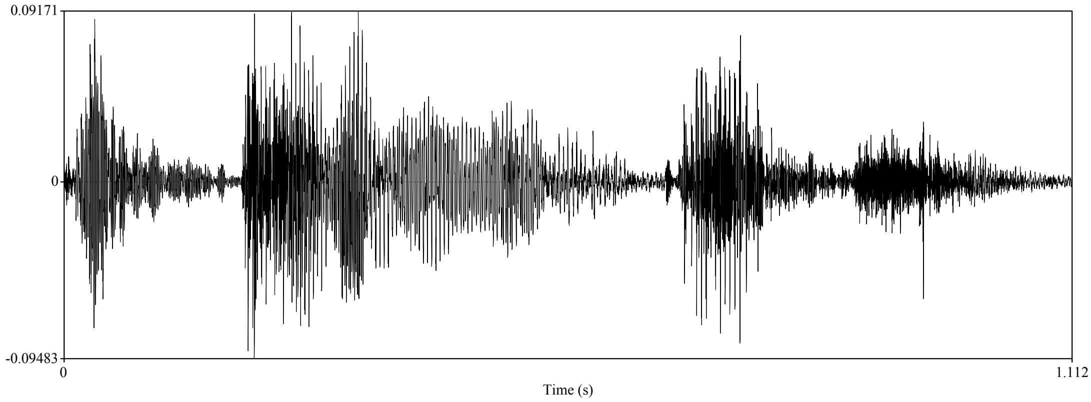

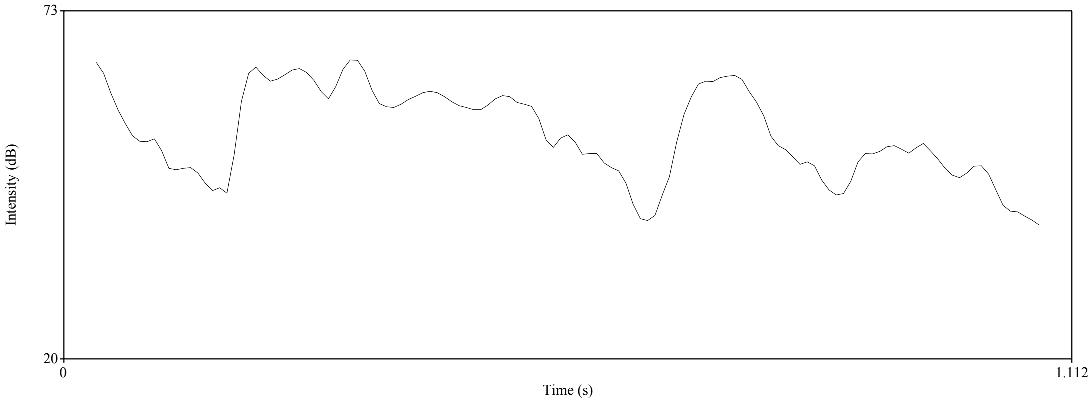

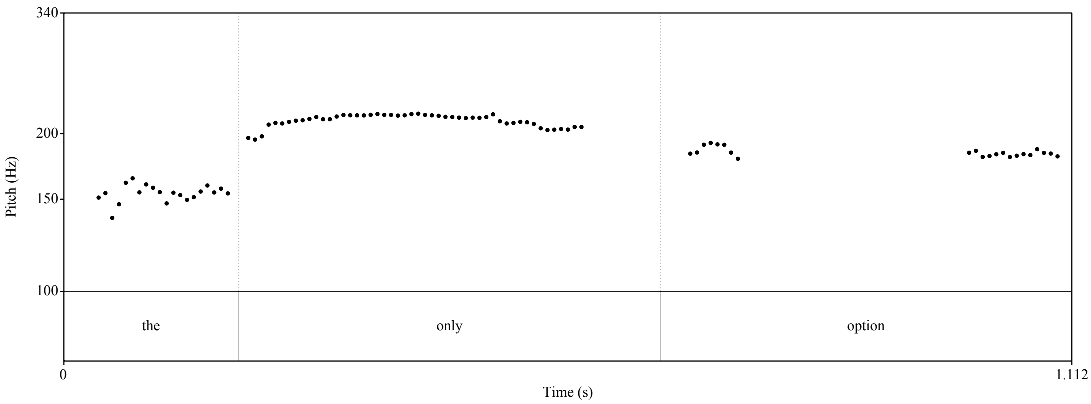

The word 'only' also sounds more prominent in relation to the surrounding utterance. This prominence seems to be achieved in a number of ways. Firstly, the speaker elongates the first vowel in this word, to approximately 0.27 seconds in duration. In addition, as shown in the pitch trace, the word 'only' is produced with an increased and mostly level f0.

This phonetic production might emphasise to a listener that there is now no other option available to the target other than to close down the farm. In other words, by placing stress on the word 'only,' the target might infer that he has no other choice but to give in to the threatener's demands. As before, it is not known whether this speaker intended the suggested interpretation based on her use of typography and/or prosodic stress.

5.8 Underlined text

This section explores the phonetic production of text where specific words have been underlined by the author. This typographical feature was seen in the written threat texts provided by Speakers 18, 38, and 40.

5.8.1 Speaker 18: direct threat data

There is one example of a typographical feature in Speaker 18's threatening texts. This can be seen in her direct threat:

'...you can expect to feel the pain <u>very</u> keenly.'

In this sentence, the word 'very' has been underlined by the author. Figure 76 shows a waveform, intensity trace and pitch trace of Speaker 18's production of the phrase '...pain <u>very</u> keenly.'

Figure 76: Waveform, intensity trace and pitch trace for Speaker 18's production of the phrase: 'pain very keenly'

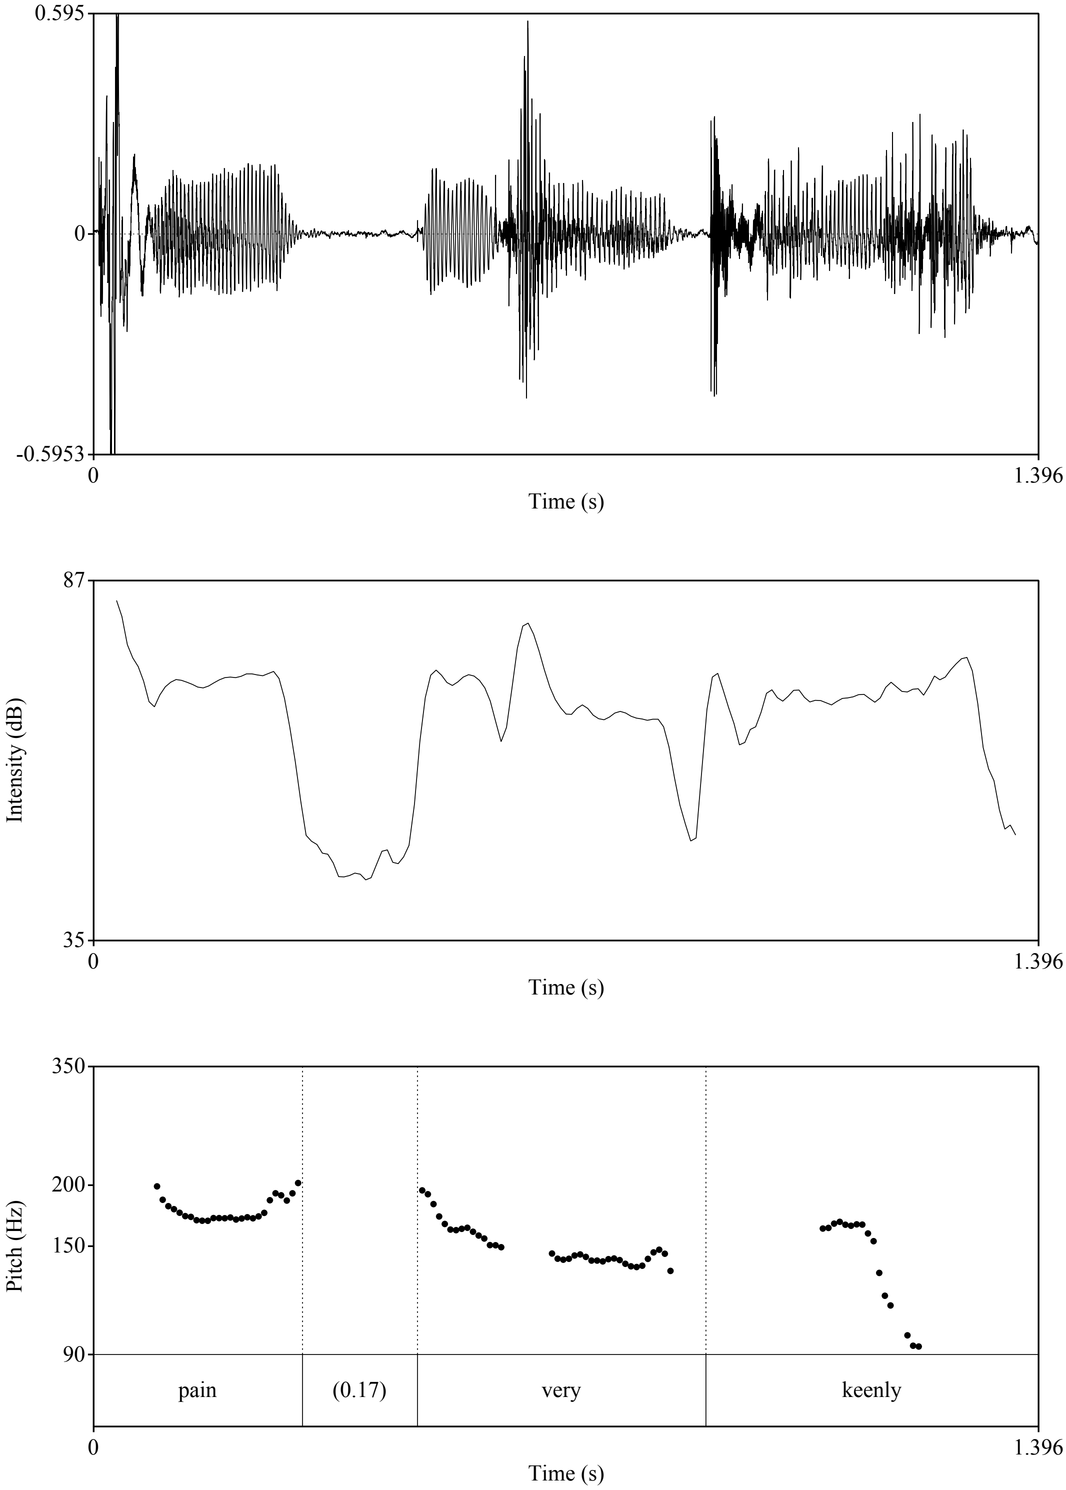

Auditorily, speaker 18 appears to place emphatic stress on the word ‘very.’ There seems to multiple pieces of evidence to support this impression. Firstly, there is a pause before the production of the word ‘very.’ of approximately 0.14 seconds. Following this pause, there is a relatively long period of frication (of around 0.12 seconds) before the release of the initial consonant [v]. This can be seen in the section of periodicity present in the waveform at the beginning of the production of ‘very.’ In addition, there is a pitch trace which further supports that there is voicing around this time.

Upon releasing this initial [v] consonant, there is a spike in intensity and f0. This word is then produced on a falling tone. There is also a short pause of around 0.06 seconds after this word is produced before next word (‘keenly’) is produced. These factors appear to contribute towards the impression that the word ‘very’ being auditorily prominent in this utterance.

As before, it can only be speculated about why this word, specifically, is underlined. It is possible that the reader might infer that the harm specified is even more severe compared to just writing the adverb ‘very.’

5.8.2 Speaker 38 data: indirect threat data

Speaker 38 underlines the word ‘everyone’ in her indirect threat text in the following sentence.

‘You are a sick human being and <u>everyone</u> will know it.’

The corresponding waveform, intensity trace and pitch trace are shown in Figure 77 on the following page.

Auditorily, there appears to be some emphasis on the word ‘everyone,’ particularly on the first syllable. This emphasis might be achieved, in part, by a short pause of approximately 0.07 seconds between the preceding word ‘and’ and the underlined word ‘everyone.’ From examination of the intensity and pitch traces, however, there seems to be little evidence of an increase in intensity or f0 in the vicinity of this word’s production.

By choosing to underline ‘everyone,’ the speaker might be suggesting that the target will now be at risk of harm from innumerable people and not just the speaker in isolation. Claiming that there is an undefined or unlimited number of people who could harm the target could be an effective strategy in inducing fear in a target. This is another area of threat assessment that could benefit from further linguistic research.

Figure 77: Waveform, intensity trace and pitch trace for Speaker 38's production of the phrase '...and everyone will know it.'

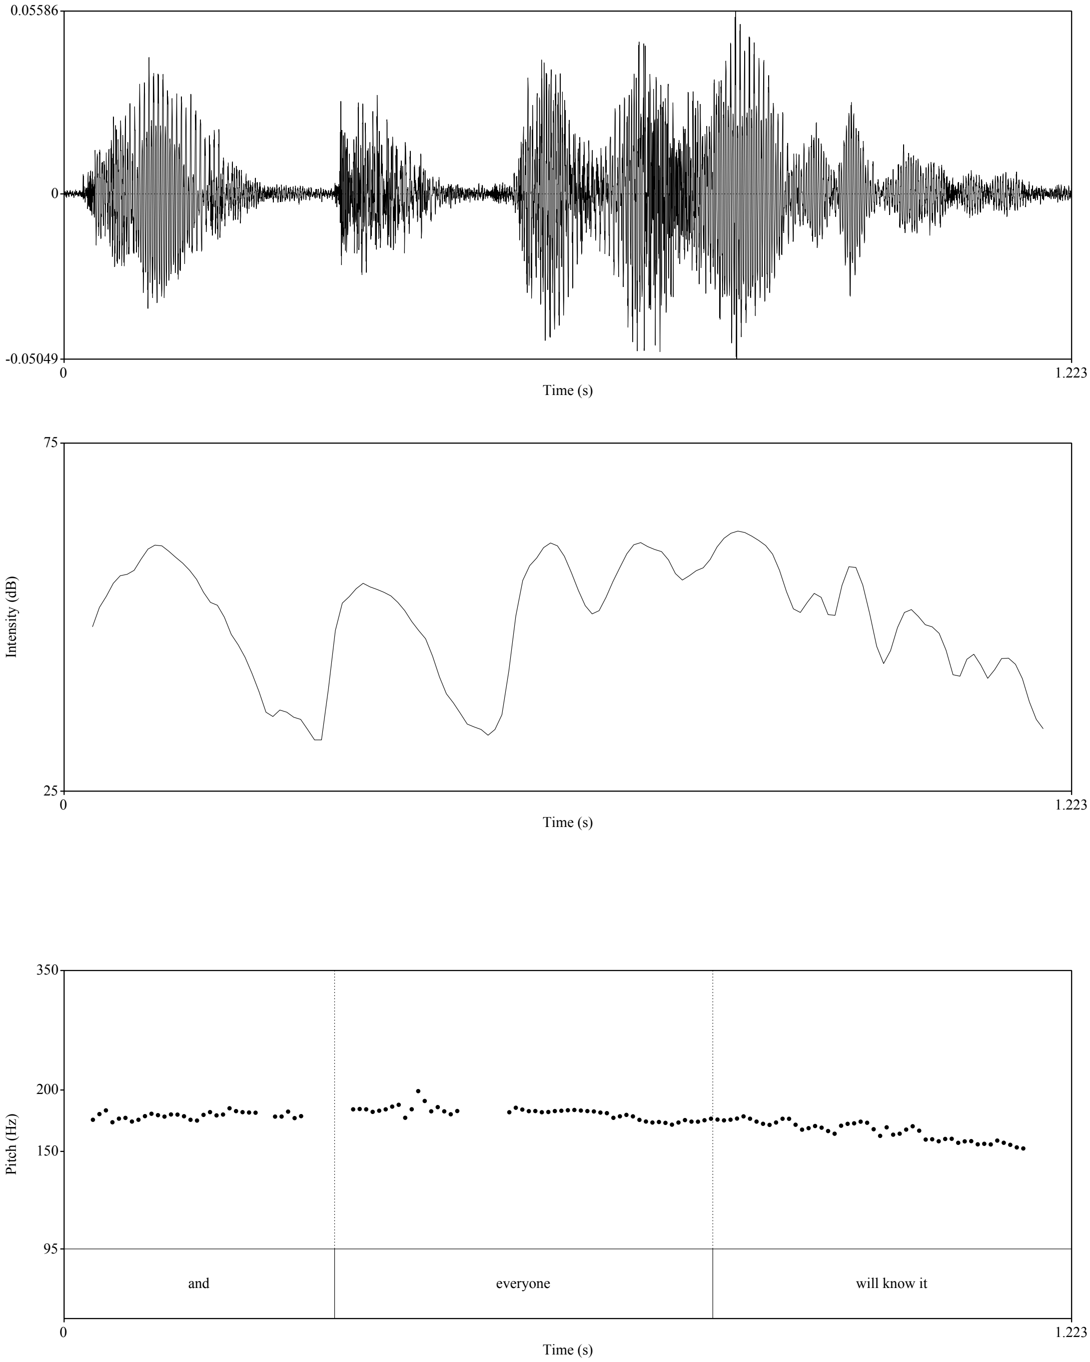

5.8.3 Speaker 40 data: indirect threat data

There is one use of a typographical feature in Speaker 40's threatening texts. This can be seen in the following extract taken from Speaker 40's indirect threat:

'...and believe me - you <u>will</u> get punished.'

Figure 78 on the following page shows the waveform, intensity trace and pitch trace of the phrase 'you will get punished.'

In this example, the underlined word 'will' is auditorily prominent. This impression is supported by an increase in both intensity and pitch, as shown in Figure 78. In addition, the following word 'get' is also similarly prominent. This prominence seems to be achieved by an increase in intensity and f0. In addition, the word 'get' is preceded by a short pause of around 0.24 seconds. Notably, this word is not marked typographically, but is produced with comparable impressions of increased loudness and pitch.

As such, similarly to Speaker 37, Speaker 40 appears to use a typographic feature to relate to some phonetic realisation. However, there is also evidence that Speaker 40 produces other words in the same recording that are prosodically similar without underlining the corresponding text. It is plausible that this speaker chose to underline 'will' to further emphasise the likelihood or certainty that her threatened 'punishment' will occur. As before, the speaker's actual motivation for writing this word with an accompanying typographical feature is unknown.

Figure 78: Waveform, intensity trace and pitch trace for Speaker 40's production of the phrase '...you will get punished.'

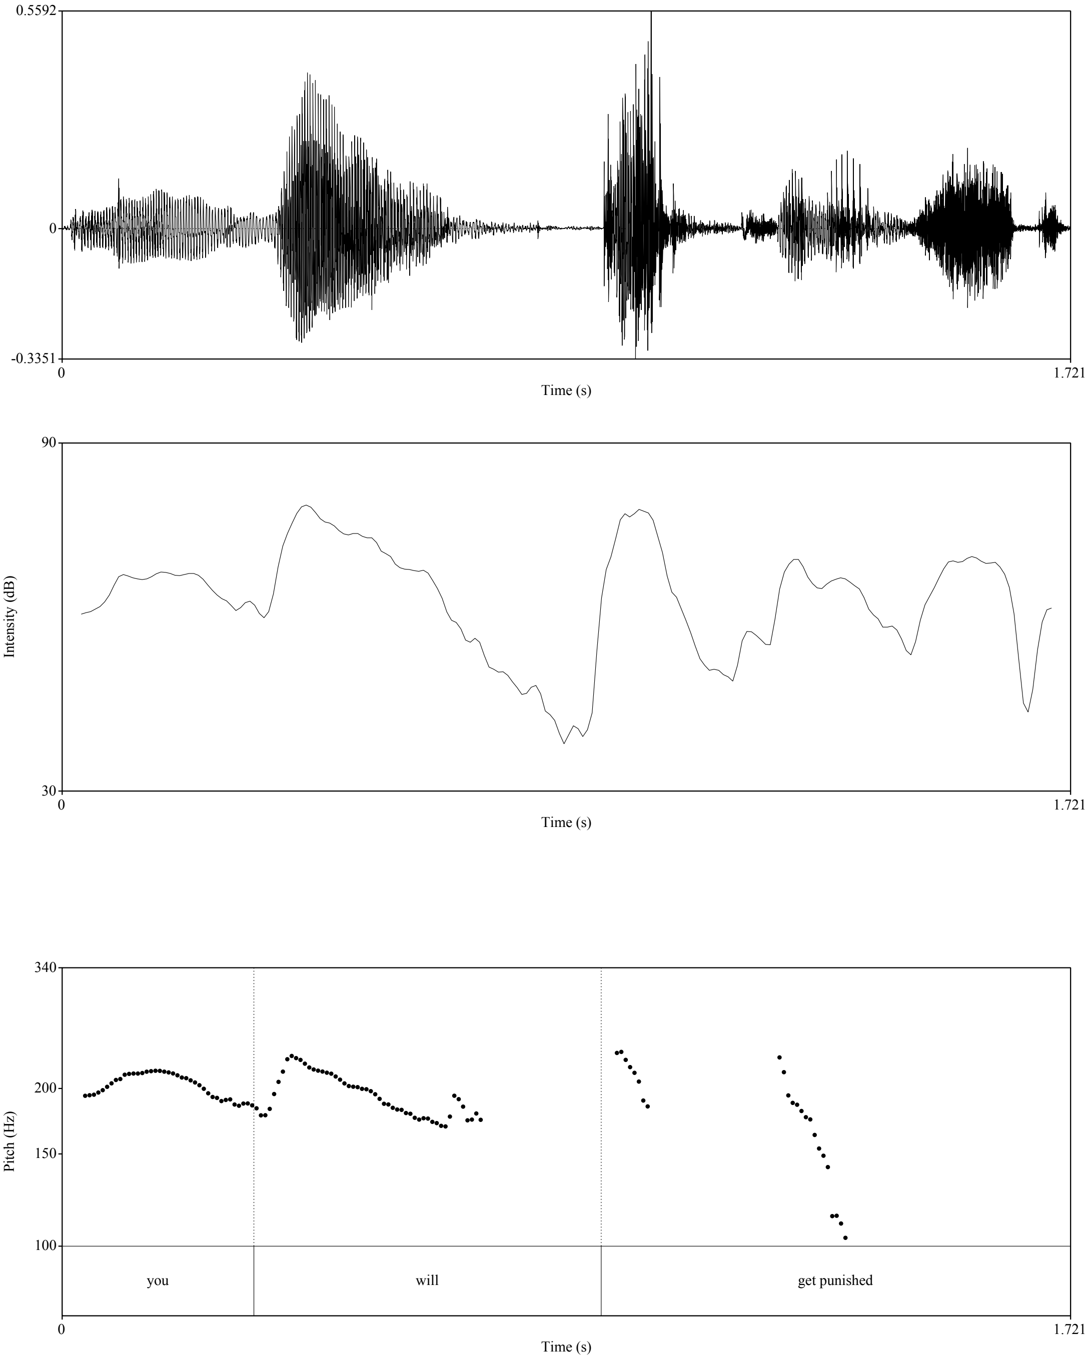

To summarise, for the speakers who included underlined texts in their written threat scripts, there is some evidence that the use of this typographical feature led to (or inspired) a difference in phonetic production (relative to the surrounding utterance). For both Speaker's 18 and 40, the underlined text seemed to correspond with increases in intensity (as well as f0 in the case of Speaker 40). However, for Speaker 38, there does not seem to be much evidence to support a comparable rise in either intensity or f0 during the production of the underlined word.

5.9 Ellipsis

The final typographic feature seen across the written data collected for this research was the use of ellipsis in Speaker 37's indirect threat script.

5.9.1 Speaker 37 data: indirect threat data

This participant uses a series of ellipses during her written indirect threat. These occur in the following sentence:

'I'm sure you, for instance, would hate to be trapped in a shed all day...no light...no food...tortured and cold...imagine it Mr West.'

It was expected that these ellipses would correlate with unfilled pauses, in particular pauses which were notably longer in duration compared with the surrounding speech data (that did not include ellipses). The following transcription provides the approximate pause duration between each of these phrases (measured in seconds):

'I'm sure you, for instance, would hate to be trapped in a shed all day (0.64) no light (0.27) no food (0.27) tortured and cold (0.68) imagine it Mr West.'

From this transcription, there appears to be some evidence that the use of ellipses corresponded to unfilled pauses. However, in this recording there are numerous other examples of unfilled pauses that are comparable in duration. This includes (but is not limited to) the following example, which occurs earlier in this same recording:

'...and I have just one question? (0.63) What if the roles were reversed? (0.47) I'm sure you Mr West...'

These data suggest that while Speaker 37's use of this typographical feature is consistent with pauses in speech, it is not consistent with *all* of the pauses she makes that are similar in duration.

The use of pauses between each of these phrases could feasibly be a way for the speaker to emphasise each of her points in turn by allowing more time to process them as questions. It is also plausible that a different speaker reading aloud this text might choose to include longer pauses where ellipses are present, in relation to where they are absent. The effect of this typographical feature on the reproduction of forensically-relevant speech data is not known.

5.10 Summary and discussion of the effect of typographical features on phonetic production

This section sought to align the written form of a threatening text to its subsequent phonetic realisation. When examining the threatening texts, it was apparent that some participants had chosen to include additional typographical features. It was possible that these authors included such features in their written texts as a means to direct their subsequent phonetic production.

On the basis of auditory and acoustic analysis, it appears that these speakers differed in terms of how words written with these typographical features were phonetically produced. In other words, for these speakers, it would appear that the same typographical feature (e.g. underlining text) was not associated with the same properties of phonetic realisation. In addition, for a number of speakers (Speakers 3, 25, 31 and 38), there was not a clear difference in the phonetic realisation of text written with typographical features, relative to their reading of text written without these features. For the remaining speakers (Speakers 18, 33, 37 and 40), there is more compelling acoustic evidence to suggest that typographical features were linked with a particular vocal production.

It should be stressed, however, that the majority of the texts collected for this research included no typographical features. It seems likely that these participants did not consider including these written details as a means of assisting (or directing) their phonetic production of the text(s). Of the few participants who did include typographical features in his/her threatening text(s), most were female actors. It is possible that had more males produced written scripts (or more detailed texts), these might also include typographical features to aid their phonetic production. In addition, there were no examples of changes to the author's font size or font boldness available in these data. As such, the effect of these factors on vocal production cannot be examined using the present data.

These data suggest, then, that there appears to be no consistent pattern in terms of how authors of threatening scripts use typographical features. This finding has considerable implications for the reproduction of written threatening scripts, as well as transcribed voice evidence more generally. As previously mentioned, it is conceivable that multiple readers of these texts might produce noticeably different phonetic realisations of the same script. This is an area of overlap between forensic linguistics and forensic phonetics which could provide ample opportunities for collaborative research. For example, it would be useful to explore how actors, jury decisions or threat assessors are influenced by changes to typography. This is because inferences about a threatening intent, for example, might be formed (in part) by impressions relating to vocal production.

6 Conclusions

6.1 Overview of research findings

This research has aimed to provide a broad overview of the phonetic and linguistic features of threatening language. The overall picture is that speakers delivering threatening speech do not appear to consistently adopt what might be loosely described as a 'threatening tone of voice.' This impression is supported by a lack of statistically significant differences between the mean fundamental frequency, intensity and articulation rate values measured from threatening and non-threatening speech data. In contrast to previous theories and research relating to threatening speech (Ohala, 1984; Mileva et al., 2017), these speakers (as a whole) instead mostly opted for non-significant increases in mean fundamental frequency and intensity. In addition, threatening speech was overall produced with an increased articulation rate (or tempo), although this difference was also non-significant. These findings then raise the possibility that speakers produce threats with individual variation, rather than being consistent with one another. However, this research has not explored the reasons or motivations that might account for this possible phonetic variation for threatening speech.

In addition, this research has examined the similarities and differences between simulated threatening texts and previous research into authentic (or real-world) threats. There were a number of similarities noted between these types of threatening data: the high usage of the second person pronoun 'you', the lack of words present which express doubt, and the expression of tense and affect. There were also several inconsistencies noted between these data and previous descriptions of threatening language as a predictor of harm. These inconsistencies included: similar word counts between direct and indirect threats, the lack of assertive statements and the lack of detail surrounding the timing, location and nature of the harm threatened.

Finally, this research sought to identify areas of overlap between the forensic phonetic and forensic linguistic domains. One such area is the vocal production of written texts; in other words, how speakers base their phonetic production of written texts using written features. In this study, a number of speakers hand-wrote their threatening messages including typographical features, such as underlined text or text enclosed in quotation marks. A series of prosodic analyses were performed on the utterances which corresponded with these typographical features. These analyses aimed to examine the vocal production of text which had been modified by typographical features. The results of these analyses suggested that these speakers were inconsistent in their vocal production of text which had been written using typographical features.

The current study posed three main research questions concerning the production of threatening language. This section will address the extent to which these research questions (and their corresponding sub-questions) have been addressed.

Research question 1:
Firstly, this research sought to examine whether there was any phonetic basis to support a 'threatening tone of voice'. Recordings were made of a large number of speakers producing baseline and threatening speech under experimental conditions. This experimental set-up aimed to collect data which represented how these speakers convey a threatening tone of voice. In other words, these data focus on how lay-people conceptualise a threatening voice, rather than selecting data which appears to be threatening. Another advantage of these data is that they provide some insight into how lay-people might perceive threatening voices. In effect, speakers might be portraying a voice that they themselves would find threatening.

This study compared samples of each speaker's baseline, non-threatening read speech, to samples of indirectly and directly-phrased threatening speech. The results presented in Chapter 5.2 were presented in terms of phonetic changes between non-threatening and threatening speech data. The phonetic features selected for analysis were chosen as they had been associated with threats in previous relevant research. In addition, these features are also known to be salient and recognisable by lay-people with respect to speech perception (Banse and Scherer, 1996). As such, these are phonetic features which would reasonably be expected to be modified by speakers to convey a particular behaviour.

For measurements relating to mean fundamental frequency and articulation rate, there was little evidence to suggest a meaningful, statistical difference between non-threatening and threatening speech. There was also no clear tendency for the usage of particular vocal tract features or vocal settings during threatening speech. However, there did seem to be a statistically-significant difference across most of these data for changes in intensity (with the exception of the male actor group data). This was manifested as an increase in intensity during threatening speech production.

This finding suggests then, that there is little basis for a broad, so-called 'threatening tone of voice' that is used across speakers. What is possible, however, is that there are numerous strategies used to signal a threatening intent. As noted in §5.2.11, there were a number of speakers who increased both their mean fundamental frequency and articulation rate when making threats. Conversely, there were speakers who decreased both f0 and AR. What appears to be most commonly observed in this research, are speakers who produce variable (typically statistically non-significant) increases to f0, AR and intensity when making threats relative to their baseline speech.

Research question 1.a:
It was noted that the data described above generally resembled previous phonetic descriptions of angry speech. This finding suggests that these speakers might associate a 'threatening tone of voice' with anger or aggression. As mentioned previously, anger and aggression are often discussed in relation to the production of threatening language. In particular, people making threats are described as acting based on feelings of anger, or as trying to convey that they are angry. In order to more closely examine how these speakers' threats relate to angry speech, Chapter 5.3 presented the results of an equivalent phonetic analysis on angry speech data. The speech data described in Chapter 3.3 consisted of short samples of read angry speech (Task 4). These speech samples were productions of the sentence: 'After what happened, I'm going to lose it tomorrow'. This sentence was chosen as the words are emotionally neutral.

Overall, the speakers recorded for this study produced angry speech similarly to what has been described in previous relevant literature. As such, angry speech for these speakers was characterised as having increased mean f0, intensity and AR relative to corresponding baseline speech. However, these quantifiable phonetic features were mostly not significantly different between angry and baseline speech. It was less clear how the phonetic features of angry speech related to those of threatening speech. Across the speakers, the difference between quantifiable phonetic features of angry and threatening speech (indirect and direct) were also not statistically-significant.

These findings suggest that for these speakers, threatening, angry and baseline speech were produced similarly to one another. This finding might suggest that speakers recorded under experimental conditions might not deliver speech in a manner which represents how these different emotional and behavioural states would be conveyed in a real-world setting. It should be noted that this point also applies to data collected from actor participants. That is to say, that actors (in general) also appeared to deliver speech with similar phonetic properties across the experiment tasks. It is also possible that this study did not examine the phonetic or prosodic features which these speakers used to convey anger and threat. As suggested earlier in this research, speakers might be signalling threat and anger using finer-grained phonetic or prosodic cues. This might involve cues such as changes to intonation or the use of pauses.

It can also be said, however, that the data collected in this research for angry speech is not ideal for comparison with the 'threatening' speech data. The threatening speech data were based on handwritten messages authored by the speaker him/herself which were directed toward a specific (fictional) person. The participants were also given a backstory surrounding their threat which aimed to provide them with some level of motivation for why they were making threatening telephone calls. The angry speech in this research consisted of a sentence being read aloud in an 'angry way', but without any additional context. Therefore, it is possible that had the angry speech been collected under more similar circumstances to the threatening speech data, this would provide a more appropriate comparison.

Research question 1.b:

The following research question centred on the possible differences in vocal production between indirect and direct threats. Overall, these types of threat appeared to be produced using similar phonetic features to one another (as well as to baseline speech). This finding indicates that speakers do not opt for specific vocal cues to communicate an intention to harm in the absence of words which do not explicitly express a threatening intent. As such, this study does not provide evidence of a compensatory effect for indirect threats. That is to say, there is no evidence to suggest that indirect threats are produced with notable phonetic differences which might serve to communicate a threatening intention in place of words where this intention is more explicitly stated.

Research question 1.c:

Chapter 5.4 explored whether authentic and simulated threatening speech differ with respect to the phonetic features analysed for this study. To allow for this comparison, male non-actor data was taken from Task 5 of the experiment. For this task, the participants read aloud typed transcriptions of the authentic threat data. These participants did not hear the original recording, nor were they provided with any descriptions of these authentic cases. For both the authentic and simulated data, the mean f0 and articulation rate of threatening speech were presented. The results of the analyses suggested that the speakers of authentic data produced threats in a more varied manner. In other words, the simulated (experimental) data was more similar across the board compared to authentic data.

This inconsistency in the phonetic realisation of authentic threats can be attributed to the varied circumstances surrounding these data. As described in Chapter 3.2, forensic phonetic materials are affected by numerous external factors, these include (but are not limited to) substance usage, poor audio quality, lack of available speech content and emotionally-charged scenarios. The phonetic differences between authentic and simulated threat data indicates that the results extrapolated from experimental data might not be fully applicable to real-life, forensically-relevant scenarios. As such, the findings of this study should be interpreted with this difference in mind.

Research question 1.d:

The simulated data presented in Chapter 5.2 were produced by two groups of speakers, one whose members had reported acting experience and one whose members reported having had no formal acting experience. The majority of the 'actors' were undergraduate students who had previously studied drama at school and/or college and were active in drama or performance clubs. The phonetic productions of the actors often differed from those delivered by non-actors across the experimental tasks. This difference appeared to be particularly evident for the female participants. For the female non-actors, there appeared to be a tendency for speech to be articulated more quickly, and for the vocal tract features and vocal settings used during non-threatening speech to be similar those used during threatening speech. In addition, the female actors typically produced threatening speech with a decrease in mean f0. In contrast, the female non-actors produced threatening speech with an increased mean f0. It is possible that these changes in speech tempo, f0, and/or vocal setting might be due to the actors being comparatively more comfortable with producing more extreme vocal changes while under experimental conditions.

Besides these phonetic differences, there were also differences in how the actors and non-actors approached the experimental tasks. For example, a number of the actor participants wrote very little, or did not write down anything at all in response to Tasks 2 and 3. As such, for these actors their performances were mostly or entirely improvised. This behaviour was not observed for the non-actors in this experiment. Also, the actor participants appeared to be more comfortable with being recorded reading aloud than the non-actors. This might not be surprising given the experience of the actors, however, this impression does have implications for the collection of speech data from actors. It is possible that the speech recorded from non-actor participants might be influenced by nervousness or anxiety that is associated with reading aloud and performance. In turn, this might affect the vocal production of non-actors (in particular).

As discussed in Chapter 2.3.4, studies into the phonetics features of emotional speech often use actors to provide stimuli. The findings of this study suggest that actors and non-actors produce speech differently under experimental conditions, indicating that caution should be taken when extrapolating results taken from actor-produced speech to a non-actor population. The current study is motivated by how lay-people understand and conceptualise threatening language, not necessarily how actors are 'taught' to appear threatening. However, it is plausible that acted performances of threats could influence lay-person threat production. For example, lay-people producing pre-scripted threats might base their delivery on performances they have seen in films, television or audio media such as audio-books or radio drama. The participants of this study were not asked about any external factors (such as exposure to media performances) which might have influenced their threat productions.

Research question 1.e:

This study also considered the possibility that male and female speakers produce threatening speech differently in terms of phonetic cues. It is common knowledge that males undergo physiological changes during puberty that results in longer vocal tracts than females. This elongated vocal tract allows for lower fundamental frequencies to be generated. In turn, low f0 is associated with conveying threatening and dominating behaviours. As described in this section, there were no statistically-significant differences between the quantifiable phonetic features present in threatening and non-threatening speech. This lack of statistical significance between the phonetic features of threatening and non-threatening speech was also seen for the female data. Therefore, the current research was not able to provide an empirical basis to suggest that these speakers (irrespective of their sex) adopt a so-called ‘threatening tone of voice’.

These findings could be interpreted in a number of ways: firstly, that there is no basis for a consistently-used ‘threatening tone of voice’ used by male speakers which resembles the previous descriptions of threatening vocalisations. Similarly, female speakers (who do not undergo comparable physiological changes around puberty) also do not appear to use particular vocal features to convey a threatening intent. It is also plausible that making threats is a ‘learned’ behaviour perhaps through exposure to threatening behaviour (real or fictional). As such, the speakers recorded for this study might have developed their own strategies for appearing threatening that are based on their own individual experience. However, this study is unable to explore this possibility any further due to a lack of relevant data. The commission of studies into the acquisition of threatening behaviour samples would be able to better assess whether threats are developed in response to evolutionary or social factors.

Research question 2 This study aimed to not only document the phonetic realisation of threats but also their linguistic content; that is to say, the words which authors use when intending to communicate an intention to commit harm. Examining the linguistic features of the threatening materials collected for this research also allowed for comparisons with other related studies. This comparison provided a sense of whether these data were linguistically distinct from what is currently known about threatening language. The participants who took part in this study were instructed to write an ‘indirect’ and ‘direct’ threat. In addition, the participants were instructed to communicate this threat with the intention to harm the owner of a farm. Otherwise, the authors were given no further details or instructions about how to write a threatening message.

A number of linguistic features present in these simulated threats appeared to be consistent with previous descriptions of authentic threatening texts, including: the high usage of the second person pronoun ‘you’, the lack of words present which express doubt, and the expression of tense and affect. As described previously in this chapter, there were also several differences between these data and the linguistic features said to predict harm.

These predictors included: direct threats being longer than indirect (or veiled) threats, assertive statements and detailed descriptions surrounding the timing, location and nature of the harm being threatened.

It is possible that the linguistic differences between the simulated and authentic data can be attributed to the design of the experiment itself. For example, the lack of detail regarding the nature of the harm threatened could be related to the lack of any information provided in the backstory concerning specific weapons. However, it is also possible that this lack of detail across the threatening data is due to the threat itself being fictional. In other words, the participants have not planned a threat based on their own actual motivations or circumstances and as such, were unlikely to be able to express much in the way of detail in their simulated threats.

Research question 2.a:
Another strand to this linguistic analysis was whether the participants of this research wrote indirect and direct threats in a manner that was consistent with the previous linguistic descriptions of different types of threatening language (as described in Chapter 2.2.5. It appeared that the direct threats produced by these participants were characterised by overt references to violence, as opposed to specific semantic or grammatical features. However, there were a number of authors who did not include any clear references to harm. It is possible that for these authors, indirect and direct threats were distinguished on separate grounds. Alternatively, these participants might have felt uncomfortable with expressing violent acts in an experimental setting. These overall findings, however, suggest that there is evidence that lay-people conceptualise indirect and direct threats in a manner similar to previous linguistic research.

Research question 3: Finally, this research examined the prosodic realisation of text which had been modified by typographical features (Chapters 5.5-5.10). When preparing the Task 2 and 3 threatening scripts for computational analysis, it was observed that some of the scripts contained typographical features. These features included underlined text and text enclosed in quotation marks. It was considered that the inclusion of such typographical features might influence the subsequent vocal production of the threats. The utterances which corresponded to the use of typographical features were extracted from their respective sound files. A series of prosodic analyses were performed on these extracted data. It was found that the speakers of these data varied in terms of whether there was acoustic evidence to suggest that the use of typographical features had influenced their vocal productions. While some speakers appeared to vocally stress the word or phrase which was typographically modified, others seemed to not vocalise any emphasis on modified text.

This finding suggested that firstly, authors of scripted speech do not necessarily include cues to the vocal production of the text. There were no instructions or directions present in these texts which explicitly-stated how the text ought to be produced (e.g. 'say this word aggressively'). However, it is possible that this related to the instructions of the experiment being to produce the threats in a 'threatening way.' Also, this prosodic analysis suggested that speakers delivered speech in a different manner despite using similar or identical typographic features. This raises the possibility that these speakers interpreted the vocal production of typographical features differently. For example, underlining text might be associated across authors with varying prosodic cues. It is also possible that for some of these authors, including typographical features in their texts was not meant to inform their subsequent vocal production. However, as the authors were not questioned on their use of typographic features, it is unclear why some of theses authors included such features in their texts.

The findings of Chapters 5.5-5.10 also have implications for other forensically-relevant texts, such as the live reading of written witness testimony, or the transcription of spoken testimony in a written form. If a speaker were to read aloud a transcript in court it is reasonable to assume that the inclusion of typographical features could inform his/her subsequent phonetic production. In turn, this phonetic production might convey different information to the court which might be absent had the transcript not contained typographical features. For example, by placing prosodic stress on a particular name, this information might be perceived to be especially meaningful or important by jurors. It is not yet known how modifying transcriptions might affect jury decision making in court. What is clearer from this study is that lay-people appear to produce different prosodic cues when reading text with the same or similar typographical features.

6.2 Limitations of the current study

There are also a number of constraints to the current research which have been identified throughout this document. This section will address these limitations in more detail. Firstly, the authentic and simulated threatening speech data collected for this study centred on very different circumstances. These ranged from spontaneous face-to-face interactions between people who knew one another, to pre-scripted threats directed toward a target not personally close to the threatener. It is plausible that threats could be produced (both phonetically and linguistically) differently toward different targets or in response to different circumstances. The data presented in the current study, largely focuses on fictional threats motivated by allegations of animal cruelty. While some of the linguistic features in these data resemble previous reports of authentic (and more varied) threats, it is possible that the findings of this study have limited applications to other threatening situations.

In addition, the simulated data could be said to represent 'bluffs' rather than so-called 'genuine threats'. In other words, because these data were in response to a fictional scenario, there was no possibility that the actions stated in the 'threats' could be realised. In some sense, it could be said that these data are in some way deceptive or untruthful. However, this study was centred on the production of a 'threatening intent'. What was important for this study, was that the participants were attempting to communicate an *intention* to commit harm, not that they would actually perform a harmful action.

Another limitation of this research is that there were no examples of non-threatening texts. As such, no comparisons could be made between the linguistic features of threatening and non-threatening texts. Therefore, the results presented in this research were likened to observations made in previous related research. Much of the previous research into threatening language or threat detection also does not compare threatening and non-threatening messages. This might be because of the difficulties of establishing what written data could serve as an appropriate point of comparison. For example, comparisons between a 'threatening' tweet and an academic essay would present linguistic differences which would be evident of their differing styles, rather than the expression of threat.

This study was also constrained by a relatively small amount of linguistic data in comparison to previous more quantitative studies on threatening language. These data, however, were collected using an experimental design based on previous research into emotional speech production. Compared to studies which examine the phonetic properties of threatening speech, this study analysed a large amount of speech data. As such, while there are a substantial amount of threatening speech data examined in this research, the quantity of written data is very small. As such, the output of the computational approaches presented in this research must be interpreted with this limitation in mind. Future research which also takes an interdisciplinary approach to data analysis should consider methods which can address different data sets.

6.3 Directions for future research

This study has also identified a number of directions for future research relating to threatening language. In this research, indirect and direct threats appeared to differ on the basis of how concepts relating to tense, affect and power were expressed. This could be examined more thoroughly using already existing threatening language corpora. However, as alluded to in Chapter 2.2.6, the data in such corpora appear to be directly-worded. Therefore, it might be necessary to collect larger quantities of indirect and direct threats in a similar manner to that which is described in this study.

This study also focused on describing the broader phonetic characteristics of threatening speech (such as the mean fundamental frequency across a threat). Future research could explore the finer-grained detail of these phonetic features. For example, by looking at the overall direction of f0 across a threat, and/or within the utterances which comprise a threat. In addition, there is also the possibility that this research did not identify (and as such did not address) the factors which do influence threat production. This research explored the relationship between speaker sex and their level of reported acting experience on threat production. However, future research should consider whether other factors might also effect threatening speech. For example, whether speakers and/or authors change their threats based on their target, or the extent to which threatening language interacts with changes to the threatener's emotional state.

There are also a number of topics alluded to in this research which were not the focus of the current study. For example, it is not known how children 'acquire' the phonetic and/or linguistic features of threatening language, or whether different languages or cultures express and interpret threatening language differently. Many of the relevant theories into threat production described in this study appear to treat threats as a universal behaviour. It is also not clear how exposure to threatening language (such as depictions in media or as a result of abusive behaviour) might influence threat production.

This area of research also lends itself to an interdisciplinary approach involving, but not limited to, the domains of forensic linguistics and forensic phonetics. This research has largely focused on the expression of threats which were not produced in real time towards a target. However, many threatening scenarios involve face-to-face interactions and as such, there is untapped potential for research which examines this area in more detail. For example, threats could be examined through the lens of conversational or interactional approaches, or by psychologists with an interest in the communication of emotions or behaviours through multi-modal channels.

In addition, there a wide range of possible applications to the findings of this research outside of academia. Most obviously perhaps, this research could inform the design of teaching or training materials which advise lay-people on threatening speech. On the basis of this research, there is evidence to suggest that individual speakers vary their phonetic productions of threats. As such, training resources could emphasise how threats might not be easily identifiable based solely on vocal cues. An approach which makes lay-people aware of the various strategies which people use to appear threatening, seems to be appropriate given the current (publicly-available) research into threat production. For example, the speakers in this study produce spoken threats with inconsistent changes to f0, articulation rate, and voice quality. The written threats collected for this research also showed that threats can (albeit rarely) be expressed as explicit warnings or promises. The plausibility of these strategies could be further tested in perception research, where participants judge 'threatening' language where these types of phonetic and lexical cues are manipulated. There are also ample opportunities for the application of this research in non-forensically relevant contexts, such as advising managerial staff on how to appear less threatening to their colleagues via speech or email communications, or to vulnerable people.

To conclude, it would appear that the production of threatening language is considerably more complex and variable than might be suggested by its treatment as legal testimony or as a predictor of violence. This research has cast doubt on the assertion of a 'threatening tone of voice'. Therefore, caution should be advised when describing a 'threatening voice' as a useful tool for predicting violent or aggressive behaviour. The linguistic analysis of this research could be taken into consideration for the development of threat detection algorithms or for training materials relating to threat assessment. It is hoped that this research serves to encourage future interdisciplinary approaches into the research of complex language crimes such as threatening language.

7 Appendix

This Appendix consists of the following items:

- Hardware and software used
- Authentic threatening data scripts
- Vocal Profile Analysis Scheme
- Praat scripts
- Ethics forms (for simulated data collection)
- Experiment instructions (for simulated data collection)
- Tests of Within-Subjects Effects (the output of an SPSS repeated-measures ANOVA for the fundamental frequency data as shown on page 161)
- Simulated threatening data scripts (With the exception of Speakers 24, 25 and 30)

7.1 Hardware

2.7GHz MacBook Pro laptop (2015) with Intel HD Audio chipset

Sennheiser HD 280 Pro headphones

Sennheiser PC 8 USB Internet Telephony On-Ear Headset

7.2 Software

Praat (Version 5.4.04)

Microsoft Word (Version 16.17)

Microsoft Excel (Version 16.17)

LIWC2015 software

Tropes software (Version 8.4)

7.3 Authentic utterances analysed for this study

Table 24: Authentic utterances analysed for this study

Speaker	Utterances
1	I'll threaten you, I'll put you in a fucking rose garden you cunt. You understand that? Because I'm capable of it. You understand that?
2	(Got your phone) you fucking arsehole, all you degrading me all round when when you got all your mates round you, and you're giving the you're giving phone calls that can't be traced. You can be traced and I will fucking trace you, and I'll tell you what you will find out if I've got fans or not (I'll) be bouncing up and down on your fucking throat you fucking arsehole. Don't ever phone me again. If you do, I'm going to kill you.
3	You know I would absolutely kill you if you ever did something like that.
4	I would beat your motherfucking ass if I ever met you.
5	Dana White. This is Dan Quinn. You messed with the wrong guy here sending your English tough guy and your Indian crazy Privesh to threaten my life and to make death threats against me and my kids. So, I just want you to know that since you're bigging up to me, you should be worried about your own safety, bitch. And you can be looking over your fucking shoulder and checking your rear view mirror in your fucking six Ferrari's, because I'm coming for you, bitch and Im fucking well packed. I'm going to shoot you right in the fucking head motherfucker. Fucking threatening my life. How about I fucking take you and the whole fucking Zuffa Crew down with you bitch. Trying to block my Stevia movement, making death threats against me. You want a piece of me? How about you fucking man up. How about you be a fucking proper man, rather than a slave master and step in the fucking (cage see), bitch. I'll drop the wind right out of you fool. You got beef? You got the fucking balls to be in the fucking (unintelligible phrase) with me? How about you give me a ring, bitch? I'll fucking do you face to face. You know the name: Dan Quinn. Waiting for your call, bitch. You're dead, bitch.

6	Hey (name removed), this is Terrence, and you've been calling my wife. If you call my wife again, I'm going to come to your house and I'm going to cut your fucking throat. Understand now. I'm going to tell you this one time. You call my wife again, I'm going to kill you.
7	So, you wanna try to bring me down? I will fucking kill you and your whole fucking family, alright? So go ahead and fuck with me, alright? Try it.
8	Don't stand there and make a fucking fuss, or else I'll twat you right there now.
9	I'm gonna go down and shoot him.
10	What it is I'm hearing voices and I'm not very well at all I'm going to kill somebody soon.

7.4 Vocal profile analysis scheme

Appendix 2: VOCAL PROFILE ANALYSIS PROTOCOL

	FIRST PASS		SECOND PASS			
				Not.	Mark.	Extr.
	Neutral	Non-Neutral	SETTING	1	2	3
A. VOCAL TRACT FEATURES						
1. Labial			Lip rounding/protrusion			
			Lip spreading			
			Labiodentalization			
			Extensive range			
			Minimised range			
2. Mandibular			Close jaw			
			Open jaw			
			Extensive range			
			Minimised range			
3. Lingual tip/blade			Advanced tip/blade			
			Retracted tip/blade			
			Sibilance			
4.Lingual body			Fronted tongue body			
			Backed tongue body			
			Extensive range			
			Minimised range			
5. Pharyngeal			Pharyngeal constriction			
			Pharyngeal expansion			
6. Velo-pharyngeal			Audible nasal escape			
			Nasal			
			Nasal w/o full release			
			Denasal			
8. Larynx height			Raised larynx			
			Lowered larynx			
B. OVERALL MUSCULAR TENSION						
10. Vocal tract tension			Tense vocal tract			
			Lax vocal tract			
11. Laryngeal tension			Tense larynx			
			Lax larynx			

C. PHONATION FEATURES

		Present		Scalar Degree		
				Not.	Mark.	Extr.
	SETTING	Neutral	Non-Neutral	1	2	3
12. Voicing type	Voice					
	Falsetto					
	Creak					
	Creaky					
13. Laryngeal frication	Whisper					
	Whispery					
	Breathy					
	Murmur					
14. Laryngeal irregularity	Harsh					
	Tremor					

7.5 Praat scripts

\AM@currentdocname .pdf

.pdf

```
form Compute intensity stats of audio files
    comment Directory of sound files
    text sound_directory Full path of audio directory
    sentence Sound_extension .wav
    comment Full path of text file with results
    text resultfile
/Users/SK/Desktop/intensityresults_test.csv
    comment Intensity analysis parameters
    positive Time_step 0.01
    positive minimum_pitch_(Hz) 100
    positive maximum_pitch_(Hz) 200
endform

# Create debug window
writeInfoLine: "Debug"

# Create listing of all sound files
Create Strings as file list... list
'sound_directory$'*'sound_extension$'
numberFiles = Get number of strings

# Check if the result file already exists
if fileReadable (resultfile$)
    pause The resultfile 'resultfile$' already exists!
Overwrite?
    filedelete 'resultfile$'
endif


# Create row with column titles to the result file
titleline$ =
"Filename','Min','Max','Mean','Median','Std','Dur','q25','q75'
,'newline$'"
fileappend "'resultfile$'" 'titleline$'

# Compute all the sound files
for ifile to numberFiles

    # Open audio file
    appendInfoLine: "---------------------------"
    appendInfoLine: "file number: ", ifile
    filename$ = Get string... ifile

    # Read file and retrieve details
    Read from file... 'sound_directory$''filename$'
    appendInfoLine: "file read"
    soundname$ = selected$ ("Sound", 1)

    dur = Get total duration
```

7.6 Ethics forms (for simulated data collection)

UNIVERSITY of York

DEPARTMENT OF
LANGUAGE AND
LINGUISTIC SCIENCE
Heslington, York, YO10 5DD, UK
Email sk720@york.ac.uk

INFORMATION SHEET

PLEASE KEEP THIS INFORMATION SHEET AND A SIGNED COPY OF THE CONSENT FORM FOR YOUR RECORDS

You are invited to take part in a research study. Before you decide whether to participate it is important for you to understand why the research is being done and what it will involve. Please take the time to read the following information carefully. If there is anything you do not understand, or if you want more information, please ask the researcher.

Title of study: Collection of threatening speech corpus

Researcher: Sarah Kelly

What is the research about?

This study aims to collect threatening speech data, alongside non-threatening speech data for use in comparative phonetic analysis and future perception experiments.

Who is carrying out the research?

The research is being carried out by a postgraduate researcher based at the Department of Language and Linguistic Science at the University of York.

Who can participate?

This study requires participants who are either:
- native speakers of English who are between the ages of 18-30.

What does the study involve?

The study will involve recording your voice as you perform different reading tasks as instructed by the researcher. The reading tasks will involve reading aloud pre-prepared sentences and 'threats', as well as simulating your own 'threats' in response to a scripted scenario. This will be immediately followed by a short online questionnaire about threatening speech, as well as your acting experience (if applicable).

The study will take place in a room located in the Language and Linguistic Science department (at the University of York), to be agreed upon by the researcher and participant.

The data will be collection in a single session, which is expected to take around an hour.

Consent form

This form is for you to state whether or not you agree to take part in the study. Please read and answer every question. If there is anything you do not understand, or if you want more information, please ask the researcher.

Have you read and understood the information leaflet about the study?	Yes ❐ No ❐
Have you had an opportunity to ask questions about the study and have these been answered satisfactorily?	Yes ❐ No ❐
Do you understand that the information you provide will be held in confidence by the research team, and your name or identifying information about you will not be mentioned in any publication?	Yes ❐ No ❐
Do you understand that you may withdraw from the study at any time before the end of the data collection session without giving any reason, and that in such a case all your data will be destroyed?	Yes ❐ No ❐
Do you understand that the information you provide may be kept after the duration of the current project, to be used in future research on language?	Yes ❐ No ❐
Do you agree to take part in the study?	Yes ❐ No ❐
Do you agree to excerpts from your audio/video recordings to be used in presentations or in teaching by the researcher, without disclosing your real name? *(You may take part in the study without agreeing to this).*	Yes ❐ No ❐

Your name (in BLOCK letters):

____

Your signature:

____

Researcher's name:

____

Date:

____

One copy to be retained by the researcher, one copy to be kept by the participant.

7.7 Experiment instructions (for simulated data collection)

Simulated data instructions (Task 1)

The North Wind and the Sun were disputing which was the stronger, when a traveler came along wrapped in a warm cloak.

They agreed that the one who first succeeded in making the traveler take his cloak off should be considered stronger than the other.

Then the North Wind blew as hard as he could, but the more he blew the more closely did the traveler fold his cloak around him;
and at last the North Wind gave up the attempt.

Then the Sun shined out warmly, and immediately the traveler took off his cloak. And so, the North Wind was obliged to confess that the Sun was the stronger of the two.

The North Wind boasted of great strength. The Sun argued that there was great power in gentleness.

"We shall have a contest," said the Sun.

Far below, a man traveled a winding road. He was wearing a warm winter coat.

"As a test of strength," said the Sun, "Let us see which of us can take the coat off of that man."

"It will be quite simple for me to force him to remove his coat," bragged the Wind.

The Wind blew so hard, the birds clung to the trees. The world was filled with dust and leaves. But the harder the wind blew down the road, the tighter the shivering man clung to his coat.

Then, the Sun came out from behind a cloud.

Sun warmed the air and the frosty ground. The man on the road unbuttoned his coat.

The sun grew slowly brighter and brighter.

Soon the man felt so hot, he took off his coat and sat down in a shady spot.

"How did you do that?" said the Wind.

"It was easy," said the Sun, "I lit the day. Through gentleness I got my way."

Simulated data instructions (Task 2 and 3)

For the next task, you will read a short story about threats made towards a fictional business.

Afterwards, you will be asked to construct a threat related to this story, and read it aloud.

You may write down what you wish to say using the paper provided.

You may also revise your 'threat' at any time during the experiment.

Backstory

Tom West set up Wests Farm in Gloucester to breed animals for medical research. Earlier this year, the mainstream media published covert video footage taken at the farm. This footage revealed that the animals were being kept in unsanitary conditions.

Subsequently, a campaign group called 'Shut Down West's Farm' was established. This group disagreed with the use of animals for medical research, as well as the conditions in which the animals are kept.

'Shut Down West's Farm' published Tom West's mobile telephone number and home address on their website. This website urges supporters of their campaign to contact Mr. West in order to convince him to close down the farm. Whilst you are not a member of this campaign group, you agree with their cause.

Despite campaigners asking Tom West to close down the farm, Mr. West insists that he will continue his business and not bow down to pressure.

You decide that simply asking Tom West to shut down the farm is not effective enough.

You decide to send Tom West a series of threatening telephone messages.

Your first message is an indirect threat toward Tom West. Your message should communicate an intention to harm him.

(Please do not turn over the page until instructed to do so)

This message proves to be ineffective. Mr. West contacts you to tell you that he has no intention of closing the farm.

Your second message is a direct threat toward Tom West. Your message should communicate an intention to harm him.

Read aloud both of your messages in a threatening tone of voice

7.8 Tests of Within-Subjects Effects

Tests of Within-Subjects Effects

Measure: f0

Source		Type III Sum	df
condition	Sphericity Assumed	191.66	2
	Greenhouse-Geisser	191.66	1.422
	Huynh-Feldt	191.66	1.583
	Lower-bound	191.66	1
condition * Sex	Sphericity Assumed	96.56	2
	Greenhouse-Geisser	96.56	1.422
	Huynh-Feldt	96.56	1.583
	Lower-bound	96.56	1
condition * Acting	Sphericity Assumed	212.337	2
	Greenhouse-Geisser	212.337	1.422
	Huynh-Feldt	212.337	1.583
	Lower-bound	212.337	1
condition * Sex * Acting	Sphericity Assumed	85.639	2
	Greenhouse-Geisser	85.639	1.422
	Huynh-Feldt	85.639	1.583
	Lower-bound	85.639	1
Error(condition)	Sphericity Assumed	6713.802	74
	Greenhouse-Geisser	6713.802	52.62
	Huynh-Feldt	6713.802	58.56
	Lower-bound	6713.802	37

Mean Square	F	Sig.	Partial Eta Squared
95.83	1.056	0.353	0.028
134.768	1.056	0.334	0.028
121.1	1.056	0.34	0.028
191.66	1.056	0.311	0.028
48.28	0.532	0.59	0.014
67.898	0.532	0.53	0.014
61.011	0.532	0.549	0.014
96.56	0.532	0.47	0.014
106.168	1.17	0.316	0.031
149.308	1.17	0.304	0.031
134.164	1.17	0.308	0.031
212.337	1.17	0.286	0.031
42.819	0.472	0.626	0.013
60.218	0.472	0.561	0.013
54.111	0.472	0.581	0.013
85.639	0.472	0.496	0.013
90.727			
127.592			
114.651			
181.454			

7.9 Simulated threatening data scripts

Tom West, You have been violating Animal right's for too long. ~~Shut down immediately~~

~~If Wests farm is not closed, there will be consequences.~~

If You don't shut down West's farm immediately, there will be consequences

Tom West, ~~You test my patience and my temper~~. Have ignored my warning.

~~Close~~ Stop what you are doing Or ~~I will hunt you down and~~
~~I will Blow your hell hole farm to kingdom King Come~~

I will hunt your sarry ass down and Stick you like one of your own fing pigs.

• Hi Mr West,
I think you should know I strongly dissagree with your practises and think you urgently should change them, especially considering the negative press ~~around~~ you are currently facing from the growing campaigne group. ~~threat~~ Which are a threat to everything you care about. Just change now before it's too late.

• Mr west,
Stop this now or I will be taking this to court. And while I'm at it you watch every thing you hold dear dissapear. Your ~~house~~ busness will burn and your family will stare down the barrel of a gun ~~until~~ ~~you~~ ~~stop~~ ~~close~~ ~~your~~ ~~business~~ ~~down~~ ~~and~~ unless you stop in the next 2 days. ~~Kind~~ ~~you~~ You are not safe.
~~Kind~~

(1) Mr. West, it seems you are incapable of taking upon yourself the responsibility of closing down your diabolical "farm". If you don't, someone else will, and I think you'd rather it didn't come to that

(2) Perhaps I didn't get through to you last time. I'll show you a level of cruelty you must be quite used to now, you fucking bastard. I'll shoot your fucking guts out ~~with~~ and let your animals feed on your ~~[illegible]~~ fucking corpse.

150

I'm calling ~~to bring~~ to politely suggest that Mr. Tom West of Wests Farm ~~[illegible]~~ consider the complaints of groups such as 'Shut Down Wests Farm' and other reasonable and compassionate people and begin the ~~release of ceasing~~ immediate ceasing of all its ~~medical~~ medical trials ~~on animals~~ and for the safe release of all non-human animals contained on site. While I understand that Mr. West is a busy man, and has also ~~rejected~~ ~~declined~~ refused to respond to similar complaints in the past, I would like to remind Mr. West that I and other like-minded, compassionate individuals have ~~his~~ more than just his phone number, and if necessary would be ready to move our protest ~~closer~~ ~~[illegible]~~ to his home, and his family. ~~Have a nice~~ I'm waiting.

2nd

S4

I don't think I was clear so
let me be fucking clear.
~~we will~~ You are going to ~~post~~
give us a guarantee of ~~closure~~
Wests Farms closure or ~~we~~ ~~are~~ I am
going to ~~come to your house,~~
~~your car~~ kill you. ~~we~~ I know
where you live, ~~we~~ ~~we~~ I know
where you work, ~~we~~ I know where
~~the~~ your victims are and ~~we~~ I have
the tools to get them. The only
thing that has kept anyone from
putting a bomb in your car and
raiding your animal Aushwitz
is the common decency and respect
for life ~~that~~ you obivously lack
and me and my friends would
be more than willing to be the
first to have the balls to
show you what it's like to
be one of your non-human victims.

I'm waiting.

Mr. West I know where you live we all do, so if I were you I'd shut down the farm as soon as possible

Mr. West, If you don't shut down the farm I'll treat you like you treat your animals, [illegible] Y.

I've got guns, I'm not afraid to use

Hello Mr West. I'm struggling to see how you could operate such a farm. It is unethical beyond means and to think that you won't take into account what others think is unforgiveable. I hope you make changes soon or we will have to take further action.

So it seems you don't care about what the rest of us think. You don't want to make too many enemies Mr. West. Worse comes to worse we'll just take matters into our own hands. Be careful.

What you're doing to those poor animals is ridiculous. We have all had enough of it and sooner or later you're gonna get what's coming.

~~Why are you~~

That's it. ~~We're~~ We are going to ~~[illegible]~~ give you a taste of your own medicine and get back at you for all those animals you've hurt. You and your farm are gonna burn and you will ~~are~~ ~~get~~ feel 10x the pain all those poor creatures have felt.

Message 1

Listen mate I've just seen that footage and all I can say is you are fucking disgusting ~~mate~~. I cant believe you ~~[illegible]~~ think you have the right to treat ~~[illegible]~~ living creatures ~~that~~ like that you twat. Who do you think...

To be fucking honest it should really be you fucking crawling around in that shithole you pig.

I'm fucking livid. You better watch your back mate because you're fucking got it coming for you. If you keep up what you doing there, ~~[illegible]~~ I know that crazier cunts than me will be knocking at your family's door very soon.

Message 2

You're got to be fucking joking me. You are one brave cunt. So your not going to shut down your farm? Then Ill fucking do it for you, starting with its twat owner.

~~And~~ Ill make sure you never feel safe again mate because every night youll be fucking thinking about me blowing your fucking home up

Protect yourself and your family while you can mate because those abused bloody animals have more use on this earth than you cunts.

If you decide not to shut I promise youll live in constant fear of everything you own going to rubble, wherever you hide.

Dear Tom West.

I Find it deplorable that you seek to make money out of Sexually exploiting animals. You are the lowest ~~of~~ Scum, ~~Scum~~ how is what you doing any different to Hitler?

You are Nazi ~~Scum~~ White privaledged scum. Use the earths resoucres for your own gred. and what do you contribute in return? I Bet you treat women as objects as well you prick

The world wouldnt miss you, you are the definition of an oxygen thief.

The sooner you die the better.

Do you think its okay to sexually exploit animals and keep them locked up?

How about I come over here with my mates and ~~t~~ take turns with your wife? Make you watch?

Then we cut her open just like the science labs do to animals. Of course we will ~~to~~ humanely put her to sleep afterwards, by slitting her throat!

Dont worry, we wont forget about you, you animals baking cunt.

I'll gouge your eyes out and leave you in an animal cage, I hope sepsis sets in before you are found.

What you're doing is disgusting. How do you live with yourself treating animals in the way you do? You ought to be ashamed of yourself, profiting from the mistreatment of those poor animals. Find a better way to make money you fuck or else.

You stupid fucking man. Say goodbye to your farm because it's about to get blown up and you're going with it. You'd better sleep with one eye open or you might not wake up the next morning!

'Shut down your farm or you'll regret the day you opened it!'

'If you continue to run your farm, I will plant bombs ~~in your~~ on your land and kill all of your livestock, ~~without consideration~~ * of whether or not your family are in the vicinity

* regardless

1st

Mr West, I am ~~going~~ calling to ~~seriously~~ warn you about the consequences to not shutting down the farm. The animals are suffering and there are many whom are furious. There are protests being held ~~in [illegible]~~ and if you do not follow my instructions, you will regret ~~farming~~ choosing to ~~open a~~ opening.

2nd

~~B~~. Mr. West, ~~I've no~~ you're had your last warning
I had told you before, there will be consequences.
I'm going to ~~f~~ find you and fucking kill you.
I will. And if I don't,
fucking many animal lovers
You better run, or close the farm - fast. will find you
You have 24 hours.

I know where you live and your personal phone number. If we could find that, consider what else we could easily find out about you and publisize. Shut down your farm or we will take action.

I warned you but you chose to ignore it, this is now on you. Ive made bombs, not just one but multiple and I won't just blow up you and your farm but also everyone you care about. You have 2 days to close down the farm - or I will. either way the animals will be put out of their misery.

Mr West, the conditions you are keeping your animals in are unsanitary and unsafe, and it's cruel to continue keeping them in this way.

If you do not close down the farm we will immediately have to take legal action defending the rights of the animals and get the police involved in ~~ma~~ ~~ou~~ moving these animals to a safer place and shutting down your farm for good

Mr West, ~~we cannot~~ ~~can't~~ ~~continue to let your farm stay open with the animals in such cruel conditions~~ you've ignored our previous messages, so now is your last chance. If you don't shut down your farm, we'll ~~take the animals and~~ blow it up. We have bombs and ~~guns~~ guns ready. We know where you live too, so unless you shut down the farm, expect us to come to you too.

~~You need to close your farm down. What you're doing is disgusting. How would you like to be treated like you're treating those poor animals. Close it down or you'll find out what it feels like.~~

① People are getting angry about this Mr West. Close the farm before the campaign gets any worse.

② Close your disgusting farm down ~~before~~ or we'll be round to close it for you. We'll come with guns and explosives and show you what it feels like to be treated like your animals.

F12

~~If you don't~~

If this farm isn't shut down before I get there, I'm going to make you regret it. Forget lawyers, or legal processes. I will come ~~to you and~~ and find you. I will ~~make you live like an an~~ treat you like an animal. ~~Pregnant~~ You will experience this life for yourself. ~~If~~ And you manage to survive it then I can assure you it won't be our last encounter.

Shooting, bombs, explosions. That's all to easy. You Death should be slow; ~~like~~ I want to see every last drip of blood leave your body. Painful, I will make sure your last words are you screaming for help. Lonely, your last sight will be me showing you, Mr. West, as a murder ~~Animal to~~ nothing but hatred in my eyes. You will be given the death you deserve. ~~death~~

① Hello Mr West,

~~W~~ I know about your filthy deplorable animal project. We've tried convincing you nicely to shut it, but that doesn't seem to be working... so we might have to convince you nastily instead. I suggest ~~How about~~ you have another think about how much you care about your farm, so that we don't have to push this any further than necessary. Is that clear? I hope so.

② Right Mr West.

Obviously you want to do this the ~~grim~~ hard way. We're not going to let you get away with what you're doing, so we'll just have to resort to something a little more violent. So get ready to see ~~some~~ your revolting farm blitzed to pieces. And as for you, ~~my~~ you disgustingly unethical swindler, ~~perhaps~~ you can expect to feel the pain very keenly. We don't have qualms about making you squeal worse than those poor animals you've been abusing. Last chance, and we mean it.

threat 1)

we know where you are and ~~were~~ we're not happy, if you dont stop we might have to pay you a visit

threat 2)

if you don't stop, we'll storm the farm with guns, and won't be stopping for anyone in the way.

521

Beast. despicable man vile vulgar

disgusting business
end the business
don't want businesses like yours abusing animals.
not going to let you get away with it.
Criminals → get you.

We've got bombs. We've got contacts. We'll ~~be there~~ coming to you. And you're going to feel the pain. the pain that you cause your animals to feel. Right now we're working our way into your factory to blow it up. When you're inside + least expecting your filth. And you're going to burn down along with your vile factory.

S22

SDWF
I've seen footage
lots of angry ppl
others want to harm
unethical , police ?

gonna join others
protest at farm
~~A~~ how would you like it

S23

I am absolutely ~~disf~~ disgusted by you, Mr. West. If I were you, I'd think very carefully about what you plan to do ~~next~~ from on your Change things immediately or face the consequences. You sicken me! Don't you recognise how much your animals are ~~suffurent~~ suffering? ~~I~~ You better change things, or ~~they~~ things will be changed for you. Mr. West.

I am going to ~~hurt~~ find you, you piece of filth. I tried to be civil, but looks like I'll have to come down there to get things changed. I'll punch the life out of you too help you change your mind about things. ~~It~~ I'll find you and I'll show how it feels to suffer like those animals are.

Good afternoon, Mr. West. I am calling to let you know that your awful treatment of the animals on your farm has not gone unnoticed by the general public, and we shall not rest until we have justice. Rest assured that if you do not discontinue this disgusting practice, it will no longer be your livestock's screams that are heard for miles around. Good day.

Mr West I tried asking nicely. It's you who deserves kept in a cage and unfortunately I don't own one. What I do have however is a very angry brother with a shedload of rather "explosive materials" I am not one man, Mr West and I swear if you don't stand down, we will blow you and your filthy little farm apart.

S27

Mr West, why don't you fucking listen man? We not getting through to your thick head or something? You need to fucking cut that shit out now before we come to your farm and beat some sense into you. When we're done, ~~[illegible]~~ you're not gonna want to even touch a animal again you piece of shit. We won't place nice man, I'm fucking military trained, ~~I~~ and I'm going to break your shitfaced jaw wide open with my sledgehammer if that's what it'll fucking ~~[illegible]~~ take

West → starting to love my fucking move with you

↳ you're a piece of shit for what you did

↳ arrogant selfish bastard who doesn't know how to listen

↳ those animals had no chance, but you do to make this right

↳ is not, we're not going to have a choice but to hurt your shit until you realise what you're doing is wrong

↳ we're not sitting by watching any longer

- Mr West
  - unsustainable
  - awful
  - arc of progress bends towards justice
  - it's not going to be good for him in Future
  - there are people that are coming for you

- Bomb your place
- Set fire to it allow you to watch your family burn
- Skin you alive
- Take you and do experiments on you - you ever had shampoo in your eyes? We're going to pour it down your mouth.
- Death is too fucking good

S29

~~How dare you, you slimy~~

Pick up the phone, you coward. You're not going to get away with this, I'll tell you that much. How dare you? Who the fuck do you think you are anyway, thinking you can play God like that. We won't stop calling, West. Not until we decide we need approach this with other means, anyway. That farm is getting closed down, one way or another, you can trust me on that. Cunt.

Right you Chuck, we've had enough.
And you know what we've decided?
Well, we've decided what'd be best for you would be to get a little taste of your own medicine. So, seeing as you don't seem to want to leave your such little torture farm, we're just going to let ourselves in one night this week. And when we're in, we're going to find you, and then we're going to cut you up, just like you did to those poor, helpless animals.
Then you'll understand what it's like.
I don't doubt you'll plead forgiveness, and it'll be admirable if you genuinely mean it rather than if you just end up saying it out of fear. Either way though, it won't matter. The guinea pigs couldn't escape and neither will you. We've had enough.

I think that what you are doing is not only wrong but sick. It is disgusting that you would breed animals, use them for experiments and then let them die in unsanitary conditions

I suggest that you shut down your farm immediately because everyone ~~a~~ now knows what you are doing and are ~~so~~ livid about the situation

How would you like to be treated like them animals? Everyone has your number now and know where you live, perhaps you need a taste of your own medicine, you're sick, twisted and deserve ~~a~~ quick karma. How about everyone in the town locks you up in a cage and start testing on you? i'm pretty sure that now everyone knows it will only be a matter of minutes before you're in the same danger as ~~you~~ the animals you abuse.

You deserve everything that you're about to get, now that the videos have gone viral you ~~can't~~ can no longer hide behind the police, everyone is going to get you in one way or another. Your life is not going to be the same now, want to go to the shop for some milk? you're going to get abuse getting the bus to town? i'd ~~make~~ wear protective clothing because now everyone knows what you have done to those poor animals, it is only a matter of time until you get what you deserve.

so, you have decided to ignore my last fucking message. so I'm going to tell you one more time and you better pay attention because I'm not messing around now!

close your fucking farm and free those animals you sick, twisted bastard. If you don't close it in the next 2 days I will create a bomb and blow up the entire farm including you, but before that I will find you and make you pay for your inhumane actions.

I've got knives and guns and I'm not afraid to use them ~~and~~ to prolong your suffering. You can't ~~[illegible]~~ hide behind "Mr. West" now, we all know your real name and where you are. I'm sure people are watching you right now, you're never going to be safe. Feel free to ignore my messages, but after 2 days if the farm still hasn't closed down you're not going to be able to ignore the attacks and torture that you're about to endure, by not only me but by the entire town.

I'd say watch your back, but it is probably too late for that now, we are onto you so a response to this message is simply not enough, take action or we will!

~~8~~532

what the fuck is wrong with you.
~~Mr~~ West, ~~your farm is morally wrong.~~ ~~E~~
~~How could you be!!~~
How can you call yourself a human?
Why is it ~~to~~ ~~the~~ allow these animals suffer but you a disgrace of a person don't? How the fuck do you sleep at night knowing ~~to~~ this.

West what the fuck is wrong with you? How do you have the audacity to call yourself human. You ~~tu~~ allow these animals to suffer whilst you yourself ~~you~~ live in comfort. I honestly don't know how the hell you fucking sleep at night. ~~I hope you stop breeding~~ Why don't you stop this you fucking bastard.

So I see that you are still ~~s~~ savagely breeding animals for your own profit. ~~How~~ would you like ~~If I~~ me to come to your house ~~a~~ with a gun and ~~stew~~ force you to mate for medical research. Better yet, why dont we just blow you up liberate the animals and leave you nothing like the fucker you are.

S33

Mr West,

As you are aware, there is a campaign group which is battling against you, which, in my opinion, is for all the right reasons.

Do you get pleasure out of seeing animals suffer? I'm sure you wouldn't like to be put under testing the way these animals are.

It's inhumane, cruel and unnecessary. There are various other ways to undergo medical research without testing on animals.

Many people are against what you are doing so for your own sake, do the right thing or there will be consequences.

S33

With regards to your message I would like to inform you that you cannot "chose" whether you close down the farm or not.

Closing down the farm is the ONLY option. If action is not taken from you soon, myself, the campaign, and everybody else who opposes your cruelty will use our weapons in order to end what you are doing, whether you have 'chosen' to or not.

S34

Mr West

We're all seen the footage of the monstrasities happen at your farm.

You should know that there are a lot of people who are very angry about this. and arent afraud to take mattey into their own hands.

If I were you I would seriously consider close your farm immediately

I think you should know.

I ~~have contacts~~ know some nasty people who cant provide me with lots of nasty chemicals with which I can make a bom.

If you don't close it, I will press a little button, destroy ~~your~~ your house, you'r car, everything

You have till the end of the week.

S35

Message 1

I am calling to ask you to shut down your farm. The animals you keep are being mistreated and it's disgusting. They should not be in your care. If you do not shut down the farm I will be forced to report you to the police and all local animal rights groups. ~~You can be assured that there will be protests~~

We know your home address and you can be assured that there will be protestors ~~protesting should~~ day and night if you refuse to comply.

Message 2

It seems you have refused to listen to my request that you shut down your farm. OK. Let's see what you'll say ~~think~~ when I arrive on ~~at~~ ~~yours at~~ your doorstep with my shotgun ~~in one hand ed~~ ~~~~ fully loaded. Or maybe I'll be more inconspicuous. I have access to bomb-making materials so you'd better sleep with one eye open from now on.

S36

You must shut down the farm. What you are doing to the animals is really unacceptable. How would you feel if you or your family were being treated like you treat the animals. You will seriously regret not closing down the farm and if you don't close it just wait and see what is in store for you.

You have just informed me that you won't be closing down the farm and you are really going to regret it. I will make your life a living hell and ~~then~~ treat you exactly how you treat the animals. I can kill you in a matter of seconds and bomb your house as well. If I were you I would reconsider your decision to not close down the farm.

537

Mr West. I have heard about your farming practices and I have just one question
What is the roles were reversed? I'm sure you, for instance, would hate to be trapped in a shed all day... no light... no food... tortured and cold... imagine it Mr West. Close your farm.

Mr West I am trying to be calm but I'm losing my temper. I have several tazer-based wepons and plan to make you see just what your animals are experiencing. You have 24hrs to shut down that fucking ~~farm~~ joke of ^ farm

S38

~~Kolla, You've ignored so many of us, you can't keep~~

(You can't ignore us for long. Just think about it, there are hundreds of us.) (~~You~~ are a sick human being and ~~soon every~~ everyone will know it.)

~~other have~~ we'll keep calling. We have your address. This will just keep getting bigger and bigger. ~~How will you feel when the ... ~~

~~Take a second to~~ why don't you think about the animals.

Their excruciating pain, living in their own filth, ~~Why don't you ta~~ Take a second to really imagine that. And that's exactly what you deserve

We're not stupid. We know you've heard us. ~~...~~ We told you you couldn't run away for long. We warned you. Now whatever happens ~~is your own fault~~ ~~isn't our fault, it's your~~ it's our duty.

We'll hunt you down. ~~Do you really think ... none of we're all nice? it's beyond us to hurt you like you hurt your precious animals?~~

~~It won't be hard to tear you apart, ten people on each extremity, sounds about right.~~

~~I can tell Do you really think what kinds of~~ Do you think you're the only one who knows how ~~Maybe that's a little~~ to cause pain? What do you ~~use~~ have over there? Needles? cages? Child's play. ~~We have everything we need~~

Bfr ~~Do you think you're the only one who likes the research?~~

on

~~You're a man of research. I'm sure you can~~

I've never seen ~~a man~~ inside the stomach of a man but I'm very willing to learn.

We will ~~stick~~ cut your throat. We will burn every inch of your skin bit by bit, and then we'll go deeper.

You had better shut down the farm or you'll regret it.

If you don't shut down the farm, I've got a gun that will force you to do so.

You really should close down this farm of yours. The coppers will get on to you and believe me - You will get punished. They don't take this sort of thing lightly you know. In fact, I've heard they can get quite ~~hostile~~ ~~aggh~~ ~~violent~~ passionate in ensuring justice is served if you get what I mean close it now.

Listen, mate. You want to play the big man breaking criminals like shit? Well, I can be the big man too. I will gun ~~them~~ you down and ~~your fam~~ blow your house into fucking bits. Now how about closing that farm?

- Mr West we know where you live and what you have been up to
- ~~Tom West people despise of your behaviour, we see all you do~~
- Tom West I have bombs and weapons you will not see the light of tomorrow

8 References

Archer, D., 1999. Exploring bullying culture in the para-military organisation. International journal of manpower, 20(1/2), pp.94-105.

Audacity [2012]: Audacity. Version 2.1.0. Audio editor and recorder. Retrieved at: https://sourceforge.net/projects/audacity/

Ash, R., Cook, J., Dannenburg, R., Drung, B., Johnson, V., Lucius, L., Shaw, M. (2015). Audacity (Version 2.1.0) [Mac].

Austin, J.L., 1975. How to do things with words (Vol. 88). Oxford university press.

Banse, R. and Scherer, K. (1996). Acoustic profiles in vocal emotion expression. Journal of Personality and Social Psychology, 70(3), pp.614-636.

Banse, R. and Scherer, K.R., 1996. Acoustic profiles in vocal emotion expression. Journal of personality and social psychology, 70(3), p.614.

Beck, J., 2007. Vocal Profile Analysis Scheme: A Users Manual. Queen Margaret University, Edinburgh.

Berthet, M., Neumann, C., Mesbahi, G., Csar, C. and Zuberbhler, K., 2018. Contextual encoding in titi monkey alarm call sequences. Behavioral Ecology and Sociobiology, 72(1), p.8.

Boersma, P. and Weenink, D. (2014). Praat: doing phonetics by computer.

Breakwell, G. (1997). Coping with aggressive behaviour. Leicester: British Psychological Society, pp.26-42.

Burczycka, M. and Conroy, S., 2018. Family violence in Canada: A statistical profile, 2016. Juristat: Canadian Centre for Justice Statistics, pp.1-96.

Burkhardt, F. and Sendlmeier, W.F., 2000. Verification of acoustical correlates of emotional speech using formant-synthesis. In ISCA Tutorial and Research Workshop (ITRW) on speech and emotion.

Centers for Disease Control and Prevention. 2013. Workplace Violence Prevention for Nurses. [ONLINE] Available at: https://wwwn.cdc.gov/wpvhc/Course.aspx/Slide/Unit67. [Accessed 8 November 2017].

Chandler, E.L., 2017. Animal rights extremism and the use of intimidation (Doctoral dissertation, University of Surrey).

Cohen, K., Johansson, F., Kaati, L. and Mork, J.C., 2014. Detecting linguistic markers for radical violence in social media. Terrorism and Political Violence, 26(1), pp.246-256.

Congress of Phonetic Sciences - Montreal, 1035-1039, 1972.

Criminal Damage Act 1971, s 2

Criminal Justice and Licensing (Scotland) Act 2010, s 38 (1)

Culpeper, J., 2011. Impoliteness: Using language to cause offence (Vol. 28). Cambridge University Press.

Dabbs Jr, J.M. and Mallinger, A., 1999. High testosterone levels predict low voice pitch among men. Personality and individual differences, 27(4), pp.801-804.

David Antonio Mendes v DPP [2015] EWHC 1597 (Admin)

De Becker, G., 1997. The gift of fear. Boston: Little, Brown.

Deisinger, G., Randazzo, M., ONeill, D. and Savage, J., 2008. The handbook for campus threat assessment and management teams. Stoneham, MA: Applied Risk Management.

Dietz, P.E., Matthews, D.B., Van Duyne, C., Martell, D.A., Parry, C.D., Stewart, T., Warren, J. and Crowder, J.D., 1991. Threatening and otherwise inappropriate letters to

Hollywood celebrities. Journal of Forensic Science, 36(1), pp.185-209.

Dietz, P.E., Matthews, D.B., Van Duyne, C., Martell, D.A., Parry, C.D., Stewart, T., Warren, J. and Crowder, J.D., 1991. Threatening and otherwise inappropriate letters to Hollywood celebrities. Journal of Forensic Science, 36(1), pp.185-209.

Dinakar, K., Reichart, R. and Lieberman, H., 2011. Modeling the detection of Textual Cyberbullying. The Social Mobile Web, 11(02), pp.11-17.

DPP v Richie Smith [2017] EWHC 3193

Dzogang, F., Lesot, M.J., Rifqi, M. and Bouchon-Meunier, B., 2010, March. Analysis of texts emotional content in a multidimensional space. In Int. Conf. on Kansei Engineering and Emotion Research, KEER (pp. 877-886).

Elizabeth A. Vacura and Raymond Vacura, Appellants, v. Carol L. Plott, Appellee, 666 F.2d 1200 (8th Cir. 1981)

Elonis v. United States, 575 U.S. (2015)

Evans, S., Neave, N., Wakelin, D. and Hamilton, C., 2008. The relationship between testosterone and vocal frequencies in human males. Physiology Behavior, 93(4-5), pp.783-788.

Ewan Paterson v PF Airdrie [2014] HCJAC 87

Facebook.com. (2018). Community Standards | Facebook. [online] Available at: https://www.facebook.com/communitystandards/violencecriminalbehavior [Accessed 13 Jul. 2018].

Fecher, N. and Watt, D., 2013. Effects of forensically-realistic facial concealment on auditory-visual consonant recognition in quiet and noise conditions. In Auditory-Visual Speech Processing (AVSP) 2013.

Fitch, W.T., de Boer, B., Mathur, N. and Ghazanfar, A.A., 2016. Monkey vocal tracts are speech-ready. Science Advances, 2(12)

Fraser, B., 1998. Threatening revisited. International Journal of Speech, Language and the Law (formerly Forensic Linguistics) 5(2): 159173.

Gales, T., 2011. Identifying interpersonal stance in threatening discourse: An appraisal analysis. Discourse Studies, 13(1), pp.27-46.

Gales, T., 2012. Linguistic analysis of disputed meanings: Threats. In: Carol A. Chapelle (ed), The Encyclopedia of Applied Linguistics. Oxford: Blackwell. Online resource. DOI: 10.1002/9781405198431.wbeal0711

Gales, T.A., 2010. Ideologies of violence: A corpus and discourse analytic approach to stance in threatening communications. University of California, Davis.

Gingiss, P., 1986. Indirect threats. Word 17(3): 155158.

Gottschalk, L. A. Bechtel, R. J. (2001). PCAD 2000: Psychiatric content analysis and diagnosis. Available from GB Software at http://www.gb-software.com.

GOV.UK. (2016). Bomb Threats Guidance. [online] Available at: https://www.gov.uk/government/publications/bomb-threats-guidance [Accessed 18 Sep. 2018].

Hammer, H.L., 2014, September. Detecting threats of violence in online discussions using bigrams of important words. In Intelligence and Security Informatics Conference (JISIC), 2014 IEEE Joint (pp. 319-319). IEEE.

Hammer, H.L., 2016, October. Automatic detection of hateful comments in online discussion. In International Conference on Industrial Networks and Intelligent Systems (pp. 164-173). Springer, Cham. Help.twitter.com. (n.d.). Violent threats and glorification of violence. [online] Available at: https://help.twitter.com/en/rules-and-policies/violent-threats-glorification [Accessed 18 Sep. 2018].

Hemsworth, P.H., Gonyou, H.W. and Dziuk, P.J., 1986. Human communication with pigs: the behavioural response of pigs to specific human signals. Appl. Anita. Behav.Sci., 15 : 45–54.

Jackson, D., Wilkes, L. and Luck, L. (2014). Cues that predict violence in the hospital setting: Findings from an observational study. Collegian, 21(1), pp.65-70.

James, D.V., Mullen, P.E., Path, M.T., Meloy, J.R., Farnham, F.R., Preston, L. and Darnley, B., 2008. Attacks on the British royal family: The role of psychotic illness. Journal of the American Academy of Psychiatry and the Law Online, 36(1), pp.59-67.

Johnstone, T. and Scherer, K.R., 2000. Vocal communication of emotion. Handbook of emotions, 2, pp.220-235.

Jolly v HMA [2013] SCCR 511

Kaplan, J.P., 2016. Case report: Elonis v. United States. International Journal of Speech, Language the Law, 23(2).

Karatepe, O.M., Yorganci, I. and Haktanir, M., 2009. Outcomes of customer verbal aggression among hotel employees. International Journal of Contemporary Hospitality Management, 21(6), pp.713-733.

Kienast, M. and Sendlmeier, W. (2000). Acoustical analysis of spectral and temporal changes in emotional speech. In: ISCA Tutorial and Research Workshop (ITRW) on Speech and Emotion.[online] Available at:
http://www.iscaspeech.org/archive_open/archive_papers/speech_emotion/spem_092.pdf [Accessed 12 Mar. 2015].

Kim, S., Ideker, K. and Todicheeney-Mannes, D. (2011). Usefulness of Aggressive Behaviour Risk Assessment Tool for prospectively identifying violent patients in medical and surgical units. Journal of Advanced Nursing, 68(2), pp.349-357.

Knzel, H.J. 1997. Some general phonetic and forensic aspects of speaking tempo. Forensic Linguistics 4: 48-83.

Kwan, G.C.E. and Skoric, M.M., 2013. Facebook bullying: An extension of battles in school. Computers in human behavior, 29(1), pp.16-25.

Laver, J., 1994. Principles of phonetics. Cambridge University Press. Pp. 396-423, 541-5

Leech, G., Rayson, P., Wilson, A. (2001). Word frequencies in written and spoken English: based on the British National Corpus. Harlow, Longman.

Lehiste, I. and Lass, N.J., 1976. Suprasegmental features of speech. Contemporary issues in experimental phonetics, 225, p.239.

Lewis-Webb v. Qualico Steel Co., Inc., 929 F. Supp. 385 (M.D. Ala. 1996)

Linaker, O. and Busch-Iversen, H. (1995). Predictors of imminent violence in psychiatric inpatients. Acta Psychiatrica Scandinavica, 92(4), pp.250-254

LIWC2015 [2015]: retrieved at http://liwc.wpengine.com/

MacLean, S.A. and Bonter, D.N., 2013. The sound of danger: threat sensitivity to predator vocalizations, alarm calls, and novelty in gulls. PloS one, 8(12)

McCrory, E.J., De Brito, S.A., Sebastian, C.L., Mechelli, A., Bird, G., Kelly, P.A. and Viding, E., 2011. Heightened neural reactivity to threat in child victims of family violence. Current Biology, 21(23), pp.R947-R948.

Mileva, M., Tompkinson, J., Watt, D. and Burton, A.M., 2018. Audiovisual integration in social evaluation. Journal of Experimental Psychology: Human Perception and Performance, 44(1), p.128.

Morris, S., Ward, D., Butt, R., 2006 'Jail for animal rights extremists who stole body of elderly woman from her grave, The Guardian, 12 May: accessed online at https://www.theguardian.com/uk/2006/may/12/animalwelfare.topstories3

Morton, E.S., 1977. On the occurrence and significance of motivation-structural rules in some bird and mammal sounds. The American Naturalist, 111(981), pp.855-869.

Mullen, P., Path, M. and Purcell, R. (2000). Stalkers and their victims. Cambridge: Cambridge University Press.

Napier, M. R., Mardigian., 2003. Threatening messages: The essence of analyzing communicated threats. Public Venue Security, September/October, pp. 1619.

National Association for the Advancement of Colored People v. Claiborne Hardware Co., 458 U.S. 886 (1982)

Newman, M.L., Groom, C.J., Handelman, L.D. and Pennebaker, J.W., 2008. Gender differences in language use: An analysis of 14,000 text samples. Discourse Processes, 45(3), pp.211-236.

Nini, A., 2017. Register variation in malicious forensic texts. International Journal of Speech, Language the Law, 24(1).

Office for National Statistics (2018). Crime in England and Wales: year ending March 2018. [online] Available at:
https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/bulletins/crimeinenglandandwales/yearendingmarch2018 [Accessed 17 Sep. 2018].

Ohala, J.J., 1984. An Ethological Perspective on Common Cross-Language Utilization of F of Voice. Phonetica, 41(1), pp.1-16.

Ohala, J.J., Hinton, L. and Nichols, J., 1997, August. Sound symbolism. In Proc. 4th Seoul International Conference on Linguistics [SICOL] (pp. 11-15).

O'Toole, M. (n.d.). The School Shooter: A THREAT ASSESSMENT PERSPECTIVE. [ebook] Quantico: Critical Incident Response Group (CIRG), National Center for the

Analysis of Violent Crime (NCAVC). Available at: https://www.fbi.gov/file.../stats-services-publications-school-shooter-school-shooter [Accessed 1 Oct. 2017]

Overstrom, N.A., 1983. Association between burstpulse sounds and aggressive behavior in captive Atlantic bottlenosed dolphins (Tursiops truncatus). Zoo Biology, 2(2), pp.93-103.

Packer v. Skid Roe, Inc., 938 F. Supp. 193 (S.D.N.Y. 1996)

Pennebaker, J. (2011). Using computer analyses to identify language style and aggressive intent: The secret life of function words. Dynamics of Asymmetric Conflict, [online] 4(2), pp.92-102. Available at: https://www.tandfonline.com/doi/pdf/10.1080/17467586.2011.627932.

Pennebaker, J.W., Booth, R.J., Boyd, R.L., Francis, M.E. (2015). Linguistic Inquiry and Word Count: LIWC2015. Austin, TX: Pennebaker Conglomerates (www.LIWC.net)

Pennebaker, J.W., Mehl, M.R. and Niederhoffer, K.G., 2003. Psychological aspects of natural language use: Our words, our selves. Annual review of psychology, 54(1), pp.547-577.

People v. Miller (Crim. No. 3921. Second Dist., Div. Three. Jan. 18, 1946.)

Pich, J., Hazelton, M., Sundin, D. and Kable, A., 2011. Patient-related violence at triage: A qualitative descriptive study. International emergency nursing, 19(1), pp.12-19.

Piolat, A. and Bannour, R., 2009. An example of text analysis software (EMOTAIX-Tropes) use: The influence of anxiety on expressive writing. Current psychology letters. Behaviour, brain cognition, 25(2, 2009).

Pisoni, D.B. and Martin, C.S., 1989. Effects of alcohol on the acousticphonetic properties of speech: perceptual and acoustic analyses. Alcoholism: Clinical and Experimental Research, 13(4), pp.577-587.

Boersma, P. and Weenink, D. (2016). Praat: doing phonetics by computer [Computer program]. Version 6.0.17, retrieved from http://www.praat.org/

Public Order Act 1986, s 5

R v Cakmak; R v Cavcav; R v Talay; R v Can; R v Karaaslan; R v Durukanoglu [2002] EWCA Crim 500

R v Colin Albert Jennison [2000] 2 Cr. App. R. S. 213

R v Flynn and St John [2008] Crim LR 799

R v JDP [2010] EWCA Crim 3224

R v Russell (Robert William) [2010] EWCA Crim 778

Reid Meloy, J., Hoffmann, J., Guldimann, A. and James, D., 2012. The role of warning behaviors in threat assessment: An exploration and suggested typology. Behavioral sciences the law, 30(3), pp.256-279.

Robinson, A., Pinchevsky, G. and Guthrie, J. (2016). A small constellation: risk factors informing police perceptions of domestic abuse. Policing and Society, 28(2), pp.189-204.

Ronson, J., 2015. How one stupid tweet blew up Justine Saccos life. New York Times, 12.

Rooney v Brown [2013] SCCR 334

Scherer, K. and Oshinsky, J. (1977). Cue utilization in emotion attribution from auditory stimuli. Motivation and Emotion, 1(4), pp.331-346

Scherer, K.R., 1977. Effect of stress on fundamental frequency of the voice. The Journal of the Acoustical Society of America, 62(S1), pp.S25-S26.

Schrder, M., Cowie, R., Douglas-Cowie, E., Westerdijk, M. and Gielen, S. (2001). Acoustic Correlates of Emotion Dimensions in View of Speech Synthesis. In: INTERSPEECH. [online] Available at: http://ftp://134.96.191.190/papers/local/schroeder_etal2001.pdf [Accessed 11 Mar. 2015].

Searle, J.R., 1989. How performatives work. Linguistics and philosophy, 12(5), pp.535-558.

Sell, A., Bryant, G.A., Cosmides, L., Tooby, J., Sznycer, D., Von Rueden, C., Krauss, A. and Gurven, M., 2010. Adaptations in humans for assessing physical strength from the voice. Proceedings of the Royal Society B: Biological Sciences, 277(1699), pp.3509-3518.

Shami, M. and Verhelst, W. (2007). An evaluation of the robustness of existing supervised machine learning approaches to the classification of emotions in speech. Speech Communication, 49(3), pp.201-212.

Shuy, R., 1993. Language Crimes: The Use and Abuse of Language Evidence in the Courtroom. Oxford: Blackwell.

Simons, A. and Tunkel, R. (2014). The Assessment of Anonymous Threatening Communications. In: J. Meloy and J. Hoffman, ed., International Handbook of Threat Assessment. New York: Oxford University Press, pp.205-211

Smith, S., 2008. From violent words to violent deeds? Assessing risk from FBI threatening communication cases. In: Meloy, J.R., Sheridan, L. Hoffmann, J. (eds.), Stalking, Threatening, and Attacking Public Figures: A Psychological and Behavioral Analysis. Oxford: Oxford University Press. pp. 435455.

Smith, S.S., 2006. From Violent Words to Violent Deeds?: Assessing Risk from Threatening Communications (Doctoral dissertation, Georgetown University).

Solan, L.M. and Tiersma, P.M., 2005. Speaking of crime. The language of Criminal Justice. Chicago.

Soltis, J., King, L.E., Douglas-Hamilton, I., Vollrath, F. and Savage, A., 2014. African elephant alarm calls distinguish between threats from humans and bees. PLoS One, 9(2)

Spitzberg, Brian H. and Gawron, Jean Mark (2016) "Toward Online Linguistic Surveillance of Threatening Messages," Journal of Digital Forensics, Security and Law:

Vol. 11 : No. 3 , Article 7.DOI: https://doi.org/10.15394/jdfsl.2016.1418 Available at: https://commons.erau.edu/jdfsl/vol11/iss3/7

SPSS (2015) SPSS Statistics. [Software] Version 22. Chicago: SPSS

State v. McDowell 620 A.2d 94 (1993)

State v. Milner 571 N.W.2d 7 (1997)

Storey, K., 1995. The language of threats. International Journal of Speech, Language and the Law (formerly Forensic Linguistics) 2(1): 7480.

Takefuta, Y., Jancosek, E. G., Brunt, M. (1971) A statistical analysis of melody curves in the intonation of American English, Proceedings of the 7th International Congress of Phonetic Sciences -Montreal, pp.1035-1039, 1971.

Thompson, W. and Balkwill, L. (2006). Decoding speech prosody in five languages. Semiotica, 2006(158).

Threat Assessment Procedures Manual. (2017). [ebook] The School Board of Broward County, Florida. Available at:
http://bcpsmentalhealth.com/downloads/ThreatAssessmentManualUpdate2017.pdf [Accessed 8 Oct. 2017].

Tompkinson, J. (2015). Accent evaluation and the perception of threatening speech acts [Presentation given at the University of York].

Tropes software [2014] (Version 8.4.4): retrieved at
http://www.semantic-knowledge.com/tropes.htm

Tsantani, M., Belin, P. and Mcaleer, P., 2016. Low vocal pitch preference drives first impressions of trustworthiness and dominance in non-contextual scenarios. Perception.

Ultan, R. (1978). Size-sound symbolism. Universals of Human Language, 2, 525568.

United States of America, Appellee, v. Antonio Clemmons, Appellant, 461 F.3d 1057 (8th Cir. 2006)

United States v. Dinwiddie, 885 F. Supp. 1286 (W.D. Mo. 1995)

Van Dijk, J.J., 2012. The International Crime Victims Survey: Latest results and outlook. Newsletter European Society of Criminology, pp.24-34.

van Kesteren JN, Mayhew P and Nieuwbeerta P (2000) Criminal Victimisation in Seventeen Industrialised Countries; Key Findings from the 2000 International Crime Victim Survey. The Hague: Onderzoek en Beleid, No 187, Ministry of Justice/WODC.

Vartia, M. and Hyyti, J., 2002. Gender differences in workplace bullying among prison officers. European Journal of Work and Organizational Psychology, 11(1), pp.113-126.

Virginia v. Black, 538 U.S. 343 (2003)

Watt, D. Kelly, S. and Llamas, C., 2013. Inference of threat from neutrally-worded utterances in familiar and unfamiliar languages. York Papers in Linguistics (Series, 2, pp.99-120).

Watts v. United States 394 U.S. 705, 89 S.Ct. 1399, 22 L.Ed.2d 664 (1969)

West, M. 1953. A General Service List of English Words. London: Longman.

Wester, A., vrelid, L., Velldal, E. and Hammer, H.L., 2016. Threat detection in online discussions. In Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (pp. 66-71).

Williams, C. and Stevens, K. (1972). Emotions and Speech: Some Acoustical Correlates. The Journal of the Acoustical Society of America, 52(4), pp.1238-1249.

Wilson, C.M., Douglas, K.S. and Lyon, D.R., 2011. Violence against teachers: Prevalence and consequences. Journal of interpersonal violence, 26(12), pp.2353-2371.

Winstanley, S. and Whittington, R., 2004. Aggression towards health care staff in a UK general hospital: variation among professions and departments. Journal of clinical nursing, 13(1), pp.3-10.

Xu, Y., Kelly, A. and Smillie, C., 2013. Emotional expressions as communicative signals. Prosody and Iconicity, pp.33-60.

Yamanaka, N., 1995. On indirect threats. International Journal for the Semiotics of Law 8(2): 3752.

Yildirim, S., Bulut, M., Lee, C., Kazemzadeh, A., Busso, C., Deng, Z., Lee, S. and Narayanan, S. (2004). An acoustic study of emotions expressed in speech. In: INTERSPEECH. [online] Available at: http://ecs.utdallas.edu/research/researchlabs/msp-lab/publications/Yildirim_2004.pdf [Accessed 11 Mar. 2015].

Zetterholm, E. (1998). Prosody and voice quality in the expression of emotions. In: International Conference on Spoken Language Processing. [online] Available at: http://www.mirlab.org/conference_papers/International_Conference/ICSLP%201998/PDF/AUTHOR/SL981043.PDF [Accessed 29 Mar. 2015].