

## Series Preface

The Elements in Forensic Linguistics series from Cambridge University Press publishes across four main topic areas (1) investigative and forensic text analysis; (2) the study of spoken linguistic practices in legal contexts; (3) the linguistic analysis of written legal texts; (4) explorations of the origins, development, and scope of the field in various countries and regions. *Spoken Threats from Production to Perception* by James Tompkinson is situated in the second of these categories and presents results from a series of corpus-based and experimental studies that investigate how verbal threats and ‘a threatening tone of voice’ are both produced and perceived in a variety of contexts.

As an academic working in the field of forensic phonetics, James Tompkinson brings a balanced perspective to the analysis of spoken threats. Drawing on categories proposed by French and Watt (2018), Tompkinson explains that previous work on spoken threats has largely been *responsive* in nature. That is, it has occurred in response to a research question posed to address a particular forensic case. In this Element, Tompkinson provides much-needed research from an *anticipatory* perspective. That is, research that investigates questions with potential for real-world applications in a wide variety of contexts, including forensic cases.

Thus, the contents of this Element provide several important contributions to the study of threatening communications. First, Tompkinson introduces us to a new Corpus of Spoken Threats (CoST), which is now available for further research purposes upon request to the author; second, he provides a comparison of the linguistic features produced in written vs spoken threats using existing literature on written threats and findings from the new CoST corpus; and third, he offers a synthesis of recent experimental work on the perception of spoken threats, furthering our understanding of what a threat actually is and how they are perceived. We hope this innovative work will encourage additional research on this context-based language crime.

Tammy Gales  
Series Editor

## Prologue

In the time, I have spent conducting academic research, I have often found that some of the most interesting projects begin almost by accident. One example of this is how the research on which this contribution to the Elements in Forensic Linguistics Series is based came to exist. In the autumn of 2013, I had just finished my undergraduate degree in English Language and Linguistics at the University of York. Around this time, I met with Dominic Watt in the back bar of the York

Brewery pub to discuss postgraduate study options. Dom happened to mention a criminal case that he was aware of, where the use of a so-termed ‘threatening tone of voice’ was a factor in a custody officer perceiving a spoken utterance from a detainee as a reiteration of a threat to kill. At first mention, this seemed unproblematic to me; a threat delivered in a threatening tone of voice. So what? But it quickly became apparent that Dom and I could not begin to distil what a ‘threatening tone of voice’ would be in terms of definable features of speech. As our discussion progressed, we questioned whether it was even possible to define what a ‘threatening tone of voice’ was, despite its existence in common parlance, or how spoken threats might differ from threats delivered in written form. In fact, our only source of agreement was that this topic required further research, and that it would take a team of researchers with equal interest in both forensic linguistics and forensic phonetics to undertake the task. Over the coming years, Dom would supervise both me and Sarah Kelly as we pursued MSc and PhD research into the language of spoken threats and tried to set about addressing some of the questions we had about this common type of language crime. The aim of this Element is to provide some of the answers to those questions we first posed nearly 10 years ago.

## 1 Introduction

In 2015, the UK Parliamentary Office of Science and Technology produced a report entitled ‘Forensic Language Analysis’ (Bunn and Foxen, 2015). This report was designed to provide an introductory and accessible review of research conducted in the academic fields of forensic linguistics and forensic speech science, aimed at a non-specialist audience. The authors of the report provide a review of different topics including authorship analysis, speaker comparison and transcription. They also highlight the scepticism and concerns that forensic phoneticians have raised over claims that the voice can be used as a tool for procedures such as deception detection. However, the most interesting part of the report for readers of this Element was the observation of a disjointed relationship between non-linguists (in this case jurors) and linguistic experts. The authors state that ‘jurors expect certain procedures to be possible which experts assert are not, such as personality analysis, determining truth and falsity, and assessing threat in speech intonation (although this is a research interest)’ (Bunn and Foxen, 2015: 3).

Bunn and Foxen’s (2015) report highlights that non-linguists believe that it is possible for linguistic experts to do things that they cannot in relation to certain aspects of language analysis, including in the analysis of spoken threats. Indeed, a question I am frequently asked when I tell non-linguists that I work in forensic speech science is ‘does that mean you can tell if someone is lying from their

voice?’. If I am ever asked to explain my work on spoken threats, non-linguists will often ask ‘so does that mean you can tell the police when a threat is serious based on the way that someone speaks?’. The idea that experts could be asked to make such important, categorical decisions based on voice and speech patterns alone is, of course, unrealistic. However, the expectation from non-linguists that these kinds of decisions can be made by language experts (not helped by fictional TV programmes which show ‘experts’ making *exactly* these kinds of decisions), is highly problematic because it creates a disparity between non-linguistic expectations and linguistic reality.

Having considered this key misconception about the work of the forensic linguist in relation to spoken threats, it is equally important to examine what forensic linguists *can* offer to the study of threatening language. To do this, we must first consider where the study of threats falls within the overall field of forensic linguistics. In their textbook ‘An Introduction to Forensic Linguistics: Language in Evidence’, Coulthard, Johnson and Wright (2017) break the field down into two broad categories. The first of these is ‘The Language of the Legal Process’, which covers areas such as the language of the law, courtrooms and police interviews. The second, ‘Language as Evidence’, focusses on areas such as authorship attribution, forensic phonetics, plagiarism detection, and expert witness evidence. This split is similar to both Larner’s (2015) and Nini’s (2019) discussions of ‘descriptive’ and ‘investigative’ forensic linguistics. Generally speaking, a distinction can be drawn between work which describes linguistic phenomena which have forensic relevance, and work which provides evidential linguistic analysis to assist criminal or civil trials.

Coulthard, Johnson and Wright (2017) provide one of the most comprehensive overviews of forensic linguistics as an academic discipline, and yet it is not automatically clear where the analysis of threats should be placed within the two broad areas of forensic linguistics that they propose. This is arguably because the linguistic analysis of threatening language transcends the boundary between these research areas, depending on the particular question being asked. Perhaps a more useful categorisation for research on spoken threats is the split between ‘anticipatory’ and ‘responsive’ research, as discussed by French and Watt (2018) in relation to research impact. French and Watt (2018: 153) categorise anticipatory research as research that has potential for real-world applications but doesn’t respond directly to a specific case or an urgent, immediate need for data. Contrastingly, responsive research does the opposite and aims to answer a specific question about a specific case through linguistic research or analysis. Relating this to the study of threatening language, general research on threats as a type of language crime would be largely *anticipatory*, whereas attempting to use linguistic research to help answer (or provide legitimate reasons not to

answer) questions such as ‘is this particular utterance a threat?’ would be *responsive*. However, in order to answer a specific question about a particular threatening utterance, it is essential to understand how threats work at a more general level. In other words, we need anticipatory research to be able to respond appropriately to specific questions about threatening language.

I first encountered this issue in 2016, when I was asked to write an article for *The Conversation* about spoken threats. The request for this article came following news that a series of hoax bomb threats had been made to schools across the United Kingdom.<sup>1</sup> Rather than comment on the specifics of this case in a completely responsive way, I opted to discuss why, following research which showed that there are no consistent phonetic cues to deception (see, e.g., Kirchhübel, 2013), attempting to identify whether a threat was a hoax or not based on a speaker’s voice was unrealistic. I also discussed some research which highlighted why threat perception based on a speaker’s voice can be linked to unhelpful linguistic stereotypes. My concluding words in that article were as follows:

*Just because a listener may be inclined to think a speaker sounds more threatening based on different aspects of their voice, there is no basis to say this makes a speaker any more likely to commit any action they threaten. And the less we rely on stereotyped impressions of speech in potentially high-stakes situations, the better. (Tompkinson, 2016)*

The work in this Element has been designed to offer more research to help expand on the kinds of issues I first discussed in the above article. The work presented here should, according to French and Watt’s (2018) definitions, be classed as anticipatory rather than responsive research. The work in this Element does not respond to a specific problem or case, but instead presents more general analyses to help illuminate areas of interest and contribute to knowledge about threatening language. Of course, my contribution here is far from the first in this area. Linguistic research on threatening language stretches back several decades and offers many complimentary and contrasting perspectives. This existing body of research is discussed and evaluated in Section 2, but there are two specific aspects which I argue are under-researched and where this Element provides a more specific contribution.

The first of these under-researched areas concerns the treatment of spoken and written threats. The balance of research on threatening language is much more heavily weighted towards the written modality. There are also some studies which, while offering useful perspectives on the topic, do not

<sup>1</sup> <https://theconversation.com/bomb-scares-can-you-judge-a-threat-from-the-voice-on-the-phone-60073>

separate spoken threats from written ones. This area is identified by Gales (2021) as an area which requires further research, and Section 3 of this Element addresses this directly. I firstly present an analysis of key linguistic features within a newly created spoken threat corpus followed by a comparative analysis with previous research on written threats. The work in Section 3 begins to bridge the gap between linguistic research on threats delivered in the two modalities, while also contributing to our understanding of the similarities and differences between spoken and written threats in a way which has not yet been done.

The second under-researched area that this Element addresses is whether the way in which a spoken threat is uttered can affect perceptions of the speaker. In Section 4, I describe a programme of experimental research which was conducted to critically examine how both speech and speaker factors can contribute to someone being perceived as sounding threatening. This ties into the notion of a ‘threatening tone of voice’ and whether there is a linguistic or phonetic basis for such a label. I also highlight both the strengths and weaknesses of experimental research which examines perceptions of spoken threats, as well as the dangers of directly applying experimental results to specific legally relevant situations.

The overall aim of this contribution to the Elements in Forensic Linguistics series is to advance knowledge and promote further debate over the central question of what we can, and perhaps more importantly what we cannot, say about threatening language.

## 2 Threatening Language: A Research Review

### 2.1 Defining Threats

Threats form a substantial part of our everyday language use. There are many possible reasons why a speaker may threaten someone, and we all make threats from time-to-time to achieve a specific course of action. In most cases, these kinds of everyday threats do not express any criminal intent and are not illegal. Consider a mother who threatens her child that their favourite toy will be taken away unless the child puts their shoes on and leaves the house quickly. Although this interaction is not illegal, a clear threat is made by the mother towards her child in the guise of ‘if you don’t do the thing I want you to do [put shoes on and leave the house] then something bad will happen [a toy will be taken away]’. Another example of an authentic but non-illegal threat, discussed by Solan and Tiersma (2015: 223), is of a boss threatening an employee that they will be fired if they do something wrong at work. The use of threats in everyday language brings together the notions of actions and consequences, binding language use

to a person's ability to ensure that unfavourable things can happen if certain conditions are not met. Storey (1995: 74) goes as far as to say that threats are simply 'a way of life', with Milburn and Watman (1981: 2) commenting that they provide speakers with a way of exerting personal and social control in unpredictable situations or environments.

However, when threats become illegal, they can serve as both standalone crimes and form part of other serious crimes such as robbery and extortion (Yamanaka, 1995: 38). Solan and Tiersma (2015: 224) state that threats are often used to accomplish serious crimes, with Greenawalt (1989: 92) also explaining that criminal acts frequently involve threats which aim to get an innocent victim to commit to an unfavourable course of action. Threats can also be illegal if they are directed towards certain people, such as the President of the United States (Danet et al., 1980) or members of the United Kingdom's royal family (Solan and Tiersma, 2015: 233). The dual nature of threats as standalone crimes and as an integral part of other crimes is captured by the definition of 'threat' provided by the Oxford Dictionary of Law (Law and Martin, 2009), which states that a threat is 'the expression of an intention to harm someone with the object of forcing them to do something' and that threats are 'an ingredient of many crimes'. The Oxford Dictionary of Law provides a more detailed definition for 'threatening behaviour', which is listed as the use of 'threatening, abusive or insulting words or behaviour' towards another person (Law and Martin, 2009). The statement and expression of intention in a threat is sufficient to uphold its status as a threat, even if the speaker has no actual intention to carry out the threatened action. One example of this is the case of Seif Eldin Mustafa, who hijacked EgyptAir flight MS181 in March 2016 and threatened to blow up the aircraft using a belt containing explosives (BBC News, 2016). It was subsequently revealed that the belt contained no explosives and therefore Mustafa could never have intended to blow up the aircraft, but the threat was considered real by security staff and those on board the plane and was therefore valid.

Despite such definitions, Gales (2016: 3) has previously warned that there is a lack of understanding about what threatening language 'actually is', and of the potential dangers when those tasked with assessing linguistic aspects of threats rely on personal or stereotypical assumptions rather than evidence-based approaches. This potential problem is further compounded when the modality of a threat is spoken rather than written. Spoken threats provide an additional problem in that unless a recording of a threatening utterance exists, they are momentary and are therefore more heavily reliant on listeners' perceptions of the speaker's intentions. There is also something of a shortage of research examining how spoken threats are perceived by listeners (Watt, Kelly and

Llamas, 2013), although some steps have been taken in more recent years to address this (e.g., Kelly, 2018; Tompkinson, 2018; Tompkinson et al., 2023).

There are several cases which highlight the need for further research into listener evaluations of spoken threats. One such example is documented in Watt, Kelly and Llamas (2013) and comes from a 2012 crown court trial where the defendant was accused of reiterating a previously unrecorded threat to kill by uttering the words *‘I will do summat [a northern English dialect term for ‘something’] about it when I get out and it won’t be with guns or anything like that’*. This utterance was produced following a situation where the defendant had been held in a police cell and was remonstrating to a custody officer that he wanted to be released. This custody officer was the hearer of the alleged reiteration of the previous unrecorded threat to kill.

This case provides one example of what Gales (2010) terms an indirect threat, where a threat is judged to have been uttered, yet the wording of the utterance does not explicitly signal intent-to-harm on the part of the speaker. Indirect threats do not overtly make clear that a threat is being made, and could, on wording alone, be classified as other types of speech acts including warnings, insults, complaints or promises. In the example above, the vague nature of the phrase *‘I will do summat about it’* meant that listener interpretation was required to determine what that ‘something’ was, and by extension whether the speaker had criminal intentions or not. The interpretation that *‘I will do summat about it when I get out and it won’t be with guns or anything like that’* constituted a serious threat would require listener inference of the speaker’s intentions. The speaker’s words in this case, if taken in their most literal interpretation, specifically ruled out the use of guns or similar weapons, and yet the utterance was still interpreted as a reiteration of a serious death threat. Watt, Kelly and Llamas (2013) also point out that during the subsequent trial, the custody officer’s testimony identified that the defendant’s behaviour, the surrounding context and the fact that he used an aggressive tone of voice, served as evidence which supported the interpretation of the utterance in question as a serious threat.

Another example where perceptions of a speaker’s voice had a role a trial involving spoken threats is taken from the Danish Supreme Court (case number U.2016.1939 H – TfK2016.491H)<sup>2</sup>. In this case, a man was accused of threatening to cut a fellow employee’s throat. As part of the defence, the accused threatener stated that because he had a low-pitched voice, he was often perceived as sounding angry. The translated and original text from the court report is produced below:

<sup>2</sup> I am grateful to Professor Tanya Karoli Christensen for alerting me to this case, and for providing the relevant background information and translations.



**English**

*The defendant is very careful with how he phrases things since he is sometimes misunderstood and perceived as angry because he has a very deep voice. He never raises his voice since nothing good comes from it anyway. He can, however, be somewhat direct in his demeanour.*

**Danish**

*Tiltalte passer meget på, hvordan han formulerer sig, idet han sommertider bliver misforstået og opfattet som sur fordi han har en meget dyb stemme. Han hæver aldrig stemmen, da man sjældent får noget ud af det alligevel. Han kan dog somme tider godt være lidt kontant i sin fremtræden.*

Here, the defendant's perception of his own voice was offered as a mitigating circumstance in court. Furthermore, throughout the case, the defendant was described by the hearer of the threat as sounding both angry and frustrated. Ultimately in this case, the defendant was found guilty and sentenced to a fine and 30 days imprisonment. But this example does highlight the complexity of using aspects of voice as evidence in criminal trials involving spoken threats.

## 2.2 Types of Threats

Shuy (1993) classifies a threat as a type of language crime. There are two broad types of verbal threats: direct and indirect (Gales, 2010). A direct threat overtly states that something unfavourable will happen and potentially also include information about the time, place and people that will be involved in the threatened action. By contrast, indirect threats are more problematic because they involve the speaker communicating more information than is contained in the words alone (Searle, 1979: 30). The potential for misinterpretation or misunderstanding is heightened when indirect threats are made, owing to a lack of expressed clarity on the part of the speaker.

Consider, for example, an utterance such as '*I know where you live*'. Based on wording alone, there is no expression of intention to perform an unfavourable act. However, as Watt, Kelly and Llamas (2013) discuss, the utterance could plausibly be interpreted as a threat given the right context and conditions. Sentences of every type of syntactic form can count as indirect threats (Fraser, 1998: 169), and these are often masked as other type of speech acts including questions ('*Are you sure you want to do that?*'), promises ('*I promise you'll get what's coming to you*'), and warnings ('*I'm warning you, I'll never forget this*').

It is also possible for utterances like those detailed above to have multiple interpretations. Take an utterance like '*Are you sure you want to do that?*'. If interpreted literally, this would be a question which would evoke a yes/no response from the hearer. However, it could equally be used to warn if the



goal is not to question the hearer but rather to get them to reflect on whether to do something potentially unfavourable. It could also be used as a threat if the unfavourable action was to be performed by the speaker to the hearer's detriment. With these kinds of utterances, the interpretation is left for the hearer to infer. Additionally, Searle (1979: 7) points out that speakers very rarely threaten by stating 'I threaten X', where 'threaten' is used performatively. This contrasts with warnings and promises, where it is perfectly plausible to declare 'I warn/promise X'.

However, despite the categorisation of threats as either direct or indirect, there is a great deal of fluidity both within and between the two categories. Consider, for example, the two hypothetical examples of potential threat utterances, below.

*I'm warning you about a bomb at York Station. It will go off this afternoon.  
I know where you live.*

Both utterances can be classified as indirect threats. The first could be classified as an indirect threat owing to the possible interpretations as either a warning or a threat, depending on the speaker's intention. Gales (2017, personal communication) classifies this type of utterance as a direct performative warning, but an indirect threat. If the utterance is interpreted literally, then it is a direct and clear warning owing to the use of 'warn' as a performative verb, whereas the threat interpretation requires listener inference as to the speaker's intentions. The speaker could simply argue that they are providing a warning which was helpful to the hearer, rather than threatening something unfavourable. However, given the severity of the action mentioned (a bomb exploding) and the mention of a clear time and a place, this utterance is more direct than 'I know where you live'. An utterance such as 'I know where you live' clearly requires a greater level of listener inference to arrive at a threat interpretation. The labelling of a threat as either direct or indirect can be seen as a method of to provide a base level of classification, with more nuanced and fluid classifications present within these overarching categories.

One such sub-category within the umbrella classifications of 'direct' and 'indirect' is the conditional threat. Both direct and indirect threats can be worded conditionally, and these types of threats are created through the incorporation of a conditionality clause into the wording (Gales, 2010: 9). Linguistically, this conditionality can be expressed in various ways but commonly takes a form such as 'if you don't do X, then I will do Y' (Milburn and Watman, 1981: 11), 'do X and we won't do Y' or 'do X or I'll do Y'.

The use of conditionality within the wording of threats relates to the relative position and control of both speakers and hearers. When there is no conditionality

clause in the wording of a threat, the speaker presents information and the listener has no control over the outcome. The speaker therefore remains in the position of power over the threatened action throughout. A hypothetical example of this would be an utterance such as *'I'm going to kill you'*, where if the speaker has the intention to commit the threatened action, there is nothing in the wording of the utterance that gives any control to the recipient. Contrast this utterance with a conditional version of the same direct threat – *'If you don't pay me the money I'm owed, I'm going to kill you'* – and the message in the utterance becomes somewhat ambiguous. The conditionality suggests that the purpose of the utterance is to get the hearer to pay the owed money to the speaker. However, as Gales (2010: 11) correctly points out, just because a speaker factors a condition into the design of their threat, there is no obligation on the speaker's part to adhere to the stated condition. This is because the speaker remains in a position of power over the hearer throughout. However, for a conditional threat to be successfully communicated, the target of the threat must believe that either the stated unfavourable action will not take place if they agree to the condition, or that the chances of avoiding the unfavourable action will be increased as a result of compliance with the condition. Taking the example provided above, this would mean that the addressee believes that if they pay the money, then the speaker is less likely to, or will not, commit the act of killing. The key factor for conditional threats is, therefore, whether the addressee believes they have control over the outcome, rather than whether they actually have any control or not.

### 2.3 The Roles of Speakers and Hearers in Threat Communication

In their review of a series of cases involving threats made towards the President of the United States, Danet, Hoffman and Kermish (1980) show that the majority of judgements rested on the so-called 'reasonable person' test. This is that if a reasonable person would interpret an utterance as a threat, then a threat has been made. However, the 'reasonable person' notion is rejected by Gingiss (1986) on the grounds that it does not attempt to define a threat, nor does it highlight the grounds upon which a so-called 'reasonable person' would interpret an utterance as being threatening. Furthermore, it is legitimate to question what a 'reasonable person' is, and what criteria would qualify someone as being 'reasonable' in the context of threat perception. Gingiss (1986: 153) argues that the assumption that both a speaker and a hearer will 'know a threat when they hear one' is insufficient for courtroom purposes, despite its status as 'the majority view'. These issues are particularly problematic with respect to indirect threats, which require a greater amount of interpretation on the part of