INFERENCE OF THREAT FROM NEUTRALLY-WORDED UTTERANCES IN FAMILIAR AND UNFAMILIAR LANGUAGES

DOMINIC WATT, SARAH KELLY & CARMEN LLAMAS
University of York

Abstract

Although verbal threats are a very common kind of language crime, the ways in which listeners interpret ostensibly 'neutral' utterances as threats are currently poorly understood. We present the results of an experiment in which monolingual English-speaking listeners were exposed to the same innocuously-worded phrase spoken in a 'neutral' or a 'threatening' way. They heard translations of the phrase in four unknown foreign languages as well as the original form in English. The listeners were asked to rate the utterances with respect to two perceived properties: (a) how threatening they thought the utterances sounded, and (b) how much 'intent' to carry out a harmful act the listeners inferred from the talker's speech. As predicted, the listeners assigned higher threat and intent ratings to the 'threatening' utterances in both English and the foreign languages than they did to the neutral ones. However, the listeners' ratings were considerably higher for both threatening and neutral utterances spoken in English than they were for the foreign language utterances. In the English condition there was also a much larger difference between the neutral and the threat utterances with respect to the overall perceived threat and intent ratings than there was for the foreign language utterances. This suggests that correctly interpreting the threatening utterances as threats is dependent upon familiarity with the language in which they were spoken. A gender effect was also found, whereby male listeners assigned higher threat and intent ratings than did women. It is suggested that men and women may respond differently to speech cues associated with threatening behaviour.

1. Introduction

Among the set of offences that are committed using only words – a set which, *inter alia*, includes blackmail, extortion, bribery, blasphemy, profanity, fraud, impersonation, perjury, defamation, and incitement to racial hatred (Shuy 2005; Tiersma & Solan 2012) – is making a verbal threat. A threat may loosely be defined as a form of expression that communicates that some undesirable state of affairs (physical injury, for instance) may or will befall the recipient or a third party as a consequence of another's purposeful actions. Existing research on this topic has tended to focus on the inference of threat from spoken utterances or written texts on the basis of their content: that is, the meanings of the words themselves, alone and in combination, and their illocutionary force: that is, what the speaker intended by uttering them (see Austin 1962; Fraser 1976, 1998; Gales 2012; Searle 1969; Shuy 1993).

Gales (2012) gives examples of three main kinds of threats: direct threats, conditional threats, and indirect threats. Direct threats assert that something bad will come about, and describe how, when and to whom it will happen, while conditional threats specify what will happen unless some condition is met. The third kind, the indirect threat, is the type of threat that is most difficult to identify as such, as it may masquerade as some other kind of speech act. It may simultaneously be interpretable as a warning, as helpful advice, or as a complaint, so its classification as a threat depends heavily upon how the listener reacts to it (Gingiss 1986; Napier & Mardigian 2003; Smith 2008; Storey 1995; Yamanaka 1995). Part of what makes

this sort of oblique utterance a threat is the supposition on the hearer's part that the speaker is potentially capable of performing some action that would materially disadvantage the hearer, or someone else. Its interpretation as a threat draws upon knowledge of relevant circumstances which is not encoded in the utterance. For instance, a comment such as 'you should be careful, you don't want to end up like person X', which might under some circumstances be interpreted as an altruistic warning, could reasonably be considered a threat if the addressee was aware that person X had recently been the victim of a violent beating. Whether the above comment is construable as threatening, and thereby has the desired effect of intimidating the recipient, is therefore contingent upon both speaker and recipient having a shared understanding of relevant background information (see further Austin 1962).

It can nonetheless be very difficult to prove that an indirect threat was *meant* as a threat, even if it was interpreted as such. The person accused of making the threat may very easily deny that he or she had any actual intention of harming someone else, and may say that no such implication was intended either. Even if the precise wording of the utterance is not in dispute, it could be claimed that it was meant in an advisory way, or simply as a joke. The wording of the utterance may in some cases look so neutral or harmless that the debate reduces to the level of irresolvable gainsaying. An illustration of this scenario is a British court case from 2012 that revolved around whether the phrase "When I get out of here [a police cell] I'm going to do something about this" constituted a reiteration of an earlier (but unrecorded) verbal threat to shoot a legal official. The defendant vociferously denied that the utterance, which was recorded by a CCTV camera in the cell and which was witnessed by the custody officer who gave evidence against the defendant, was a threat of any kind. Indeed, it is quite easy to imagine that the remark could perhaps have been a statement of the speaker's intent to lodge a formal complaint, or to write to his Member of Parliament. The custody officer was, however, adamant that because he (the officer) interpreted the remark as a reiteration of a threat it *was* therefore a threat, an assertion he defended by stating his impression that the defendant's behaviour at the time was agitated and belligerent, and that the tone of voice he used sounded aggressive.

It is rather obvious that what we might call the speaker's 'tone of voice' may play a key role in how ostensibly neutral utterances can be interpretable as threats to harm the addressee or a third person. Yet to the best of our knowledge no previous research whatever has been done on what we might call 'the phonetics of threat': the ways in which speakers convey menace, ill-will, intention to harm, and so forth, using speech cues that listeners interpret in the expected way. This gap in the research literature is surprising, given how frequently samples of allegedly threatening speech that are obtained via interception of telephone calls or from covertly-recorded conversations are encountered in forensic speech analysis casework. There is a considerable literature on the phonetic properties of other sorts of emotional speech, and that on 'angry' speech (Gobl & Ní Chasaide 2003; Polzehl *et al*. 2011; Xu & Kelly 2010) is probably most relevant here. However, we must avoid conflating angry speech with threatening speech. Wishing to threaten someone does not presuppose that the threatener is angry with the recipient. We should also distinguish between different sorts of anger (say, what are popularly known as 'hot' and 'cold' anger). However, it seems reasonable to expect some overlap in the clusters of phonetic cues that speakers might exhibit when angry and when communicating a threat, and we must allow in any event for the fact that listeners who believe that a threat is being made by a speaker might assume that the speaker is also angry; this heightened emotional state might be thought to be the reason for the threat to be made in the first place. All the same, in the perception experiment we describe in subsequent sections we aimed to keep the two areas distinct from one another, because while being angry with

another person is not an offence, threatening to harm him or her certainly is. In the next section we consider verbal threats from a legal perspective.

2. Legal aspects

So as to contextualise the experiment described below, it is important to consider further the legal implications of the use of verbal behaviour classifiable as 'threatening speech'. It should be stressed from the outset that for an offence to have been committed it is not necessary that the speaker *have* the capacity to carry out the actions that he or she implies or explicitly identifies, nor must it be demonstrated that the speaker has any *actual* intention of carrying them out. People very frequently make empty threats towards one another, after all (Shuy 1993). It is the act of producing a statement which leads an individual to believe that the words used encode a plausible intention to cause harm that is considered in and of itself to be damaging to the victim.

In England and Wales, the Public Order Act (1986) provides definitions and guidance relating to the status of the verbal threat as a criminal offence. In particular, Section 4.1 of Chapter 64 of the Act ('Fear or provocation of violence') specifies the following:

> (1) A person is guilty of an offence if he—
>
> (a) uses towards another person threatening, abusive or insulting words or behaviour,
>
> or
>
> (b) distributes or displays to another person any writing, sign or other visible representation which is threatening, abusive or insulting,
>
> with intent to cause that person to believe that immediate unlawful violence will be used against him or another by any person, or to provoke the immediate use of unlawful violence by that person or another, or whereby that person is likely to believe that such violence will be used or it is likely that such violence will be provoked. (Public Order Act 1986, Ch. 64, Section 4.1).

The subsequent section relating to 'Intentional harassment, alarm or distress' (4A) states that an offence has been committed when, 'with intent to cause a person harassment, alarm or distress', the offender 'uses threatening, abusive or insulting words or behaviour… thereby causing that or another person harassment, alarm or distress' (Public Order Act 1986, Ch. 64, Section 4A).

Demonstrating that there was intent on the speaker's part to cause psychological or physical harm to another person when the threat was made is crucial if the utterance is to qualify as a criminal offence under Section 4 or 4A ('A person is guilty of an offence under section 4 only if he intends his words or behaviour, or the writing, sign or other visible representation, to be threatening, abusive or insulting, or is aware that it may be threatening, abusive or insulting' (Public Order Act 1986, Ch. 64, Section 6.3). However, Section 5 of Chapter 64 of the Act also allows for the fact that the accused may not necessarily have directed the offending language at the person who believes the speaker to intend him or her harassment, alarm or distress: the linguistic material in question simply has to be produced within the hearing or sight of the recipient.

We cannot engage here very fully with the ways in which the wording of these sections of the Act might be interpreted generally, let alone in any specific case. These are tasks best handled by specialists in legal language, and linguists who concern themselves with the

pragmatic functions of different sorts of speech acts, including forensic linguists (see further Austin 1962; Shavell 1993; Rothchild 1998; Searle 1969, 1975, 1979; Smith 2008; Storey 1995; Salgueiro 2010; Tiersma & Solan 2012). However, we deal briefly below with some of the principal themes of earlier research on threatening language, which has largely dwelt upon discussion of what lends a linguistic construction the status of a threat and whether it is to be treated as a 'legal' threat versus an illegal one. Not all utterances that are classifiable as threats violate any law (e.g. 'if you don't pass your exam, then I'll remove you from this course of study'), but they still count as threats insofar as they are statements formulated in such a way as to imply that the person uttering them has at least the will, and perhaps also the capacity, to bring about a state of the world that is unfavourable to the recipient (Fraser 1976). Fraser also contends that when a statement is considered by the recipient to be unfavourable to him or her (e.g. to that person's health or safety) it thereby acquires the status of a 'successful' threat, irrespective of whether the person who produced the communication is the same individual who intends to carry out the unfavourable act.

In spite of the lucidity of the definitions of threats and threatening verbalisations provided by Fraser and other researchers working in this area, there are still extensive grey areas in relation to what kinds of words or speech acts might be construed as threatening, abusive, insulting, harassing, alarming or distressing. In the real world, it seems safe to say that direct, explicit threats produced 'performatively' (e.g. 'I'm threatening you with violence') will be hugely outnumbered by those worded less directly. Conversely, utterances which look like threats may not be intended as such. An ostensibly threateningly-worded utterance such as 'I'm going to kill you' might on the one hand be used as a playful rebuke between friends during a board game, but under other circumstances (a whispered telephone call received in the dead of night, for example) it could cause the person hearing it feelings of intense fear, dread, and distress. Though the same words in the same order are being used, the circumstances in which they are uttered are pivotal; as mentioned earlier, the offence is committed when justifiable alarm, distress or fear for personal safety is experienced on the part of the addressee. Decisions about whether an utterance constitutes a threat or not must therefore rely to a considerable degree on the ability of the recipient of the threat, a police officer, or a legal professional to judge whether any intent to harm was insinuated by the speaker (Gales 2012; Smith 2008). Whether the speaker is in a position to carry out the action(s) implied by the threat – i.e., whether Austin's 'ability condition' is satisfied (Austin 1962) – is also relevant.

It can also be difficult to distinguish objectively between what counts as a threat versus an insult or verbal abuse, but under some circumstances the distinction is unimportant: the target of insulting language may feel very threatened by it, even if the abuser has not used words or phrases that would usually be interpreted to denote desire or intent to cause the recipient harm. As before, insulting or abusive language serves as a threat if it can be demonstrated that there was intent on the speaker's part to cause alarm or emotional distress to another person. Abusive or insulting terms accompanying a more explicit threat may also affect how severely or seriously the threat is perceived.

Individuals making verbal threats towards others often wish to hide their identities by delivering their messages over the telephone. They may know, however, that there is a possibility that their speech is being recorded, and they may also be aware of the penalties associated with being found guilty of threatening behaviour. It is therefore in the interests of the threatening party to phrase the threat as indirectly as possible, by trying to strike a balance between framing the threatening message in such a way that on one level it looks innocuous, while on another level the recipient will interpret the threat as the speaker intended it. This leaves the threatener, if challenged, with the escape route of claiming that the recipient

simply misinterpreted the words that were used, as per the scenario in the case described in Section 1. A veiled threat may in some circumstances therefore be the most useful type from the threatener's point of view: written down, the message might look neutral or harmless, but if delivered in spoken form in the right way, it might have precisely the intended impact on the listener.

In the following sections, we consider what this 'right way' might be. We do not attempt in the present paper to say what kinds of phonetic strategies a speaker might use when trying to make an innocuously-worded phrase sound threatening to a listener, nor do we seek to identify specific cues in the talker's speech that listeners might respond to such that they infer a greater or lesser level of threat or intent to carry out some harmful action. Rather, we draw our initial focus on whether listeners act consistently in how they assign threat and intent ratings to a set of controlled utterances, in line with whether the utterances were produced by talkers in a neutral way versus a way that was intended to sound threatening.

3. Background to the experiment

To help us to gain some more clarity with respect to the above issue, we use an experimental approach to examine listeners' reactions to sentences read aloud by talkers who were instructed to adopt a threatening tone of voice while reading. We label these productions 'induced-threat' sentences. The sentences were worded so as to seem innocuous if taken at face value, such that if listeners were to infer threat from them when hearing recordings of them being spoken aloud, their inferences would, in principle, be drawn on the basis of the way the sentences were spoken, instead of the semantic properties of the words and phrases of which the sentences were composed.

We are not seeking here to investigate the acoustic properties of these phonetic cues themselves, though we recognise that pursuing this objective would be a useful contribution to our knowledge in this area. Rather, our experiment is designed to gather some initial observations about listeners' judgments of speech samples produced by talkers who had been asked to adopt a 'threatening tone of voice', in respect of how listeners chose to rate the samples for 'level of threat' and perceptions of what we might informally call the speaker's 'sincerity', i.e. the extent to which listeners believe that the speaker might actually carry out some (unspecified) harmful action.[1] Since the classification of a verbal threat as an offence in UK and US legislation depends so heavily upon the intimidatory effect that the perpetrator's utterance has on the listener, we thought it appropriate to focus principally upon how listeners react to verbal stimuli designed as threats. That is, as experimenters we did not wish to assume that our speakers' utterances would be heard as threatening by our impartial listeners simply because we had instructed the speakers to adopt a threatening tone when reading them out loud; the key thing was to establish how the listeners reacted to the samples. Moreover, as a way of drawing a sharper focus on how the sentences were delivered, we circumvented any interference from semantic content by presenting listeners with stimuli spoken in four foreign languages which were unknown to them as well as stimuli spoken in English. By including utterances in multiple languages our results might lend support to claims made by, for

[1] It should be stressed that in this research we are not seeking to attempt to identify phonetic cues to sincerity, or anything of that kind; we are conscious that unless we are careful to make this point explicit we might run the risk of being seen to contravene item 9 of the 2004 Code of Practice of the International Association for Forensic Phonetics and Acoustics (www.iafpa.net/code), viz., 'Members should not attempt to do psychological profiles or assessments of the sincerity of speakers'. Our goal here is simply to report our listeners' subjective judgments in a way that is relevant to the questions at hand.

example, Bolinger (1989), Gussenhoven (2004) or Xu *et al*. (2013) about the universality of associations between prosodic cues and social/emotional affect.

We also set out to obtain responses from both men and women so as to allow us to test whether there is a consistent effect for listener sex, in line with results reported in earlier literature which indicate that men and women process and respond to emotional speech differently (Bonebright *et al*. 1996; Schirmer *et al*. 2004, 2005; Schirmer & Kotz 2003, 2006). Though threatening speech has not previously been investigated, it is conceivable that sensitivity to threat cues in speech might be higher in one sex than the other. While we would hesitate to classify threatening speech as a form of 'emotional speech' in any straightforward way, sex-dependent differences between listeners in the present study would accord with the findings of existing research on vocal correlates of emotion, which has often included anger among the set of emotions tested for. For instance, Bonebright *et al*. (1996), who investigated sex differences in the perception of vocal affect using recordings of acted portrayals of fear, anger, happiness, sadness and a 'neutral' emotional state, found that 'male actors were better at portraying anger than females, and… anger was the only emotion which males were better at identifying than females' (1996: 440). They account for this asymmetry by pointing to differences in the socialisation of children, asserting that 'males are encouraged to express anger and to control the display of other emotional states', and citing anthropological research on hunter-gatherer societies in which 'males were taught to express threats and show no fear' (Miller 1928, cited in Bonebright *et al*. 1996: 441).

Neurophysiological studies reviewed by Schirmer & Kotz (2006) provide evidence of sex differences in the degree of automaticity with which listeners respond to 'emotional-prosodic information' in speech processing, leading them to conclude that '[b]ased on these findings one can assume that emotional expressions and thus social interactions are of greater significance to women than to men' (2005: 27), and that women's 'greater interest in affiliation might make [them] more dependent than men on the emotional state of others, which might in turn enhance their perception of emotional cues' (2005: 27). On the behavioural side, it is argued by Lochman *et al*. (2006) that females exhibit more of a tendency than do males towards 'relational aggression', defined as '[harming] others not through the use of physical violence, but instead through acts that damage peer relationships or threaten to do so' (2006: 116). Thus, we might reasonably predict that in the present experiment the responses gathered from men and women will differ, and will do so in ways that could be accounted for in terms of gender socialisation factors (e.g. that from a young age boys engage in physically threatening behaviour more often than girls do and are thus more highly sensitised to the signals associated with an impending physical attack, or that girls are more attuned to emotional-prosodic information than boys are, and can read these cues more accurately).

With these considerations in mind, we proceed to a description of the experiment itself, in which the above variables and others we controlled for are discussed in more detail.

4. Methods

4.1. Materials

The recordings used for the experiment were collected from 12 male speakers (8 native British English speakers and 4 native or near-native speakers of Norwegian, Finland-Swedish, Hebrew, and Arabic). The sentences read by the speakers of the languages other

than English were direct translations of the English versions (see below). No attempt was made to control for habitual voice quality across the talkers.

Owing to the heavy demands that having to listen to a full set of stimuli would make on our participants, we did not have every speaker perform every reading task. Table 1 shows which reading tasks were carried out by which speaker. The 'threatening script' mentioned in the rightmost columns in Table 1 is described in the text below the table.

Speaker	**Language**	**Task**			
		A	**B**	**C**	**D**
		Say the phrase 'I know where you live' (or foreign language equivalent)	Say the phrase 'I know where you live' (or foreign language equivalent) in a threatening way	Read extract from a non-threatening script (or foreign language equivalent)	Read extract from threatening script (or foreign language equivalent)
1	English	x	x		
2	English	x	x		
3	English	x	x		
4	English	x	x		
5	English			x	x
6	English			x	x
7	English			x	x
8	English			x	x
9	Finland-Swedish	x	x		
10	Norwegian	x	x		
11	Hebrew			x	x
12	Arabic			x	x

Table 1: Reading tasks performed by each of 12 speakers (8 English speakers, 4 speakers of other languages).

The phrase 'I know where you live' (and its translated equivalents) was chosen as although it would be hard to argue that it intrinsically constitutes a threat *per se*, it nevertheless has ample potential as an indirect threat, under certain circumstances. Listening to a speaker reading the sentence on its own in a 'neutral' way (Task A) ought not to suggest to listeners that they are hearing the speaker threatening someone. A higher threat rating for the same isolated sentence read in a threatening way (Task B) would therefore indicate that participants had responded as expected to changes in the speaker's articulation of the utterance. Similarly, a consistent difference between listener's ratings for the neutral and

induced-threat versions of the test sentence in some or all of the other four languages would suggest that it is changes in the phonetic form – i.e., the segmental and prosodic features of the utterances – that are responsible for changes in the way the utterances are rated.

As shown in Table 1, half of the speakers (four of the English speakers, the Finland-Swedish speaker and the Norwegian speaker) were given the instructions *Say the phrase 'I know where you live'* and *Say the phrase 'I know where you live' in a threatening way*, in that order. These speakers read from 'cold': they were offered no other explicit directions, and were free to read the phrases as they saw fit. They had not been informed in advance that the purpose of the study was to investigate perceptions of threatening speech, so had no reason to assume that the first version of the phrase should be read in anything other than a neutral or default manner. Giving more guidance on what was meant by 'a threatening way' might also have been leading, and it was left to the speaker to decide how to produce the phrase as though intended as a threat.

The other six speakers (i.e. the other four English speakers, the Hebrew speaker and the Arabic speaker) were presented in Tasks C and D with an imaginary backstory and then asked to read aloud the sentences (1)-(3) in each task, as follows:

Task C (neutral)
Answerphone message/voicemail in which you and some friends are inviting another friend for an afternoon at the park. The friend's house is situated on the way there, so you offer him/her a lift in your car.

(1) We're going to the park for a picnic, should we come and get you? [pause]
(2) I know where you live [pause]
(3) So we'll come and pick you up

Task D (induced-threat)
Answerphone message/voicemail threatening someone who owes you a large sum of money, and refuses to pay it back to you.

(1) You better watch out [pause]
(2) I know where you live [pause]
(3) If you don't pay me back, I'll know where to find you

For the foreign language samples, the scripts were translated by native or highly-proficient users of the language in question. The three sentences in Task C were designed to be produced by the speaker in a non-threatening manner, without any suggestion of animosity. By contrast, while there is nothing in the wording of Task D that directly encodes a threat of violence, the speaker had been encouraged by the backstory to develop a sense of being justified in leaving the threatening message, as though the money had been 'stolen' from him. At no point was the reader of Task D instructed to try to sound angry. Both scenarios were designed to represent realistic situations that participants could imagine themselves being involved in.

The sentence *I know where you live*, and its foreign language equivalents, were extracted from the Tasks C and D readings, such that listeners would hear the sentence in isolation from its accompanying context. Extraction of the sentence was facilitated by the pauses left by speakers on either side of it.

The purpose of implementing these two differing strategies for obtaining the speech samples was to establish whether embedding the target sentence in a script would result in more consistent perceptions of threat in the induced-threat sentences (Task D) by listeners than in the parallel case where the sentences were read without an accompanying context. Providing speakers with scenarios in which the sentence might plausibly be spoken under real circumstances would, it was thought, enable them to draw upon their own experiences and perhaps also media depictions of threatening behaviour, and to produce the utterances in a more naturalistic way.

The recordings were made using the audio capture capabilities of a Toshiba Satellite C850–10C personal computer, with the speaker positioned approximately 60cm from the microphone. Because recordings in forensic casework are generally of poor technical quality it was in the interests of realism not considered a priority to obtain the best possible quality recordings, but all the same the recording sessions took place in environments with low levels of background noise. The recorded sentences were then embedded in an online questionnaire via which participants' responses were logged.

4.2. Participants

A total of 30 participants (13 male, 17 female) were approached via online social media to complete an online survey. All were native English speakers aged 18 years or over with little or no proficiency in a language other than English, and no reported hearing difficulties. The majority of participants had no formal linguistic or phonetic training.

4.3. Procedure

Participants were asked to complete a questionnaire in which they would rate the following: (a) the degree to which a sample of speech seemed threatening (the 'threat' rating), and (b) the degree to which they believed the speaker intended harm (the 'intent' rating). They responded using an on-screen horizontal sliding scale running between 0 (for parameter (a), 0 represented 'not at all a threat') and 100 ('very likely a threat')), allowing participants a degree of precision in their responses. By default, the slider was set at 0 when presented to the participant. Each listener rated 24 sentences apiece (16 English, 8 foreign language), yielding responses for a total of 720 sentences.

Participants were also offered the opportunity to give a description of the voice they had heard. As the majority of participants had had no previous linguistic training, these responses were expected to resemble typical lay-listener responses. For the foreign language utterances, participants were also asked whether they knew what language had been spoken, and if so, to name it.

Foreign language samples were played in randomised order before listeners heard any English tokens (also randomised), in an attempt to ensure that participants' responses to the foreign language stimuli were not biased by their having heard the English material beforehand.

The results collected from the online survey were subjected to ANOVA testing in order to establish the strength of the effects of the independent variables (e.g. language spoken) on the threat and intent ratings. Pearson's product-moment correlation coefficients were also calculated to gauge whether listeners' threat and intent scores were correlated. Although it is recognised that individual respondents may have used the scale in different ways (e.g. some

may have avoided giving responses at the extreme ends of the scale), we did not attempt to normalise the responses in any way, for example by expressing listeners' individual response values relative to the range of values each listener used. Correcting for individual variation in this way may yield a different picture and it is worth considering applying such a method in future research on this topic.

5. Research hypotheses

The hypotheses we are seeking to test in the current research are as follows:

1. Induced-threat utterances will be identified as such – that is, listeners' ratings for these sentences on the threat scale (and probably also the intent scale) will be higher than those for their 'neutral' counterparts.
2. Male and female listeners will rate the utterances differently, in line with previous experimental literature showing gender differences in the perception of emotional speech.
3. Utterances produced by speakers who were provided with the scenario plus three-sentence script to read from will yield consistently higher threat and intent ratings than will induced-threat utterances that were read from cold, i.e. without any supporting backstory/context.
4. Foreign language utterances will be rated by listeners for threat and intent in the same way as their English equivalents are, because the set of phonetic cues used to express threat will be approximately the same across the five languages. Put another way, the linguistic content of the utterance being rated for threat and intent is immaterial: the listener will rate the utterance on the two scales solely on the basis of its phonetic form.
5. The threat and intent ratings will be correlated, because the more threatening a utterance is perceived to be, the more listeners will believe that the speaker intends to carry out the threat.

In the following sections we discuss the level of support given to each of these hypotheses by the data collected for the purpose.

6. Results

Data collected from the online survey were separated into two groups: *English* and *foreign language*. We look first at the differences that can be observed among the threat and intent ratings for the neutral and induced-threat utterances in the *English* condition.

Figure 1 summarises the results for the English utterances where listeners' responses on the threat and intent scales are presented by listener sex and by production mode ('neutral' – Tasks A and C readings pooled – versus 'induced-threat' (Tasks B and D) sentences pooled). It can immediately be seen that the average scores are all rather low on the 0–100 scale, showing that listeners were, overall, not strongly inclined to hear the utterances as threatening. This is to be expected for the neutral utterances produced for Tasks A and C, though more surprisingly they were still rated some way above zero ('not at all a threat'). Another conspicuous trend is that, in both the neutral and the induced-threat conditions, male listeners awarded higher threat and intent ratings than did females, more particularly in the induced-threat condition. This difference is supported by the result of ANOVA testing, which reveals there to be a highly significant effect for listener sex ($F = 20.289$; $p < .0001$). In all cases, the average scores for threat are marginally higher than those for intent, except in the case of the male listeners' responses to the induced-threat utterances, for which the average

level of perceived threat is more than 6% higher than that for intent. The male listeners' average score for threat among the English induced-threat utterances attains nearly 50%, indicating that the male listeners appear to have perceived these utterances to be considerably more threatening than did the female listeners ($F = 55.527$; $p < .0001$). Listener sex is shown also to have a strong effect upon intent scores ($F = 10.642$; $p = .001$), and for intent the effect of production mode (neutral vs. induced-threat) is also highly significant ($F = 31.381$; $p < .0001$).

6.1 English stimuli condition

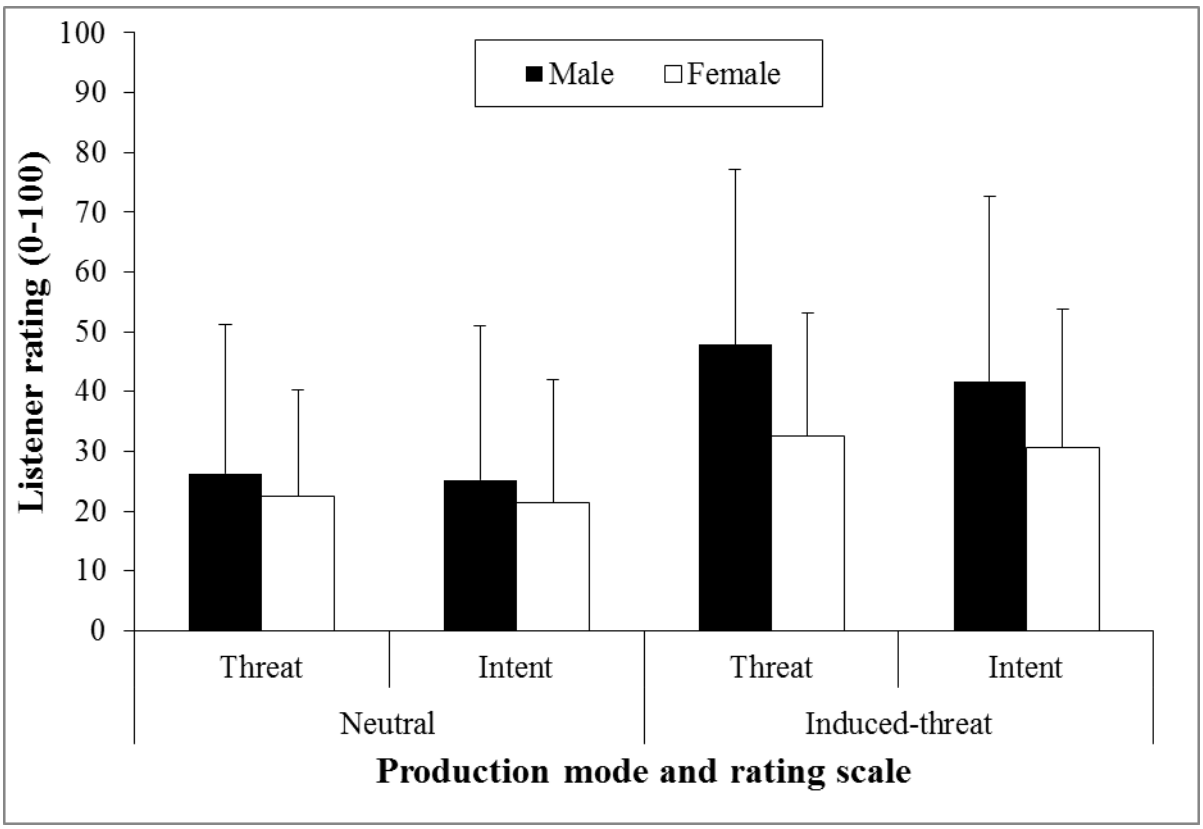

Figure 1: Threat and intent ratings (%) for Task A-D sentences in the English condition, by listener sex. Error bars represent one standard deviation.

The error bars in Figure 1 indicate that there is considerable variation in the individual scores, and a good deal of overlap between the ratings given by the men and the women is also evident. However, two of our hypotheses have been supported by the English data: it appears that listeners can distinguish neutral from induced-threat utterances, and that male listeners consistently give higher ratings (albeit often just marginally higher ones) for both threat and intent than females do, even if the sentences they were rating are not intended to sound threatening.

6.1.1 The influence of scripted context

In this section we investigate whether listeners gave different ratings on the threat scale for the same target sentence *I know where you live* when the sentence is read cold (Tasks A and

B) versus when it is embedded in a three-sentence script accompanied by a contextualising backstory (Tasks C and D; see Section 4.1).

Four of the eight English-speaking readers were given the Task C and D materials to prepare for the recordings. It was thought that being given brief details of a plausible scenario which might prompt a person to say the target sentence would help the speaker to produce the neutral and induced-threat utterances more 'naturally'.

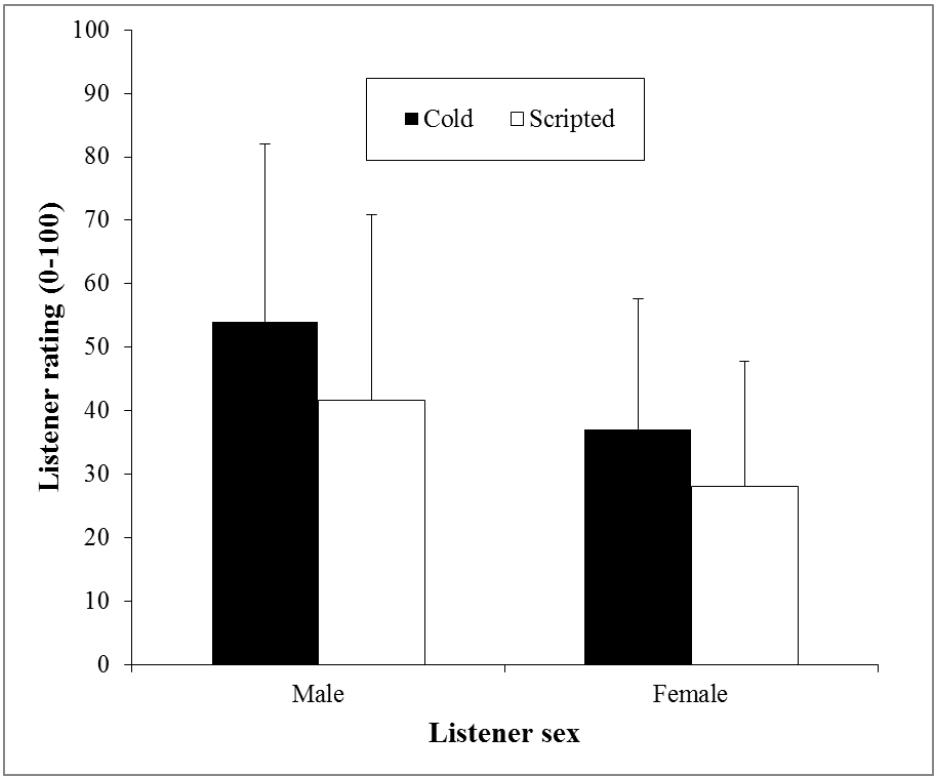

Figure 2: Average perceived threat ratings (%) for the induced-threat target sentence spoken 'cold' (Task B) and in a scripted context (Task D), by listener sex (English condition). Error bars represent one standard deviation.

An ANOVA test on the listeners' threat ratings of the English sentences showed there to be a significant effect for scripted context versus cold reading ($F(8.541) = 1$, $p < .004$). At first glance, this might appear to bear out our third hypothesis. However, the sentences produced by speakers who performed Task B (cold reading, induced-threat) were on average rated higher on the threat scale than those spoken by Task D readers (scripted context, induced-threat), as can be seen in Figure 2.

Owing to the subjective nature of the listening task we cannot identify an obvious explanation for why the pattern in Figure 2 is the reverse of that expected, but it is possible that when reading the sentence in the absence of accompanying context the Task B speakers exaggerated the phonetic cues they used to indicate that the sentence was intended as a threat, whereas in Task D they relied more upon the context to communicate to the recipient that *I know where you live* was meant to represent an indirect threat to harm him or her (note that speakers were not told that the second of the three sentences would be excised from the recording and played to listeners in isolation, so may have assumed that listeners would hear all three sentences). It is also possible that the readers felt more constrained by the context in

Task D, such that they felt less inclined to adopt what they would consider a threatening tone of voice.

When the data are viewed as shown in Figure 2 we can again see that male listeners tended to give higher threat ratings than did female listeners. For the Task B 'cold' readings, the male listeners' average score exceeded the 50% mark, indicating that they believed the utterances to merit ratings closer to the 'very likely to be a threat' end of the scale than to the 'not at all a threat' end.

We next evaluate the listeners' judgments of the foreign language sentences, to see whether the data give support to hypothesis (4), which asserts that in the current experiment the language in which translations of the neutral and induced-threat utterances are delivered will have no influence on listeners' threat and intent ratings.

6.2 Foreign language condition

The foreign language judgments were initially aggregated into a single group on the grounds that as the listeners had had little foreign language experience, there was no *a priori* reason to assume that the foreign language sentences would be treated differently (other than that they were spoken by different individuals). Therefore, they were considered collectively for the purposes of comparison with the English samples.

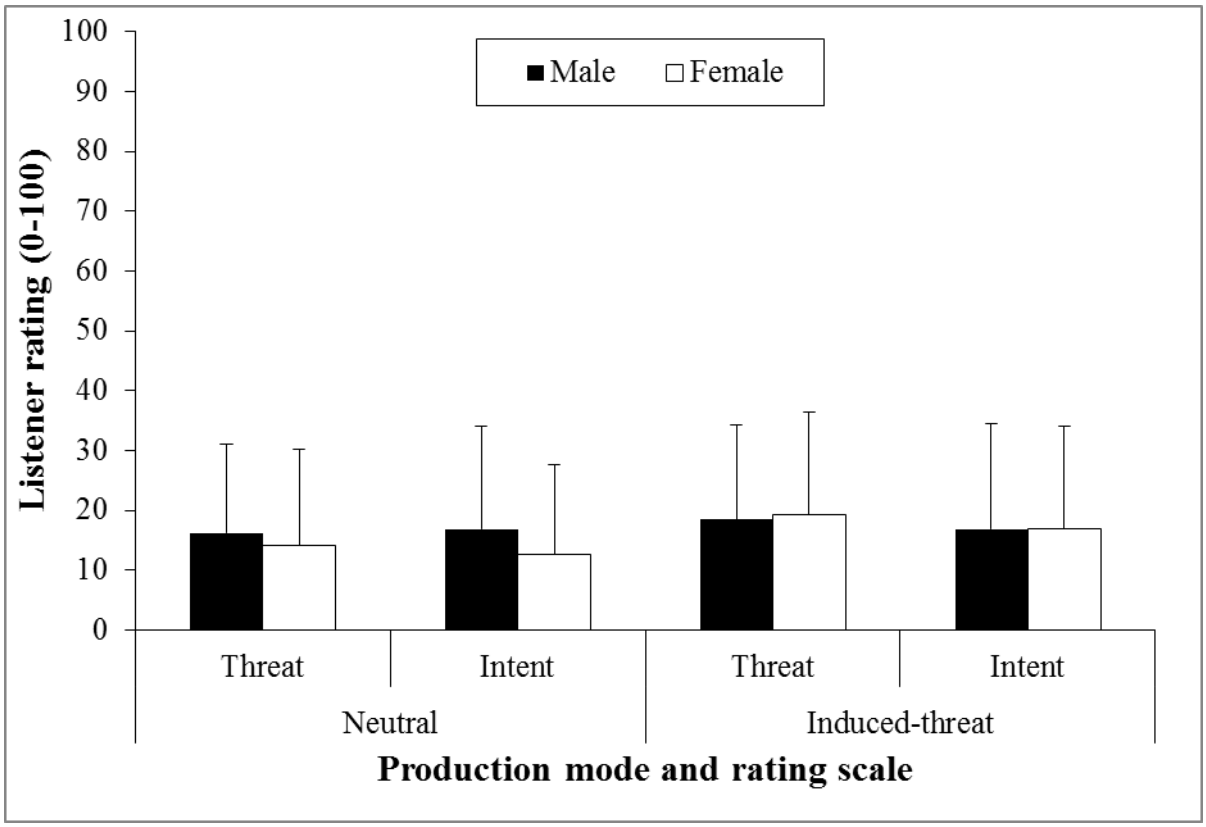

Figure 3: Threat and intent ratings (%) for Task A-D sentences in the foreign language condition (results for all 4 languages pooled), by listener sex. Error bars represent one standard deviation.

Figure 3 shows there to be practically no difference in the listeners' threat and intent ratings across the neutral and induced-threat conditions. The ratings are low throughout; no average rating exceeds 20%, and the error bars imply that even the highest ratings for the induced-threat utterances did not often surpass 50% on either the threat or the intent scales. Listeners, in other words, apparently did not hear much difference between the neutral and the induced threat utterances, though such differences as there are go in the expected direction for perceived threat in the induced-threat condition. Males give slightly higher average ratings for females in the neutral condition, but the reverse is true in the induced-threat condition.

6.2.1 Influence of scripted context

As we saw in Figure 3, the threat ratings for the induced-threat utterances (those spoken in Tasks B and D) in the foreign languages are generally very low. It appears to make practically no difference to listeners' ratings whether the test sentences are read 'cold' in isolation or embedded in a script when read aloud (Figure 4; note that the Task B readings were for Norwegian and Finland-Swedish only, while those for Task D were in Hebrew and Arabic only).

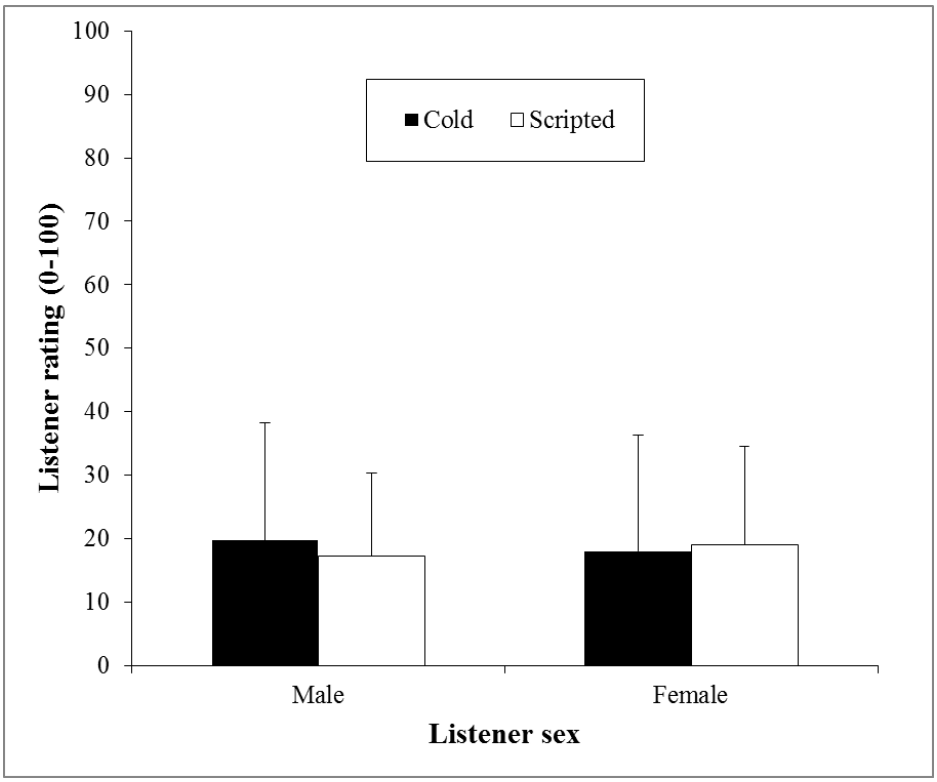

Figure 4. Average perceived threat ratings (%) for the induced-threat target sentences spoken 'cold' (Task B; Norwegian and Finland-Swedish only) and in scripted contexts (Task D; Arabic and Hebrew only), by listener sex.

The pattern for the male listeners, whereby the cold reading is rated higher for perceived threat than its scripted equivalent, recalls that seen for the English language data. However, the difference is only marginal. That for the female listeners goes in the opposite (but initially expected) direction. Neither difference is found to be statistically significant, and the male listeners' average scores are not significantly different from the females' scores either.

The general tendency overall, then, is for listeners to treat the foreign language utterances as lacking in perceived threat. Perceived intent is correspondingly low. The latter is understandable, because it is hard to imagine a scenario in which intent to harm would consistently be rated higher than perceived threat for any given utterance or set of utterances. It seems probable, nonetheless, that scores on the two scales will be correlated: an utterance which is perceived to be low for threat will tend to be rated low for intent, while utterances judged to sound highly threatening to listeners will be more likely to yield high ratings for intent. We examine the evidence for this assumption – our hypothesis (5) – in the following section.

6.3 Correlations between threat and intent

There was found to be a very strong positive correlation between the scores on the threat and intent scales overall (Pearson's $r = 0.883$, df = 718, $p < .0001$),[2] with the great majority of threat ratings (73.9%) being either equal to or larger than the corresponding intent rating for the same utterance rated by the same listener.

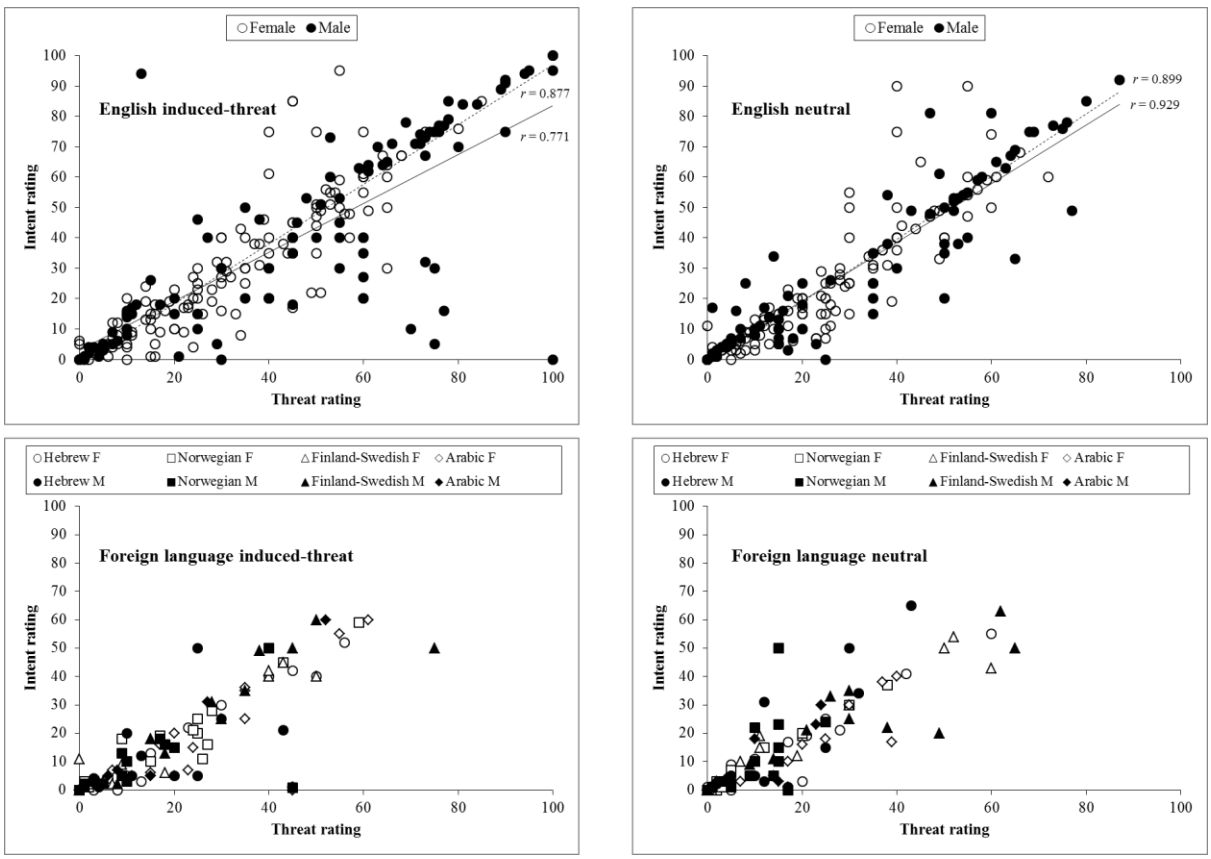

Figure 5. Correlations of threat and intent ratings for induced-threat and neutral utterances in English and four foreign languages, by listener sex. Solid trendlines in the English condition are for male listeners, dashed lines are for female listeners.

Figure 5 shows all 720 individual pairs of ratings divided by language (English in the top two panes, and the foreign languages in the lower two panes), by listener sex (where male

[2] Correlation coefficients were calculated using Wessa (2012).

listeners are represented by black markers, and female listeners by white ones). The two left-hand panes represent the results for the induced-threat conditions, and those on the right show the data in the neutral conditions. The results for the four foreign languages are displayed separately in the lower two panes. Regression trendlines and correlation coefficients (Pearson's *r*) are shown in the upper two panes (English conditions) for the male and the female listeners separately.

Positive correlations are evident in all of the data sets represented in Figure 5. For the English data (top panes) the match between the ratings for threat and intent appears to be tightest for high and low values, with more variability in the middle of the range. In the English induced-threat condition, it tends to be male listeners who award high ratings on both scales. For given utterances in this condition, male listeners tend to give higher ratings for threat than for intent in the mid-range area, while for female listeners it is the other way round. In the English neutral condition, the scores given tend to fall lower on each scale (cf. Figure 1), and there is greater agreement within each listener group, as reflected by the higher correlation coefficients. The correlations between the sex-group scores for threat and intent in both the induced-threat and neutral conditions were for English very highly significant (induced-threat, male listeners, $r = 0.771$, df = 102, $p < .0001$; induced-threat, female listeners, $r = 0.877$, df = 134, $p < .0001$; neutral, male listeners, $r = 0.929$, df = 102, $p < .0001$; neutral, female listeners, $r = 0.899$, df = 134; $p < .0001$).

Similarly, when the foreign language data are pooled the correlations between threat and intent scores for the male and female listeners groups were very highly significant (induced-threat, male listeners, $r = 0.858$, df = 50, $p < .0001$; induced-threat, female listeners, $r = 0.96$, df = 66, $p < .0001$; neutral, male listeners, $r = 0.79$, df = 50, $p < .0001$; neutral, female listeners, $r = 0.959$, df = 66, $p < .0001$). As we saw in Figure 3, the foreign language sentences were overall rated lower than their English counterparts for both threat and intent, but for individual utterances listeners still consistently gave scores on each scale that agreed with one another quite closely. Owing to a lack of space it is not possible to show regression lines for each of the languages in the plots themselves, but the correlations for all language/listener sex group/production mode pairings of threat and intent ratings (e.g. female listeners' ratings for the Norwegian neutral utterances) were shown to achieve significance at the 5% level. The loosest correlation was for the male listeners' threat and intent ratings for the Hebrew induced-threat utterances ($r = 0.578$, df = 11, $p = .039$). The responses for this particular condition show that listeners mostly rated the utterances low for threat, and that in the main the listeners' ratings on the threat scale were higher than those on the intent scale. While listeners tended to perceive threat in these utterances, then, we might say that they did not perceive them to be delivered with much conviction.

7. Discussion

The results discussed in the preceding section bear out the majority of our initial research hypotheses. The first of these, which proposed that listeners will show that they can distinguish induced-threat utterances from their neutral counterparts by assigning utterances of the former type higher ratings for threat and intent than those they gave to the latter, is clearly supported by the English data. The differences in the average scores for the English data are, however, not as big as we might have expected, for two reasons: the neutral utterances were perceived as threatening and intentful to a greater degree than anticipated, and the induced-threat utterances were apparently not perceived to be as threatening and intentful as we had thought they would be. Male listeners returned higher scores on both

scales than female ones did, on average, particularly for those sentences produced in the induced-threat mode.

The corresponding scores for the foreign language utterances, by contrast, exhibit practically no differences of this kind. The induced-threat utterances are rated only fractionally higher on the two scales than the neutral ones, and the scores are in any case very low. The responses for the two listener sex groups are closely comparable with one another. It appears, then, that English-speaking listeners hearing target sentences spoken in an unfamiliar language cannot decide whether they are hearing a neutral or a threatening utterance. The average scores on the two scales for utterances in both production modes are in the region of 15–20%, so it is not the case that they are heard as completely lacking in threat or intent, but the scores may just reflect the fact that listeners were aware of the nature of the experiment and so were more strongly inclined to assign threat and intent ratings to the stimuli than they might have been had it not been suggested to them that the criteria of interest related to the making of verbal threats.

Our second hypothesis, that men and women would respond differently to the stimuli when rating the utterances for perceived threat and intent, is also supported by the English data, but not to anything like the same degree by the foreign language data. These patterns are clearly visible in Figures 1–4, and to a lesser extent in Figure 5. We might tentatively conclude that male listeners are more sensitive than women to threatening content in speech produced by other males, in line with the arguments about sex differences in the perception of emotional speech put forward by Bonebright *et al*. (1996), Schirmer & Kotz (2005), Lochman *et al*. (2006), and others. Males are more likely than females to be the targets of physical attacks (Archer 2004) and on balance it is therefore in their interests to be alert, whether consciously or subconsciously, to signals of hostility or potential harm expressed by others, whether verbally or otherwise. It is notable that following the recording session, several of the speakers recruited for that part of the experiment reported that they had produced the target utterances as though they were addressing another man. Whether the phonetic strategies they adopted to signal threat would have differed had they been asked specifically to speak as though the target of the threat were a woman rather than a man would present an interesting line of inquiry. Future research in this area may also cast light on whether any parallel disparity in sensitivity to same-sex utterances exists, whereby female listeners might respond more strongly than male ones to sentences produced in the two modes by female talkers.

When the threat scores were partitioned according to whether the English sentence *I know where you live* or its foreign language equivalents (induced-threat mode only) were produced in isolation (‘cold’) or embedded in a script with an accompanying backstory (‘scripted’), the pattern predicted by hypothesis (3) was not observed. Rather, we found that – with the exception of the female listeners’ ratings for the pooled foreign language utterances – higher threat scores were returned for the unscripted sentences. This suggests that in the ‘cold’ readings the English speakers may have produced the target sentence using a greater number of phonetic cues that the listeners associated with threatening speech, and/or using more extreme deviations from their ‘neutral’ pronunciation setting.[3] They may have believed that the wording of the context sentences flanking the target sentence in the induced-threat mode (Task D) would make it clear that the target sentence was meant as a threat. They had not been told that the listeners would only hear the target sentence from Task D. Alternatively, the speakers may have chosen to focus more effort on making the first or third sentences of the three in the script sound threatening, such that the second was closer to a neutral reading for that speaker than it would have been when the sentence was read in isolation.

[3] We intend to investigate the nature of the phonetic cues themselves in a future study.

In the English condition, the male listeners returned significantly higher average scores on the threat scale for both cold and scripted induced-threat sentences than did the females. This pattern was not observable in the foreign language condition, however, nor was there any difference between the average scores for the cold and scripted utterances. It should be remembered, on the other hand, that the data for the four foreign languages were pooled, and that averaging the scores in this way might disguise effects that obtained for one language but not others. This question awaits further investigation.

The foregoing discussion should by now have made it apparent that the responses on the threat and intent scales differ markedly according to whether the language heard is English (the listeners' native language) or an unfamiliar foreign language. This finding runs counter to hypothesis (4). It seems that perceiving threat and intent to harm in a talker's speech depends on being able to understand the words being spoken, and is not solely a matter of attuning to a set of phonetic cues that listeners might associate with threatening speech. As before, one might argue that rolling the four foreign languages together in this fashion ignores the differences between the languages and the talkers who spoke in them. In Figure 5, one can see that in the two lower panes (representing the individual responses for foreign language induced-threat and neutral utterances) the Finland-Swedish speaker, denoted by filled and open triangles, has a tendency to be rated relatively high on both the threat and the intent scales. In the same way, we have not in the current paper paid regard to differences between the individual English speakers, although some are consistently perceived to sound more threatening and intentful than others. Confidence in reading aloud for a sound recording, experience with acting, and differing levels of exposure to real and simulated threat in the media (violent films and TV programmes, computer games, etc.) may all have had an influence on the way in which the utterances were produced in the induced-threat mode. Space constraints preclude a more nuanced treatment that would reveal how far the languages chosen and the phonetic behaviour of individual talkers contribute to the differences in listener scores, but we believe that it is important to investigate these issues in subsequent research.

The fifth of our hypotheses, which concerns the correlation between listeners' ratings for threat and intent, is strongly supported by the present data. There is indeed a close relationship between the level of threat listeners report perceiving in the test utterances, and the strength of their belief that the speaker intends to cause the recipient harm. It may be, of course, that the participants did not distinguish very clearly between threat and intent. For a threat to be perceived as such may require an assumption on the hearer's part that the speaker at least potentially wishes to harm the target of the threat, so to that extent one could for a given utterance always reasonably expect to see a value on one scale that is commensurate with a value on the other. And as noted earlier, it would also be peculiar if individual utterances were frequently rated higher for intent than for threat. However, listeners were free to choose values on the two scales independently for each stimulus they heard, and in many cases they exercised this freedom, so it is not the case that the two scales are so inextricably linked to one another that a rating for threat invariably predicts an approximately equivalent rating for intent.

8. Conclusions

In this paper we set out to investigate how English-speaking listeners respond to neutrally-worded test sentences produced by speakers of five different languages (English, Hebrew, Norwegian, Finland-Swedish and Arabic), where the perceptual criteria being tested for are

the level of perceived threat, and the level of intent on the part of the speaker to perform an act that would be harmful to the target of the utterance.

The results show, first and foremost, that the listeners identify utterances designed as threats as such much more accurately if they are spoken in English than in an unfamiliar foreign language. Given that making a threat constitutes an offence, we might conclude that in real-world scenarios we should be especially careful when evaluating witness testimony in which it is claimed that a threat was being made, where the language used is one that is unfamiliar or unknown to the witness.[4] The results of this study also suggest that we should also remain open-minded in respect of the factors that underpin a witness's claim that a threat was made where the language being used is English, in light of the fact that there is so much variability in how the individual English utterances tested here were rated for threat and intent (Figure 5).

The second key finding of the current research is that male and female listeners respond differently to the stimuli, with men tending to perceive a substantially higher level of threat and intent in the English utterances. This listener sex difference is neutralised in the foreign language results. This finding confirms some of the generalisations made in earlier research on sex differences in the perception of vocal expressions of emotion, and it links the perception of threat – which has to our knowledge not been investigated using this kind of experimental approach before – more closely to the perception of related emotions, in particular anger.

From the forensic viewpoint we might also learn usefully from the fact that the perception of threat in a given utterance seems to correlate so tightly with the perception of intent to harm. On the one hand this is not at all surprising, but it is easy to conceive of scenarios in which a defendant accused of having threatened a second party might find it difficult to claim that s/he had only meant the utterance as an empty threat. For example, if jurors are allowed to listen to a recording of the incriminating utterance and believe it to sound highly threatening, they are more likely to perceive it to signal a high level of intent on the speaker's part to harm the addressee, thereby increasing the probability of the imposition of a harsher penalty.

There is of course plentiful scope for expanding and improving upon the design of the current study. Varying the test utterance beyond a single sentence per language would be one way in which to do this. For example, it would be possible to gain a sense of how much the semantic content of the utterance contributes to its perception as a threat, regardless of its phonetic characteristics. It would be uncontroversial to say that a direct threat like 'I'm going to stab you with this knife' could be interpreted as a threat no matter how it was spoken, but a sentence as neutrally worded as 'I put the book on the shelf' would depend on some other source of information to elevate it to the status of a threat. Use of professional actors rather than untrained volunteers might also yield greater consistency between speakers and listeners in terms of how the test utterances were produced and perceived, though there might be a risk that the utterances would be performed in an excessively stylised way that would detract

[4] Interestingly, some of the listeners reported hearing foreign language utterances *as* English ones, when they were asked if they were familiar with the language being used. Eight of 17 listeners believed that the Hebrew utterances were English or 'Scottish', with three transcribing the Hebrew target sentence /ʔani jodeʔa ʔefo ʔata ɡaʁ/ as the English phrase 'And you're there for the girl'. Four listeners heard the Arabic utterance as English. One of these, who had also heard the Hebrew sentence as English, offered the transcription 'An ad for the secretary' for the Arabic sentence /ʔana ʕa:rif inta sa:kin fe:n/).

from their forensic realism (cf. Roberts 2011). The majority of people who threaten one another with harm are not trained actors, after all.

Benefit might also be gained by using foreign languages which are sufficiently familiar to listeners that stereotypes of the language or of people who speak it would come into play. Additionally or alternatively, languages which listeners might have difficulty identifying by their spoken forms alone might be named to one half of the listener group, to see whether knowing the name of the language biases listeners towards higher or lower threat/intent ratings. Showing pictures of individuals from different ethnic or religious groups during exposure to the auditory stimuli might provoke similar changes in threat perception, while a greater range of regional and social accents of English would also no doubt provide insights into how some sociolects are associated more strongly with criminality and other anti-social tendencies (e.g. Dixon *et al*. 2002).

The subjectivity of emotion perception makes it a difficult field in which to draw firm conclusions, but without systematically-collected data on which to ground a more extended programme of research we cannot progress from the situation we are currently in, whereby the classification of utterances as indirect threats in courts of law is entirely a matter of one person's opinion against another. It is important to reiterate that we do not wish to claim that we yet have any objective grasp of what phonetic cues turn an innocuous, neutral statement into an utterance that makes the recipient fear for his or her safety. Rather, we see the results reported in this paper as a first step towards gaining a better, less simplistic understanding of how the interplay of intentions, actions and responses on the part of speakers and listeners conspire to create this particular kind of verbal crime.

References

ARCHER, JOHN. 2004. "Sex differences in aggression in real-world settings: A meta-analytic review". *Review of General Psychology* 8(4): 291–332.

AUSTIN, JOHN L. 1962. *How to Do Things with Words*. Cambridge, MA: Harvard University Press.

BOLINGER, DWIGHT L. 1989. *Intonation and its Uses: Melody in Grammar and Discourse*. Stanford: Stanford University Press.

BONEBRIGHT, TERRI L., THOMPSON, JERI L. & LEGER, DANIEL W. 1996. "Gender stereotypes in the expression and perception of vocal affect". *Sex Roles* 34(5/6): 429–445.

DIXON, JOHN A., MAHONEY, BERENICE & COCKS, ROGER. 2002. "Accents of guilt: Effects of regional accent, race and crime type on attributions of guilt". *Journal of Language and Social Psychology* 21(2): 162–168.

FRASER, BRUCE. 1976. "Warning and threatening". *Centrum* 3(2): 169–180.

FRASER, BRUCE. 1998. "Threatening revisited". *International Journal of Speech, Language and the Law* (formerly *Forensic Linguistics*) 5(2): 159–173.

GALES, TAMMY. 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

GINGISS, PETER. 1986. "Indirect threats". *Word* 17(3): 155–158.

GOBL, CHRISTER & NÍ CHASAIDE, AILBHE. 2003. "The role of voice quality in communicating emotion, mood and attitude". *Speech Communication* 40: 189–212.

GUSSENHOVEN, CARLOS. 2004. *The Phonology of Tone and Intonation*. Cambridge: Cambridge University Press.

LOCHMAN, JOHN E., POWELL, NICOLE R., CLANTON, NANCY & MCELROY, HEATHER K. 2006. "Anger and aggression". In: George G. Bear & Kathleen M. Minke (eds), *Children's Needs III: Development, Prevention, and Intervention*. Washington, DC: National Association of School Psychologists. pp. 115–133.

MILLER, NATHAN. 1928. *The Child in Primitive Society*. New York: Brentano.

NAPIER, MICHAEL R. & MARDIGIAN, R. STEPHEN. 2003. "Threatening messages: The essence of analyzing communicated threats". *Public Venue Security*, September/October, pp. 16–19.

POLZEHL, TIM, SCHMITT, ALEXANDER, METZE, FLORIAN & WAGNER, MICHAEL. 2011. "Anger recognition in speech using acoustic and linguistic cues". *Speech Communication* 53(9-10): 1198–1209.

PUBLIC ORDER ACT. 1986. London: Her Majesty's Stationery Office. Online resource: http://www.legislation.gov.uk/ukpga/1986/64/pdfs/ukpga_19860064_en.pdf

ROBERTS, LISA. 2011. "Acoustic effects of authentic and acted distress on fundamental frequency and vowel quality". *Proceedings of the 17th International Congress of Phonetic Sciences*, Hong Kong, August, pp. 1694–1697.

ROTHCHILD, JOHN. 1998. "Menacing speech and the First Amendment: A functional approach to incitement that threatens". *Texas Journal of Women and the Law* 8: 207–223.

SALGUEIRO, ANTONIO B. 2010. "Promises, threats, and the foundations of Speech Act Theory". *Pragmatics* 20(2): 213–228.

SCHIRMER, ANNETT & KOTZ, SONJA A. 2003. "ERP evidence for a gender specific Stroop effect in emotional speech". *Journal of Cognitive Neuroscience* 15(8): 1135–1148.

SCHIRMER, ANNETT & KOTZ, SONJA A. 2006. "Beyond the right hemisphere: Brain mechanisms mediating vocal emotional processing". *Trends in Cognitive Sciences* 10(1): 24–30.

SCHIRMER, ANNETT, ZYSSET, STEFAN, KOTZ, SONJA A. & VON CRAMON, D. YVES. 2004. "Gender differences in the activation of inferior frontal cortex during emotional speech perception". *Neuroimage* 21: 1114–1123.

SCHIRMER, ANNETT, KOTZ, SONJA A. & FREDERICI, ANGELA D. 2005. "On the role of attention for the processing of emotions in speech: Sex differences revisited". *Cognitive Brain Research* 24(3): 442–452.

SEARLE, JOHN R. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.

SEARLE, JOHN R. 1975. "Indirect speech acts". In: Peter Cole & Jerry L. Morgan (eds), *Syntax and Semantics, vol. 3: Speech Acts*. New York: Academic Press. pp. 59–82.

SEARLE, JOHN R. 1979. *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge: Cambridge University Press.

SHAVELL, STEVEN. 1993. "An economic analysis of threats and their illegality: Blackmail, extortion and robbery". *University of Philadelphia Law Review* 141: 1850–1903.

SHUY, ROGER W. 1993. *Language Crimes: The Use and Abuse of Language Evidence in the Courtroom*. Oxford: Blackwell.

SHUY, ROGER W. 2005. *Creating Language Crimes: How Law Enforcement Uses (and Misuses) Language*. Oxford: Oxford University Press.

SMITH, SHARON 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. 435–455.

STOREY, KATE. 1995. "The language of threats". *International Journal of Speech, Language and the Law* (formerly *Forensic Linguistics*) 2(1): 74–80.
TIERSMA, PETER M. & SOLAN, LARRY M. 2012. "The language of crime". In: Peter M. Tiersma & Larry M. Solan (eds), *The Oxford Handbook of Language and Law*. Oxford: Oxford University Press. pp. 340–353.
WESSA, PATRICK. 2012. *Pearson Correlation (v1.0.6) in Free Statistics Software (v1.1.23-r7), Office for Research Development and Education*. Online resource: http://www.wessa.net/rwasp_correlation.wasp
XU, YI & KELLY, ANDREW. 2010. "Perception of anger and happiness from resynthesized speech with size-related manipulations". *Proceedings of Speech Prosody 2010*, Chicago, 100027: 1–4.
XU, YI, KELLY, ANDREW & SMILLIE, CAMERON. 2013. "Emotional expressions as communicative signals". In: Sylvie Hancil & Daniel Hirst (eds), *Prosody and Iconicity*. Amsterdam: Benjamins. pp. 33–60.
YAMANAKA, NOBUHIKO. 1995. "On indirect threats". *International Journal for the Semiotics of Law* 8(2): 37–52.

Dominic Watt
Department of Language and Linguistic Science
University of York
Heslington
York
United Kingdom
YO10 5DD
email: dominic.watt@york.ac.uk

Sarah Kelly
Department of Language and Linguistic Science
University of York
Heslington
York
United Kingdom
YO10 5DD
email: sk720@york.ac.uk

Carmen Llamas
Department of Language and Linguistic Science
University of York
Heslington
York
YO10 5DD
United Kingdom
email: carmen.llamas@york.ac.uk