Language and Law / Linguagem e Direito, *Vol. 11*(2), 2024, pp. 94-114
DOI: doi.org/10.21747/21833745/lanlaw/11_2a4

The power of uptake: Responses to claims to power on anonymous online fora

Felicity Deamer [1], Michaela Wolf[2], Robert Davies[3], and Tim Grant [4]

[1] *Aston Institute for Forensic Linguistics*, f.deamer@aston.ac.uk
[2] *Aston Institute for Forensic Linguistics*, m.wolf@aston.ac.uk
[3] *Aston Institute for Forensic Linguistics*, r.davies6@aston.ac.uk
[4] *Aston Institute for Forensic Linguistics*, t.d.grant@aston.ac.uk

Abstract
Previous work in this area has identified powerful individuals on anonymous online fora as those who draw on multiple power resources in their interactions, and we were able to show how power operates differently across the different fora (Newsome-Chandler & Grant, 2023). Here we explore the same three datasets from a parenting forum, a white supremacist forum, and a Dark Web CSEA forum, but here we have shifted our focus to track those who are liable to be influenced or persuaded by claims to power. Using a purpose specific speech act coding framework, we have devised a taxonomy of responses to claims to power, and we provide a small-scale analysis of the effects of high- power resource individuals on those they are interacting with. We explore whether such a taxonomy of response types that can be scaled through automation to pick out individuals who may be more vulnerable to persuasion and influence. As well as the taxonomy and coding framework itself, we present findings from a close examination of all the posts and responses, and we explore whether persuadability remains constant or fluctuates according to situation or context, the interactant, or any other observable factor(s).
Keywords: *Power, Influence, Speech Acts, Online Fora.*

Resumo
Trabalhos anteriores nesta área identificaram como "indivíduos de poder", em fóruns online anónimos, aqueles que recorrem a múltiplos recursos de poder nas suas interações, e conseguiram demonstrar como o poder funciona de forma diferente nos diferentes fóruns (Newsome-Chandler & Grant, 2023).

Neste trabalho, exploramos três conjuntos de dados de um fórum de pais, um fórum de supremacia branca e um fórum CSEA da Dark Web; porém, aqui mudámos o nosso foco para seguir aqueles que são suscetíveis de serem influenciados ou persuadidos por reivindicações de poder. Utilizando um quadro de codificação de atos de fala com objetivos específicos, criámos uma taxonomia de respostas a reivindicações de poder e apresentamos uma análise em pequena escala dos efeitos de indivíduos com recursos de poder elevado sobre aqueles com quem estão a interagir. Investigámos se essa taxonomia de tipos de resposta pode ser adaptada através da automatização para selecionar indivíduos que possam ser mais vulneráveis à persuasão e à influência. Para além da taxonomia e do próprio quadro de codificação, apresentamos os resultados de uma análise minuciosa de todas as mensagens e respostas e, além disso, exploramos se a persuasão se mantém constante ou flutua de acordo com a situação ou o contexto, o interagente ou qualquer outro fator observável.
Palavras-chave: *Poder, Influência, Atos de fala, Fóruns Online.*

1. Introduction

Motivated by the need to better understand criminal online spaces, earlier work enabled us to identify powerful individuals on online anonymous fora as those who draw on multiple power resources in their interactions. Earlier findings demonstrate how power is wielded differently across three different fora; a general discussion forum, a white supremacist forum, and a Dark Web CSEA forum (Newsome-Chandler & Grant, 2023). Newsome-Chandler and Grant deliver a nine-resource model (see figure 1 below), and by evaluating how users assert authority—from community to technological expertise—the framework enables nuanced understanding of hierarchies and influence in hidden digital spaces.

Figure 1. Power resources

1. Personal experience
2. Broad topic expertise
3. Accredited expertise
4. Community expertise
5. Technological expertise
6. Veteran power
7. Subject of law enforcement
8. Private knowledge
9. Citing a secondary source

By developing a resource-based model of enacted power in interaction, Newsome-Chandler and Grant (2023) were able to observe both the commonalities and differences between the online communities presented in the three datasets. The analysis clearly showed that each forum demonstrates a distinct culture of power. This supports the idea that power is not a uniform or static concept, but one that varies significantly by context. The findings suggest that an individual's ability to enact power depends on their access to relevant power resources in a given context. This access may be influenced by personal history, experience, or by situational factors such as topic or content. Even individuals with a broad set of power resources may find themselves less effective in certain interactions. Thus, a "powerful" individual is best understood as someone who possesses a diverse repertoire of resources and is situated in an interaction that allows them to deploy those effectively.

This previous work focused on how users make claims to power to direct or influence the beliefs or behaviour of others. Here we explore the same three datasets, but we have shifted our focus to the other side of the interaction; onto those who are liable to be influenced or persuaded by such claims to power. The objective is to provide a small-scale analysis of the effects of claims to power on recipients – to create a pragmatics-based descriptive taxonomy of their responses to claims to power. Using a purpose-specific speech act coding framework, we have devised a taxonomy of response types that could potentially be scaled through automation to determine how influence, persuasion, and potentially radicalisation work within distinct anonymous online networks, and to potentially pick out individuals who may be more vulnerable to persuasion and influence.

2. Background and theoretical underpinnings

Utterances (and by extension, online posts) can be described as speech acts, meaning that they not only convey information but also perform a particular function or action (Austin, 1962; Põldvere, Felice, & Paradis, 2022). Every utterance can be analysed as simultaneously realising three distinct speech acts; the locutionary act is the production of a particular string of words in a particular grammatical structure encoding a specific meaning, the illocutionary act is the act that is performed in producing the utterance (i.e. asserting, promising, questioning) and the perlocutionary act is the resulting psychological or behavioural effect it has on any interlocutors.

The speaker can control the illocutionary speech act but not its perlocutionary effect. I can 'warn' you, hoping to 'persuade you against', but may in fact have the perlocutionary effect of 'encouraging you to'. A speaker (or poster in our case) can produce an utterance with a particular communicative intention in mind, but her interlocutor has the power to render this illocutionary act a success or failure. The speaker only has control over the illocutionary act; if I say "I assert that Elon Musk…", the listener cannot deny I have asserted, but she can refuse to believe me. Whether or not she believes me or not is the perlocutionary effect, or what McDonald (2020) refers to as the ratification theory of uptake.

Any post on an online forum will contain one or more illocutionary speech acts (e.g. an assertion, a question, a request, etc.) and depending on the size of a thread, there will be any number of responses to that original post. Underlying any given response

is the perlocutionary effect (i.e. whether the responder believes the assertion they're responding to, or whether they accept a request and intend to satisfy or meet it), but it isn't typically possible to determine from a response to a post, the precise perlocutionary effect (i.e. whether an assertion or a promise has been believed).For the purposes of this study, we think about perlocution at the most basic level i.e. whether a response is either accepting or rejecting the speech act(s) performed in the post it is responding to, which McDonald (2020) refers to as 'uptake'.

As mentioned above, there has been a large body of previous work exploring hierarchies of power on these online fora (Newsome-Chandler & Grant, 2023), which has shed light on how individuals express power through recurring claims to specific resources (e.g. broad topic expertise, citing of secondary sources, technological expertise etc.), and how power manifests differently across different online fora.

Table 1. Examples of claims to power resources (taken from the dataset):

"I came across this piece: [link]"	Citing a secondary source for authority
"I honestly don't know. There has been a lot of domestic military activity over the last few years.[...]"	Broad topic expertise
"I'd completely ignore the SWs opinions on the matter (unless she has medical degree).[...]"	Community specific initialism
"[...]so it does happen but many of the posters are foster carers first and in adoption without fostering first it is very rare, in my experience."	Direct personal experience
"I know some people who have lived in various Islamic areas for years.[...]"	Indirect personal experience
"There is software, like Eschalot, to create vanity onions."	Direct technological expertise
"I work in social care and in this area, each individual assessment is unique and each family different [...]"	Direct accredited expertise
"Regular with a name change here[...]"	Community expertise

1. Their work gives us a better idea of how users attempt to assert and gain power, but in order to understand how those claims to power impact other users on the fora, we need to explore how those interactants respond or act upon those claims to power. A user's power and influence are best measured as the impact they have on others. In other words, an online user is only as powerful as other users enable them to be. With this purpose in mind, this paper sets out to create a taxonomy of responses which captures the following information: Whether or not the response was to a **claim to power** (and if so, which type of claim to power – **Personal Experience, Community Expertise, Citing a Secondary Source etc.**) or to a post that made no claim to power
2. What speech acts were being responded to (i.e. assertion, question, commissive etc.)
3. Whether the response was an acceptance or rejection of the post it was responding to (**Acceptance/Partial Acceptance/Rejection**)

4. Whether that acceptance or rejection was made **explicitly** or **implicitly**

Just as Baker, Vessey, and McEnery (2021) found that the socio-cultural contexts (e.g. function, purpose, conventions, value systems) of different Islamic texts determine the language used and the route to persuasion and radicalisation, earlier preliminary analysis of interactional styles across the three fora illustrates that the different socio-cultural contexts, at least to some extent, determine the manner in which the claims to power resources are drawn upon interactionally. On the white supremacist forum, powerful users are skilful rhetoricians who can persuade others of an ideology.

Whereas on the dark web CSEA forum powerful users are those who effectively preserve the safe space for already radicalised individuals via the use of politeness features, hedging phrases, and expressive language. In line with this, we hypothesise that the different socio-cultural contexts might also determine (1) how persuasive or influential a given claim to power might be (i.e. how readily accepted it is), and secondly, how explicitly users are prepared to offer acceptance or rejection.

3. Research questions and data

In order to explore users' responses to claims to power, we set out to manually tag the same subset of 24 threads from each forum (selected on the basis of them starting with an explicit request for advice) which were used in earlier phases of this research (Newsome-Chandler & Grant, 2023) using a taxonomy of responses to explore the following patterns of rejection and acceptance:

1. Do we see any interaction between type of power resource and acceptance and rejection rates (and does this differ across fora?)?
2. Do we see any interaction between number or variety of claims to power and acceptance and rejection rates (and does this differ across fora?)?
3. Are some speech acts more persuasive than others (and does this differ across fora?)?
4. Do explicit and implicit acceptances and rejections pattern in different ways between types of claims to power, and within and between the three fora?

4. The coding framework

Previous research has attempted to develop automated approaches to speech act annotation (Core & Allen, 1997). Although computational methods can be used to support manual annotation and is beneficial in terms of time efficiency (Archer, Culpeper, & Davies, 2008, p. 637), it is difficult to fully automate pragmatic annotation. This is due to its complex and context-dependent nature, i.e. implicitness/explicitness, which present difficulties that can be overcome by manual annotation (Yu, Li, Su, & Fuoli, 2023, p. 1; Milà-Garcia, 2018; Kohnen, 2015, p. 64; Archer et al., 2008; McAllister, 2015; McEnery, Xiao, & Tono, 2006). Milà-Garcia (2018) has previously implemented manual annotation to analyse conversational disagreement and agreement in different discursive settings, and has demonstrated its practicality in the field of speech act analysis. Similarly, Rees-Miller (2000) conducted a manual analysis of disagreement in American academic talks and courses. The focus of the study was to investigate a potential relationship between the way disagreement is expressed and the speaker's social power (professor vs. student), the severity of disagreement, and the context disagreement appears in.

Given how there is no 'off the shelf' coding framework which meets the needs of our research, it was necessary to develop our own that was tailored to our specific research data.

4.1. The overall process

Step 1: Initial Development of the coding framework
Step 2: Preliminary Coding and Discussion to refine the coding scheme
Step 3: Coding scheme finalised
Step 4: Inter-Rater Reliability measured - coders individually coded an unseen section of data. Inter-rater reliability assessed using percentage agreement and Cohen's kappa.
Step 5: Disagreements discussed and a final coding decision jointly made.
Step 6: Division of Dataset where coders began separately coding individual halves of the remaining corpora using the finalised coding scheme.

4.2. Development of the coding framework

As in earlier similar studies, our coding scheme is designed to be as broad as necessary whilst still capturing responses to a diverse range of speech acts as it is important to balance "usefulness" and "ease of consistency" within coding (Archer et al., 2008, p. 631). The first step was to manually identify the types of speech acts that were being responded to (e.g. assertion, assertion to fiction, assertion based on secondary source, request, question, suggestion- expressing a potential future action to another poster to consider, commissive – expressing one's own commitment to an action in the future which can include promises, intentions, and plans). Assertion and suggestion are both typically in declarative form; request and question are both typically in interrogative form. Commissives also tend to be in declarative form, but the difference between suggestions and commissives is that suggestions predicate a future action to the addressee and commissives to the speaker.

Secondly, it was determined whether the speech act in the post being responded to was either accepted or rejected. We also considered whether this was done explicitly or implicitly. An acceptance of a speech act was defined as cases where a responder fulfils an action (e.g. answering a question or fulfilling a request), or where the responder "shares the same point of view [...] as the previous speaker" (Milà-Garcia, 2018, pp. 274–275). On the other hand, rejection was defined as cases where a responder does not fulfil an action (e.g. not doing what is requested or stating that an action won't be fulfilled), or where "the speaker's point of view diverges from that of the previous speaker" (Milà-Garcia, 2018, p. 275). Finally, a partial acceptance was defined as cases where a reply fell between acceptance and rejection and did not fit neatly into either (e.g. replying to an assertion with "I'll think about it"). The questions we asked ourselves when constructing the framework are detailed below.

It is important to note here that this binary classification system in which a response is classed as either accepting or rejecting a claim to power is a very coarse-grained measure of users' evaluations of claims to power. We are inferring evaluations of claims to power based on responses. This is an important starting point, which will enable future research focusing on explicit evaluations of claims to power. Sinclair and Coulthard's (1992) three-part model of spoken discourse offers a useful framework to think about

this with. They observe that spoken discourse can be carved up into three parts or stages;

1. Initiation: this stage involves one speaker initiating the conversation by starting a new topic or responding to a previous topic. This can include asking questions, making statements, or giving directives.
2. Response: In this stage, the other participant in the conversation responds to the initiation made by the first speaker. This can involve answering questions, elaborating on the topic, or showing agreement or disagreement.
3. Feedback: The final stage of the model involves the first speaker providing feedback on the response given by the second speaker. This can involve acknowledging the response, asking for clarification, or further developing the conversation.

Here we are exploring parts 1 and 2 of this discourse model, and have set out to determine (via users' responses) whether the initiation has been accepted, but we haven't had access to, and haven't sought to explore the feedback stage explicitly.

Figure 2. Steps to constructing the framework

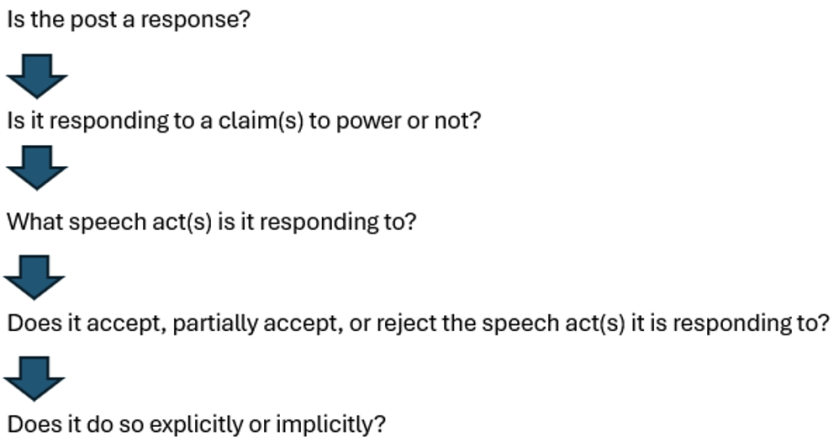

Figure 3. Coding Scheme Capturing Responses to Speech Acts

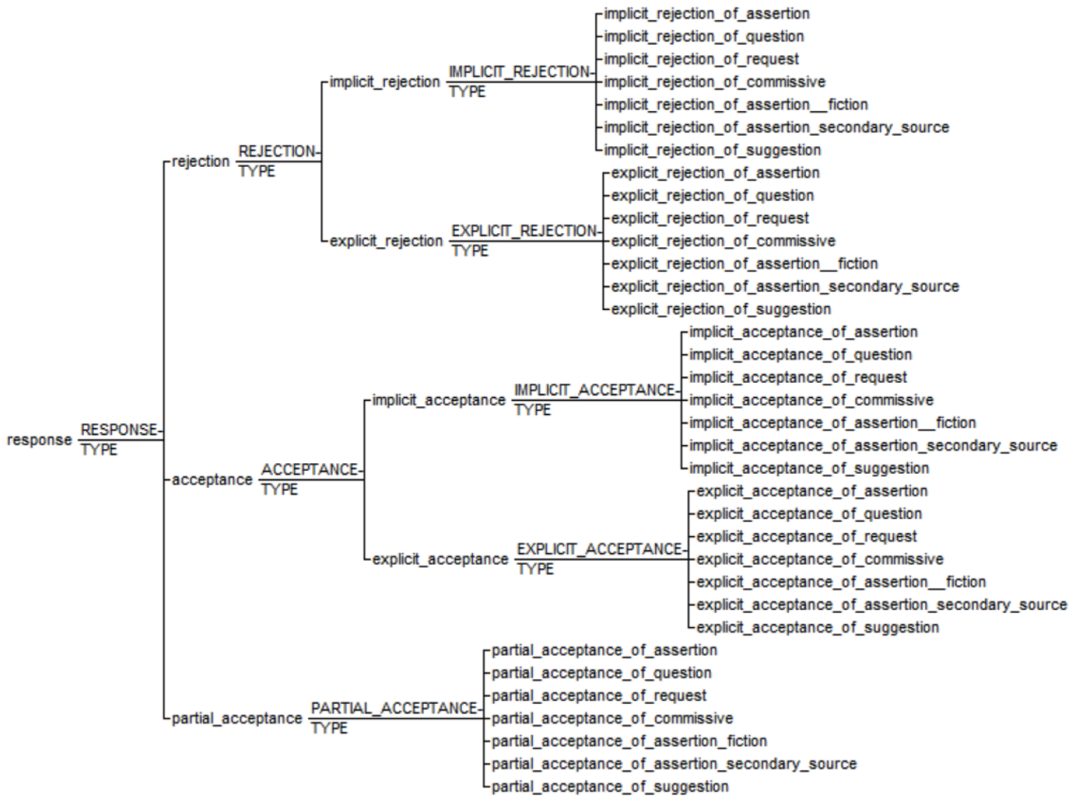

Coding examples

The table below shows demonstrative examples of the coding scheme applied.

Table 2. Demonstrative examples of Coding

Implicit Acceptance (Request)	U1: "Can someone provide me with a link to a useful website, please?" U2: "Hey, please message me privately"
Implicit Rejection (Commissive)	U1: "I'm going to talk to my social worker about this today." U2: "Maybe think about that again"
Explicit Acceptance (Question)	U1: "How old is your child? Can they speak yet?" U2: "She's 3 months old, so she definitely can't speak just yet!"
Explicit Rejection (Assertion to fiction)	U1: "[fictional story]" U2: "This story is horrible"
Partial Acceptance (Suggestion)	U1: "Maybe you should speak to your husband about this." U2: "I am going to think about it."

It is important to consider that manual analysis involves an element of subjectivity, and can therefore be prone to errors caused by human factors such as fatigue. To mitigate this, we introduced inter-rater agreement testing (Yu et al., 2023, p. 4).

4.3. Inter-rater reliability

We used two methods of evaluating interrater reliability: percentage agreement and Cohen's Kappa. Percentage agreement is known to overestimate the reliability between two coders as it takes no account of the chance agreement between the two coders. Cohen's Kappa, however, is calculated by taking the number of possible answers in each coding decision into account,but it makes the assumption that the choices are all independent of one another. In our data this assumption is not true, and this means that Kappa is likely to underestimate reliability. By presenting both figures, we get a range within which 'true' reliability might be expected to sit.

Table 3. Inter-rater reliability

		Percentage Agreement	Cohen's kappa
Coding of power presence	No Code | Ambiguous | No Claim to power | Claim to Power	89%	0.85
	No Code | n/a | Personal experience | Community Initialism | Vet. | Accredited expertise | Broad Topic Expertise | Subject to Law Enforcement | Cite | Personal Knowledge | Community Expertise	88%	0.87
Coding of replies	No Code | Accept | Reject | Partial	82%	0.76
	No Code | Explicit | Implicit | NA	71%	0.62
	No Code | Assertion | Question | Request | Commissive | Thanking | Assertion-Fiction | Assertion-sec-source | Suggestion	87%	0.86

Figure 4. Questions guiding coding framework

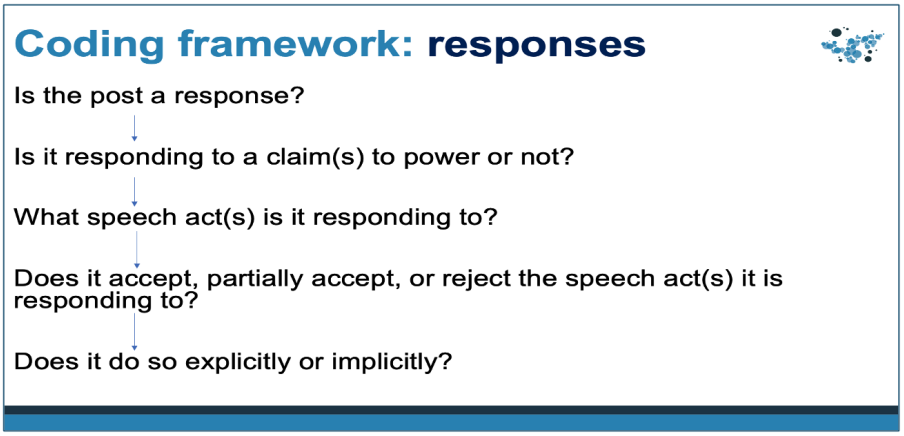

Examples of disagreement

Partial and implicit acceptance/rejection (Demonstrative Example)

U1:	"All I can speak of is (FORUM), never a member but the "sample" stuff and the leaked stuff is generally plain sick [....] someone. Not to mention I have heard elsewhere what Lux said about the content being trash"
U2:	"Now I am intrigued. What could possibly be on the site that got you so upset? The posters above seemed to describe it as the same old stuff and nothing too new."

Implicit vs. explicit (Demonstrative Example)

The responder could be asking the question out of interest; "What could possibly be on the site that got you so upset?", although the use of 'possibly' may indicate an element of doubt towards the assertion. We settled on a partial acceptance of the assertion as the meaning of the utterance may lie somewhere in-between.

U1:	"What made you join this forum?"
U2:	"Welcome to the forum! I was involved with Czech Nationalism for most of my life but they were made to join the European Union, so I had to focus my efforts elsewhere. I still can't believe our people chose destruction but somehow saw it as an improvement. "

The response could be understood as an implicit acceptance of the question due to a potential need to infer how the Czech Republic joining the EU could then lead to the poster joining a white supremacist website. However, due to the fact that the poster is still providing a reason for joining (that being because "they were made to join the European Union" and the person "had to focus [their] efforts elsewhere"), the reply was coded as an explicit acceptance of the question.

5. Analysis design

After completing the coding for all the threads across the three fora, the raw frequencies for the various categories (e.g. frequency of acceptances across all fora; the frequency of acceptances for each speech act etc.) were extracted from Corpus Tool and input into Excel sheets. Given that the data are frequency counts of categorical variables, chi-square tests were seen as a good choice (Grant et al., 2017) for the more straightforward comparisons (e.g. acceptance rates for posts containing a claim to power compared to those containing no claim to power). In corpus linguistics, chi-square tests are often used to evaluate keyness (Rayson, Archer, Piao, & McEnery, 2004) i.e., to determine whether a feature occurs significantly more often in one corpus over another (Baker, 2006). Similarly, we are exploring frequency of response type (acceptances; partial acceptances; rejections), and use chi-quare tests to determine whether a response type occurs significantly more often in one forum over another, in response to one speech act over another, and/or one claim to power or another. Because chi-square tests become too complex and difficult to interpret with multiple factors and levels (Grant et al., 2017), only two categories can be compared at a time using chi-square. When exploring

the effects of multiple factors (i.e. fora, type of claim to power, response type, degree of explicitness), we use multiway frequency tables (loglinear analysis) to explore a possible association between our factors (i.e. fora, type of claim to power, response type, degree of explicitness). This loglinear analysis makes it possible for us to determine the unique contribution of each factor and its interactions with other factors.

The findings presented and discussed below offer a snapshot of how anonymous users are responding to claims to power delivered via different speech acts, across our three fora. There are numerous ways in which this richly coded dataset could be explored, and the possibilities for future analysis and exploration are numerous, but for the purposes of this paper, we are focusing on whether or not persuasion occurs differently in different online anonymous networks, where the socio-cultural context differ, and where the interactional style and routes to persuasion and/or radicalisation are likely to differ as a result. With this in mind, we will be exploring a series of specific questions, which are detailed below.

1. *Is response type (acceptance, partial acceptance, or rejection) dependent on whether a claim to power (or indeed multiple claims to power) is made in the post that is being responded to?*

Given the pragmatic distinction between a post containing a claim to power and one that doesn't, it is reasonable to expect to find a difference between acceptance and rejection rates of these posts. Moreover, given that previous findings indicated that those making multiple different claims to power are likely to be higher up the hierarchy of power than those making fewer or multiple of the same claims to power, differences between responses to multiple claims to power and one or no claims to power are to be expected.

2. *Do those acceptance and rejection rates pattern differently across the three different fora?*

Given that we know the purpose, conventions, and requirements of the three fora differ significantly, and that some Power Resources (PwRs) are drawn upon differently in our three different fora, we might expect to find that responses to claims to PwRs (or lack thereof) might differ between fora.

3. *Is the type of response (accept, partially accept, or reject) dependent on which speech act is being responded to?*

Given that different speech acts are performing different illocutionary acts (some of which will differ with respect to politeness for example), we might expect to find that the likelihood of a particular speech act being accepted or rejected might differ.

4. *Does the relationship between response type and speech act differ across fora?*

Given that we know the function and interactional style of the three fora differ significantly, and that some PwRs (which are likely to be delivered via particular speech acts) are drawn upon more in one forum than another, we might expect to find that certain speech acts are more likely to be accepted or rejected in one forum over another.

5. *Is the type of response (accept or reject) dependent on which power resource is being responded to?*

In question 1, we explored whether response type is dependent on whether a claim to power, or multiple claims to power are made, but we also want to explore that further by unpacking whether there are differences in acceptance and rejection rates between specific PwRs.

6. *Does the degree of explicitness of response differ across fora?*

We know that the discourse context of each of our online fora differ and the function and purpose of a forum can steer the way in which users interact, as well as the rhetorical and communicative devices they employ to influence and persuade each other. On the basis of this, we hypothesised that we might see more like-minded implicit responses in the fora that facilitate normalisation (i.e. Dark Web and general discussion forum), acceptance, and support, so as to hedge, and indirectly acknowledge and manoeuvre interlocutors. However, it is equally possible that on the white supremacist forum, where users are likely to be vying for a position higher up the hierarchy of skilful rhetoricians (Deamer, Busso, Htait, & Grant, 2023), is also a space in which implicit responses are commonplace, but this time more likely as a face-saving strategy.

7. *Is the degree of explicitness dependent on whether a claim to power (or multiple) has been made, and does any such dependence differ across fora?*

We answer this by comparing the frequency of explicit and implicit responses to claims, no claims, and multiple claims to power, to explore whether responses are dependent on whether or not a claim to power has been made or not.

6. Salient findings

Across fora

1. **Users are more likely to reject posts containing a claim or multiple claims to power than if no claim to power is made at all, but they do not appear more or less likely to accept or reject a multiple claim to power than a single claim to power**

A Chi-square test of independence revealed that differences in acceptance and partial acceptance rates between posts containing a claim to power and those making no claim to power was not significant (p=.067; p=.093), but there are significantly more rejections of posts containing a claim to power than to those making no claim to power (p=.004).

Moreover, when comparing responses to multiple claims to power and posts making no claim to power, there is a significant difference in frequencies of partial acceptances (p=.002) and rejections (p=.039), but not acceptances (p=.068).

Finally, we can observe that there is no significant variation between responses to posts containing multiple claims to power and those containing just one claim to power (acceptances p=.889; partial acceptances p=.323; rejections p=.690).

Figure 5. Proportion of responses to claims to power

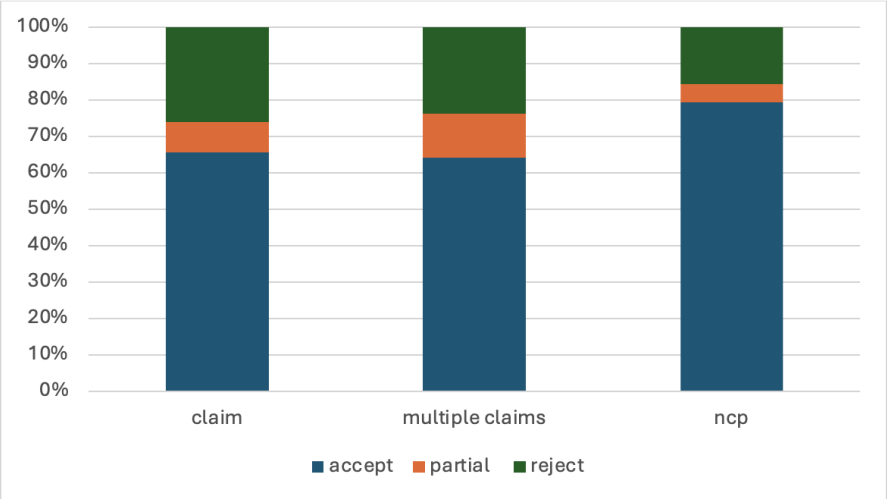

2. **Users are more likely to reject an assertion to a secondary source than a standard assertion**

Although proportions of response type do differ, users are not significantly more likely to accept (p=.17) or partially accept (p=.845) an assertion than an assertion to a secondary source, but there are significantly more rejections of assertions to a secondary source than to standard assertions (p=.049).

3. **Users are more likely to reject a request than a question**

Similarly, acceptance and partial acceptance rates of questions and requests do not differ significantly (accept p=.295; partial accept p=.686), but rejections did (p<.001).

4. **Users are more likely to accept a question than a standard assertion or an assertion to a secondary source.**

Figure 6, below, illustrates that although proportions of response type do differ, users are not significantly more likely to accept (p=.17) or partially accept (p=.845) an assertion than an assertion to a secondary source, but there are significantly more rejections of assertions to a secondary source than to standard assertions (p=.049). Similarly, acceptance and partial acceptance rates of questions and requests do not differ significantly (accept p=.295; partial accept p=.686), but rejections did (p<.001). Moreover, when comparing responses to assertions and assertions to a secondary source with questions, we see that users are significantly more likely to accept a question than an assertion (accept p<.001) or an assertion to a secondary source (accept p=.001). Rejection rates also differ significantly showing that assertions to secondary source and standard assertions get rejected more than questions (assertion/question p<.001; assert ss/question p<.001). No other comparison between speech acts is significant (likely due to low frequencies).

Figure 6. Proportion of responses to speech acts

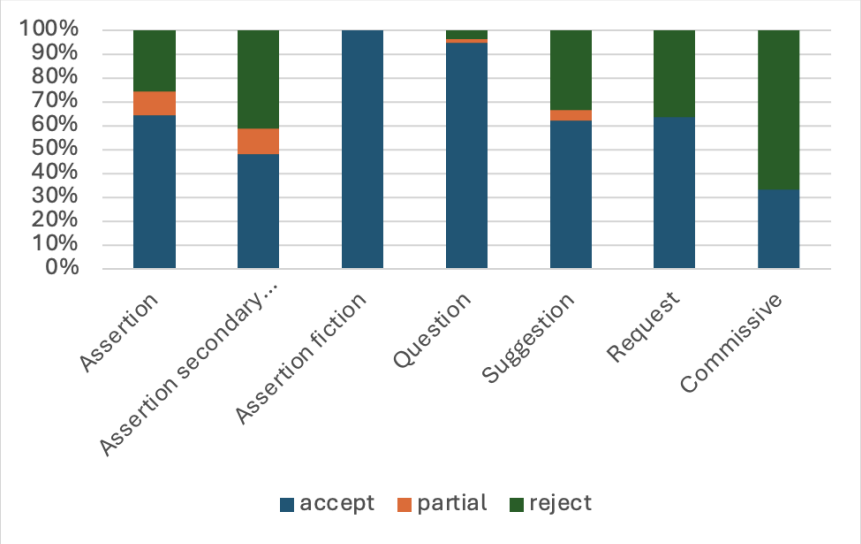

5. **Is the type of response (accept or reject) dependent on which power resource is being responded to?**

Although differences in proportion of response types to different power resources are visible in figure 7, none of these differences were significant (when comparing only the PwRs which had high enough frequency of responses). Again, this is likely due to low frequencies.

Figure 7. Proportion of responses to power resources (with high enough frequencies)

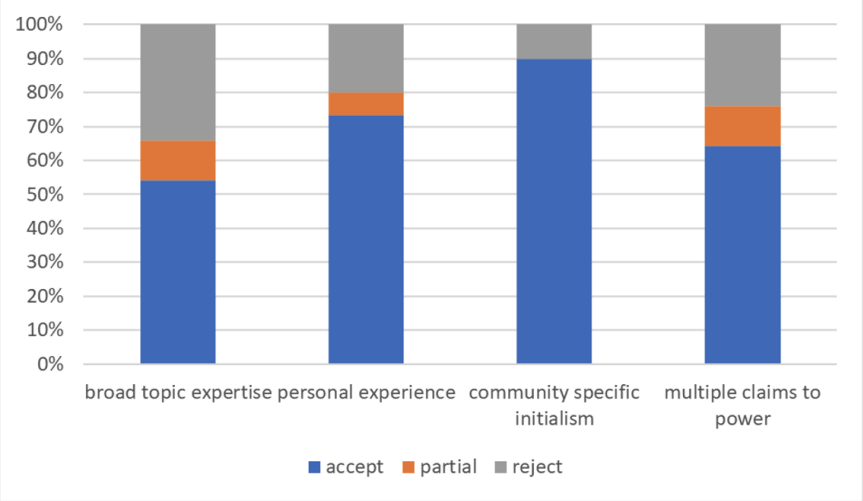

6. **Degree of explicitness is not dependent on whether or not a claim to power has been made or not across fora.**

Again, we ran chi-square tests comparing frequency of explicit and implicit responses to claims, no claims, and multiple claims to power, which show that degree of explicitness is not dependent on whether or not a claim to power has been made or not.

Figure 8. Proportions of explicit and implicit responses to claims to power

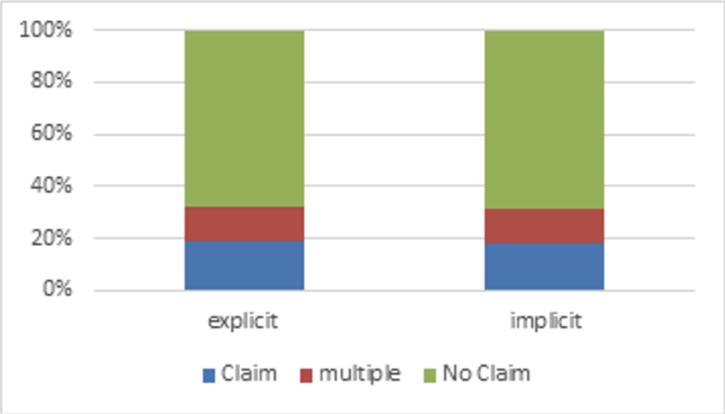

Between fora

7. **Response type is dependent on the fora the interaction is taking place on, and on whether one or multiple claims to power are made or not**

We can observe in figures 9-11 that acceptance rates appear to be higher in response to posts containing no claim to power than in response to one or multiple claims to power across all three fora. And users appear more likely to reject multiple claims to power across all three fora, particularly on the dark web. Loglinear modelling allowed us to determine whether response type is in fact dependent not just on whether a claim or multiple claims to power has been made or not (as explored in 5.1 above), but also whether it is dependent on the fora, and whether there is an interaction between those three factors. The multiway frequency analysis reveals that there is a two-way ($p<.001$) but not a three-way ($p=.33$) interaction between factors. Namely, response type is indeed dependent on which fora the interaction is taking place on ($p=.011$), and on whether one or multiple power claims are made or not ($p<.001$), but no three-way interdependence between attributes was found i.e. the association between response type and power claims does not differ between fora.

Figure 9. Proportion of responses to claims to power White supremacist forum

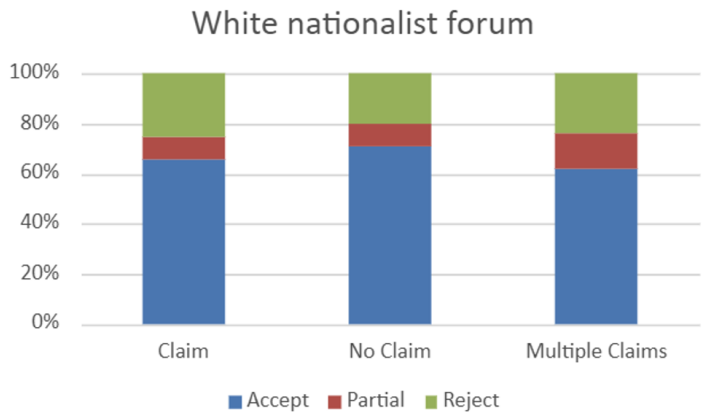

Figure 10. Proportion of responses to claims to power General discussion forum

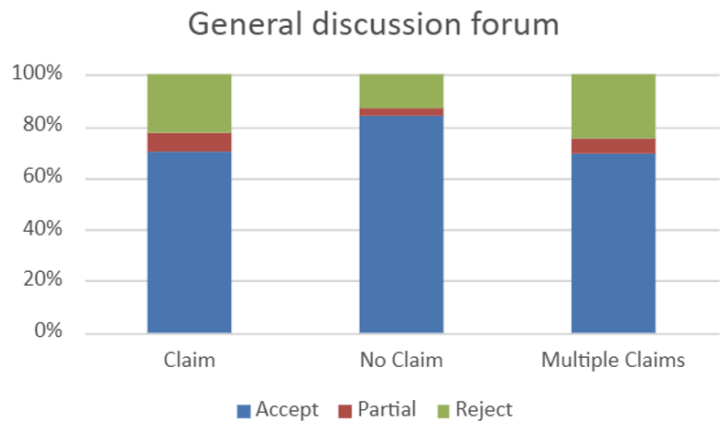

Figure 11. Proportion of responses to claims to power Dark Web

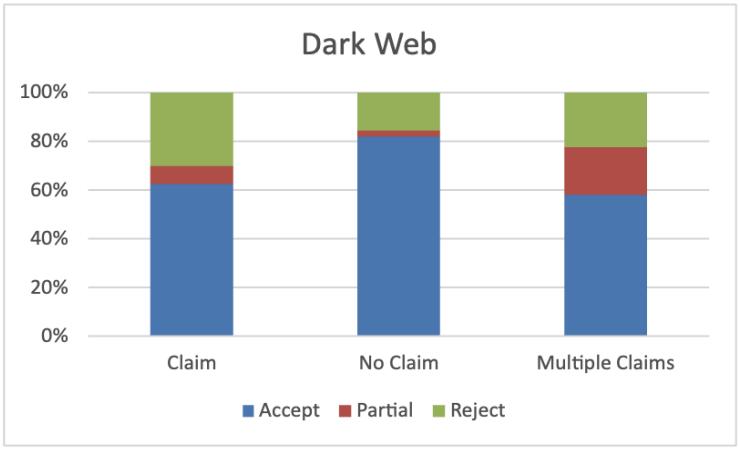

8. **The relationship between response type and speech act does not appear to differ significantly between fora**

Loglinear modelling reveals no main effect or higher order interaction between our factors (response type; speech act; fora). As discussed above, when using chi square tests to compare responses to two speech acts at a time, we do see some significant differences, but it is likely that frequencies are too low for differences and interactions to be detected in a loglinear model.

9. **Responses are more often implicit on White supremacist forum than General discussion forum and the Dark web**

We can see from figure 12 that responses are significantly more implicit on the White supremacist forum than General discussion forum and the Dark web, but when we break that down into acceptances and rejections (figure 13), and explore the differences using loglinear analysis, those differences are lost.

Figure 12. Proportion of explicit and implicit responses between fora

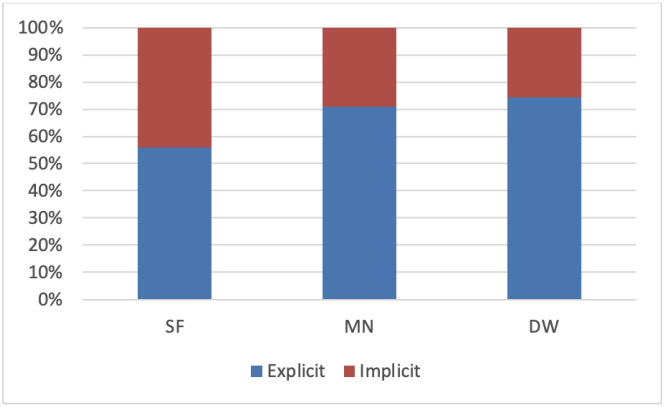

Figure 13. Proportion of explicit and implicit acceptance and rejection between fora

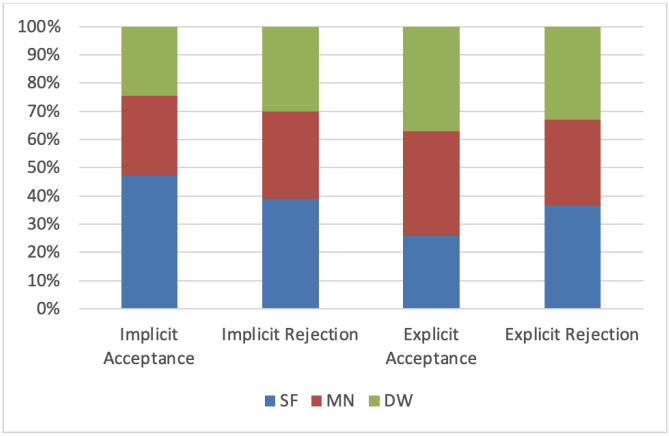

10. Relationship between degree of explicitness and whether or not a claim (or multiple) to power was made does not differ between fora

Loglinear analysis reveals that there are no interactions between factors i.e. these differences did not vary significantly between fora.

Figure 14. Proportions of explicit and implicit responses to claims to power White supremacist forum

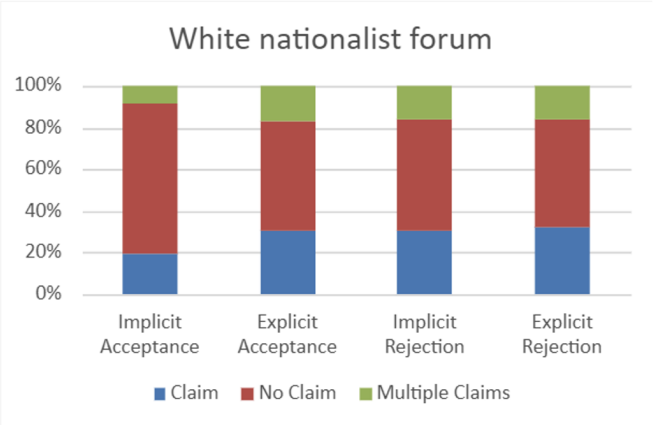

Figure 15. Proportions of explicit and implicit responses to claims to power General discussion forum

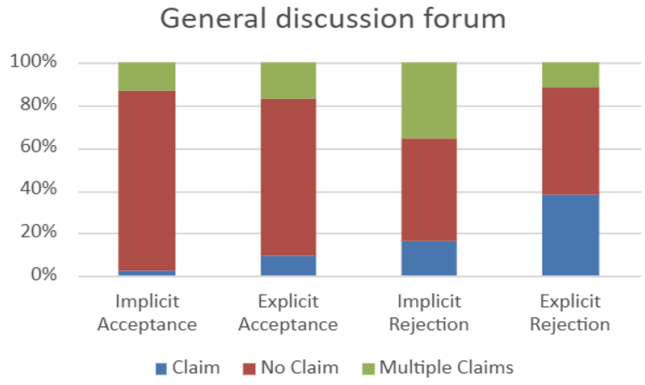

Figure 16. Proportions of explicit and implicit responses to claims to power dark web

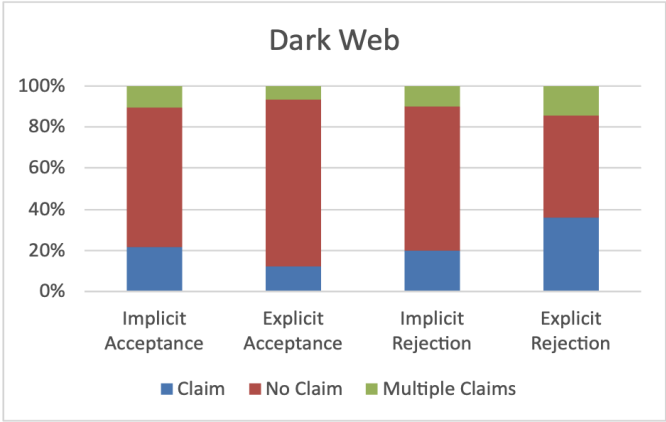

7. Conclusion

Power is hierarchical and inherently bidirectional, where one individual or group seeks to dominate and control another. Having previously identified a number of discrete resources drawn upon by anonymous users to claim power, here we have explored the other half of the bidirectional relationship. Namely, how other users respond to those claims to power, and what that says about (a) the efficacy of specific types of claims to power in different socio-cultural contexts, where routes to persuasion and radicalisation looks somewhat different, and (b) how power dynamics play out interactionally in those different contexts.

Like any hierarchy of power, the power structures on the anonymous online fora examined are unlikely to be set in stone, and some of our findings are perhaps indicative of users jostling for power, saving face, and/or displaying a reluctance to 'lend face' by recognising and acknowledging the power resources of other users. For example, we can see that users are more likely to reject posts containing at least one claim to power than if no claim to power is made at all. We can also see that assertions to a secondary source, in which a user is arguably putting forward a strong display of knowledge and expertise by referencing a robust piece of evidence, are more likely to be rejected by other users than straight forward assertions (not backed up by evidence). Moreover,

agency curtailing speech acts such as requests (a strong act of dominance) are more likely to be rejected than agency enhancing speech acts such as questions (in which the interlocutor is actually invited to make their own claim to power). These findings are perhaps indicative of a broader observation which has been made in other contexts, namely that one is more likely to make a claim to power (i.e. listen to me because.…) or be more forceful (Grant & Woodhams, 2007) when one expects some resistance, or when there is a strong possibility of refusal.

Another notable finding is that both acceptances and rejections are more likely to be implicit on the white supremacist forum than on the general discussion forum or the dark web. This is significant because we know from initial explorations of socio-cultural context and interactional style on the three fora that White supremacist forum is a space in which users who make frequent claims to multiple power resources do so in a manner that could be described as 'hard power', where claims to power are typically formulated via inductive argument without any apparent accommodation to the recipient's frame of mind (e.g. more face-threatening and less hedging). Whereas a powerful user on General discussion forum and the dark web is more likely to consider the psychological state of those they're interacting with, and moderate their knowledge and expertise with emotional awareness. It makes sense then, that recipients of claims to power on White supremacist forum (where users are jostling to move up the hierarchy and create a reputation for themselves as skilful rhetoricians) are less likely to lend face with an explicit acceptance, even if they are in agreement. Moreover, the White supremacist forum primarily functions as a platform for persuasion of an ideology, and much like in any political discourse where success is measured by the growth and popularity of the movement, it makes sense that users are likely to avoid impeding the proliferation of the ideology by explicitly rejecting the claims put forward by other users in support of that ideology.

Further exploration is needed, but response type does appear to be dependent on the fora the interaction is taking place on, as well as on whether at least one claim to power is being responded to. Although we don't see a relationship between degree of explicitness and whether or not a claim (or multiple) to power is made, or a relationship between response type and speech act, this does require further exploration, and expanding our analyses into different corpora may reveal more insights here.

There are many other questions we could ask of this data, and we have explored just a few of them. For example, it would be informative to take a closer look at whether the speech act being responded to dictates the explicitness of the response, and whether that triangular relationship of explicitness, response type, and speech act differs across fora. These analyses would also benefit from the expansion into new and different anonymous fora, beyond the three explored here.

Finally, based on our richly coded dataset from the three fora we have focused on to date, it would be straightforward to identify individuals who have high acceptance rates across the three different fora, and to closely examine all the posts and responses associated with those persuadable and/or easily influenced individuals, and to explore whether throughout their messaging their persuadability remains constant or fluctuates according to situation or context, the interactant or any other observable factor(s).

References

Archer, D., Culpeper, J., & Davies, M. (2008). Pragmatic annotation. In M. Kytö & A. Lüdeling (Eds.), *Corpus Linguistics: An International Handbook* (pp. 613–642). Mouton de Gruyter.

Austin, J. L. (1962). *How to Do Things with Words*. Oxford: Clarendon Press.

Baker, P. (2006). *Using corpora in discourse analysis*. Continuum.

Baker, P., Vessey, R., & McEnery, T. (2021). *The Language of Violent Jihad*. Cambridge: Cambridge University Press.

Core, M., & Allen, J. (1997). Coding dialogs with the DAMSL annotation scheme. In (pp. 28– 35). Cambridge, MA.

Deamer, F., Busso, L., Htait, A., & Grant, T. (2023, March). (Unpublished government report). *Hierarchies of Power*, *5*.

Grant, T., Clark, U., Reershemius, G., Pollard, D., Hayes, S., & Plappert, G. (2017). *Quantitative Research Methods for Linguists: A Questions and Answers Approach for Students* (1st ed.). Routledge. Retrieved 2025-08-11, from https://www.taylorfrancis.com/books/9781351722988 doi: 10.4324/9781315181707

Grant, T., & Woodhams, J. (2007). Rape as Social Activity: an Application of Investigative Linguistics. In J. Cotterill (Ed.), *The Language of Sexual Crime* (pp. 1–15). London: Palgrave Macmillan UK. Retrieved 2025-08-11, from http://link.springer.com/10.1057/9780230592780_1 doi: 10.1057/9780230592780_1

Kohnen, T. (2015). Speech acts: a diachronic perspective. In K. Aijmer & C. Rühlemann (Eds.), *Corpus Pragmatics: A handbook* (1st ed., pp. 52–83). Cambridge University Press. Retrieved 2025-08-11, from https://www.cambridge.org/core/product/identifier/9781139057493%23c01504-2-0/type/book_part doi: 10.1017/CBO9781139057493.004

McAllister, P. G. (2015). Speech acts: a synchronic perspective. In K. Aijmer & C. Rühlemann (Eds.), *Corpus Pragmatics: A handbook* (1st ed., pp. 29–51). Cambridge University Press. Retrieved 2025-08-11, from https://www.cambridge.org/core/product/identifier/9781139057493%23c01504-1-0/type/book_part doi: 10.1017/CBO9781139057493.003

McDonald, L. (2020). Your word against mine: the power of uptake. *Synthese*, *199*(1-2), 3505–3526. Retrieved 2025-08-11, from https://link.springer.com/10.1007/s11229-020-02944-1 doi: 10.1007/s11229-020-02944-1

McEnery, T., Xiao, R., & Tono, Y. (2006). *Corpus-based language studies: an advanced resource book*. New York: Routledge.

Milà-Garcia, A. (2018, September). Pragmatic Annotation for a Multi-Layered Analysis of Speech Acts: A Methodological Proposal. *Corpus Pragmatics*, *2*(3), 265–287. Retrieved 2025-08-11, from http://link.springer.com/10.1007/s41701-018-0037-z doi: 10.1007/s41701-018-0037-z

Newsome-Chandler, H., & Grant, T. (2023). Developing a Resource Model of Power and Authority in Anonymous Online Criminal Interactions. *Language and Law / Linguagem e Direito*, *10*(1), 110–130. Retrieved 2025-08-05, from https://ojs.letras.up.pt/index.php/LLLD/article/view/12841/12478 doi: 10.21747/21833745/lanlaw/10_1/a4

Põldvere, N., Felice, R. D., & Paradis, C. (2022). Advice in Conversation: Corpus Pragmatics Meets Mixed Methods. In P. Culpeper & M. Haugh (Eds.), *Elements in*

pragmatics (1st ed., pp. 1–80). Cambridge University Press. Retrieved 2025-08-11, from https://www.cambridge.org/core/product/identifier/9781009053617/type/element doi: 10.1017/9781009053617

Rayson, P., Archer, D., Piao, S. L., & McEnery, T. (2004). The UCREL semantic analysis system..

Rees-Miller, J. (2000). Power, severity, and context in disagreement. *Journal of Pragmatics*, *32*(8), 1087–1111. Retrieved 2025-08-11, from https://linkinghub.elsevier.com/retrieve/pii/S0378216699000880 doi: 10.1016/S0378-2166(99)00088-0

Sinclair, J. M., & Coulthard, M. (1992). Towards an analysis of discourse. In M. Coulthard (Ed.), *Advance in Spoken Discourse Analysis.* London: Routledge.

Yu, D., Li, L., Su, H., & Fuoli, M. (2023). Assessing the potential of AI-assisted pragmatic annotation: The case of apologies. Retrieved 2025-08-11, from https://arxiv.org/abs/2305.08339 (Publisher: arXiv Version Number: 5) doi: 10.48550/ARXIV.2305.08339