Power and Vulnerability: Managing Sensitive Language in Organisational Communication

Patrick G. T. Healey
Queen Mary
University of London

Prashant Khare
Queen Mary
University of London

Ignacio Castro
Queen Mary
University of London

Gareth Tyson
HKUST (GZ) & Queen Mary
University of London

Mladen Karan
Queen Mary
University of London

Ravi Shekhar
Queen Mary
University of London

Stephen McQuistin
University of Glasgow

Colin Perkins
University of Glasgow

Matthew Purver
Jozef Stefan Institute & Queen
Mary University of London

Abstract

Organisational responsibilities can bring power but also a degree of vulnerability and exposure. This leads to divergent predictions about the use of potentially sensitive language: power might license it, exposure might inhibit it. Data from a large corpus of organisational emails shows that people in positions of relative power tend to avoid potentially sensitive words suggesting that, in at least some circumstances, vulnerability is a more significant influence than power in organisational language use.

Keywords: Dialogue, Politeness, Power, Communication

1 Managing Sensitive Topics in Conversation

In their study of politeness in conversation, Brown and Levinson observe that sensitive topics such as: "politics, race, religion, women's liberation" (p.314), pose a danger to both speakers and hearers (Brown et al., 1987). Sensitive topics can threaten the self-image or *positive face* of a hearer if, for example, the talk refers to attributes of that individual. They can also pose a threat to the self-image of a speaker if, for example, the hearer challenges their remarks.

Judgements of what is appropriate are complex and vary between times, contexts, individuals and speech communities. People can introduce sensitive topics with the intention of strengthening a relationship, e.g. as a form of self-disclosure, just as much as with the intention of threatening it.

One of the most significant complicating factors in judgements of appropriateness is social power, understood as mutually agreed relative status (see p.320). As Brown and Levinson observe, differences in social power distort the equilibrium of "mutual vulnerability" to face threats.

How do we respond to changes in this equilibrium? Brown and Levinson introduce a *Power Hypothesis*. Roughly, the more powerful a person is, the less vulnerable they are to threats to their own self-esteem (*positive face*) and the more able they are to resist requests or instructions that might restrict their freedom of action (*negative face*). Moreover, they can also be less concerned with the face needs of others.

An alternative *Exposure Hypothesis* is that people of higher status are potentially more vulnerable to face threats, both to themselves and others. Intuitively, people with significant decision making responsibilities are subject to additional social pressures; they may need to work to build a consensus around decisions, they often need to manage the potential challenges to a decision, and they may be held personally accountable for those decisions.

These two hypotheses make contrasting predictions. The Power Hypothesis suggests that people with more social power should be more likely to use potentially sensitive words whereas the Exposure Hypothesis suggests that they should be less likely to use those words. We explore these predictions in a large corpus of email communication produced by the Internet Engineering Task Force (IETF) a debate-based,

consensus-driven, forum that brings together multiple stakeholders (industry, academia, civil society) to agree many of the technical standards (e.g., TCP/IP, HTTP) that ensure the Internet works.

2 Method

The IETF is structured into Working Groups (WGs), each with a particular technical focus (e.g., HTTP protocol) and a mailing list (most IETF works takes place via email). WG chairs facilitate the work with responsibilities including moderating mailing lists, organising meetings, setting the agenda, and judging consensus on major decisions.

We use two publicly available data sources: the IETF mail archives[1] and the Datatracker.[2] The mail archives cover WG activities, meetings, and administration. The Datatracker provides information about organisational roles of participants.[3]

The sample consists of all WG email communications in 2019, the last 'pre-Covid' year for which we have complete data. There are a total of 51,977 emails, from 2108 unique participants, across 176 group mailing lists. Following the approach used by (McQuistin et al., 2021) and (Khare et al., 2022) each email is coded for the WG list in which it occurs, the identify of the sender, and their organisational role.

We distinguish three organisational roles ordered according to their level of social power in the organisation. *Current* for the chair of the workgroup list; *Allo* for a chair of a different workgroup who is sending on the list; and *None* for a contributor who has not been (at the time of sampling) a WG chair.

The analysis uses three Linguistic Inquiry and Word Count (LIWC (Boyd et al., 2022)) categories as dependent variables to index potentially sensitive language use:

1. **Politics:** words commonly used in political discussions (e.g., congress, parliament, president, democratic) or legal (court, law) discourse.

2. **Ethnicity:** words that identify national, regional, linguistic, ethnic, or racial identities.[4]

3. **Religion:** use of religious words such as "church, altar, god, christmas, hell, mosque, temple".[5]

Importantly, LIWC categories encode only mentions of specific words; context is not considered, with no attempt to disambiguate words that might have alternative senses in this technical domain. This means that, for example, questions about European IP address assignment policies and comments about British humour both count as LIWC *Ethnicity*. We do not attempt to separate these cases.

3 Results

Statistical tests of the hypotheses are performed using Generalised Linear Mixed Model (GLMM) analyses with Organisational Role (Current vs. Allo vs. None) as a fixed factor. For all three LIWC categories the raw data are positively skewed so a Gamma Distribution is used.

GLMM analysis of the LIWC *Politics* measure shows an overall main effect of Role ($F_{2,790} = 16.12, p < 0.001$). Pairwise comparisons show that *None* are reliably more likely than either *Allo* or *Current* to use political language (*Current* vs. *None*: $t_{(790)} = -6.13, p < 0.001$, *Allo* vs. *None*: $t_{(790)} = -2.83, p = 0.010$). Current and Allo are not reliably different ($t_{(790)} = 1.72, p = 0.086$).

GLMM analysis of LIWC *Ethnicity* also shows a main effect of Role ($F_{2,227} = 25.7, p < 0.001$) and a similar pattern to political language. Pairwise comparisons show that *None* use more Ethnicity related language than either *Allo* or *Current*: (*Current* vs. *None*: $t_{(227)} = -3.89, p < 0.001$, *Allo* vs. *None*: $t_{(227)} = -4.478, p = 0.010$). Current and Allo are not reliably different ($t_{(227)} = 0.93, p = 0.354$)

[1] `https://mailarchive.ietf.org/`

[2] `https://datatracker.ietf.org/` – administrative database of IETF

[3] See both `https://www.ietf.org/about/note-well/` and the IETF privacy policy available at `https://www.ietf.org/privacy-statement/`. IETF leadership confirmed that our work conforms with acceptable use.

[4] Words reflecting racial or ethnic slurs are generally excluded and included as part of the LIWC swear category.

[5] This list is not exhaustive but no clear definition is provided in the LIWC documentation so we take some of this on faith [sic].

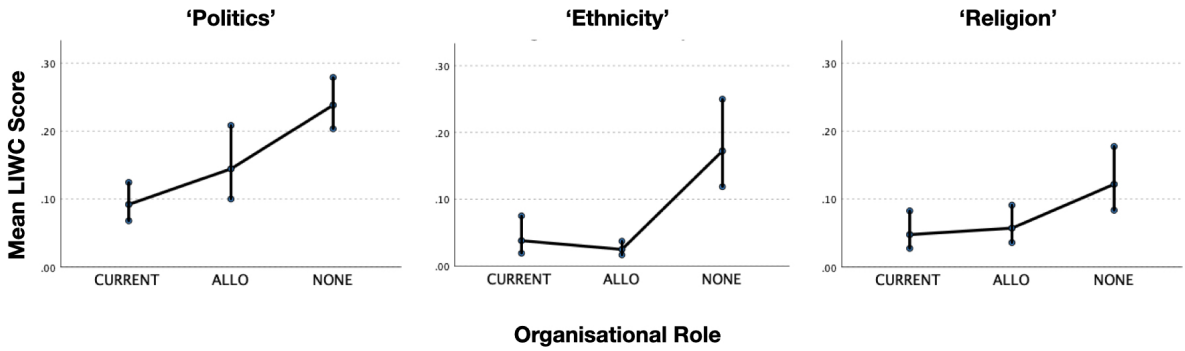

Figure 1: Language Use by Role: Estimated Marginal Means

GLMM for LIWC *Religion* shows an overall effect of Role ($F_{2,408} = 5.12, p = 0.006$). The pairwise comparisons between *None* and *Current* ($t_{(408)=-2.76,p=0.018}$) and *None* and *Allo* ($t_{(408)=-2.39,p=0.03}$) are reliable.

The results for all three analyses are illustrated in Figure 1. The consistent pattern is that people with organisational responsibilities (*Current* or *Allo*) are substantially less likely to use language potentially connected with sensitive topics than those without organisational responsibilities, even though they are all part of the same email discussions.

4 Conclusion

In absolute terms, there is very little use of potentially sensitive words. The modal LIWC score in both raw data and the aggregated figures is zero. Even in the subset of emails with non-zero LIWC scores they are only fractions of a percentage (Figure 1).

Despite the low overall incidence, and the context insensitive character of the LIWC categories, there are systematic differences in the distribution of potentially sensitive words. People who have more social power within the IETF, i.e. WG chairs, make much less use these words. Given that these interactions occur on shared mailing lists, this also implies that WG chairs also do not reciprocate or align on these expressions when they encounter them. Overall, this pattern is consistent with the Exposure hypothesis but conflicts with the Power hypothesis.

The observed effects are likely to depend on the kinds of accountability built in to a system of social power and the transparency of the decision making process. The IETF is characterised by a relatively open, egalitarian, consensus building culture that may especially promote a more careful, considered approach to organisational communication.

Although positions of social power are naturally thought of as positions of strength, the results suggest that should also be considered as positions of vulnerability and this can be reflected in patterns of language use.

References

Ryan L Boyd, Ashwini Ashokkumar, Sarah Seraj, and James W Pennebaker. 2022. The development and psychometric properties of liwc-22. *Austin, TX: University of Texas at Austin*, pages 1–47.

Penelope Brown, Stephen C Levinson, and Stephen C Levinson. 1987. *Politeness: Some universals in language usage*, volume 4. Cambridge university press.

Prashant Khare, Mladen Karan, Stephen McQuistin, Colin Perkins, Gareth Tyson, Matthew Purver, Patrick Healey, and Ignacio Castro. 2022. The web we weave: Untangling the social graph of the IETF. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, pages 500–511.

Stephen McQuistin, Mladen Karan, Prashant Khare, Colin Perkins, Gareth Tyson, Matthew Purver, Patrick Healey, Waleed Iqbal, Junaid Qadir, and Ignacio Castro. 2021. Characterising the IETF through the lens of RFC deployment. In *Proceedings of the 21st ACM Internet Measurement Conference*, pages 137–149.