

# Talk it up or play it down?

## (Un)expected correlations between (de-)emphasis and recurrence of discussion points in consequential U.S. economic policy meetings

Chenhao Tan\*

Dept. of Computer Science  
University of Washington  
<https://chenhaot.com>

Lillian Lee

Dept. of Computer Science  
Cornell University  
<http://www.cs.cornell.edu/home/llee>

Presented at Text as Data, Oct 14-15, 2016

Comments welcome; please send to [chenhao@chenhaot.com](mailto:chenhao@chenhaot.com) and [llee@cs.cornell.edu](mailto:llee@cs.cornell.edu)

### Abstract

In meetings where important decisions get made, what items receive more attention may influence the outcome. We examine how different types of rhetorical (de-)emphasis — including hedges, superlatives, and contrastive conjunctions — correlate with what gets revisited later, controlling for item frequency and speaker. Our data consists of transcripts of recurring meetings of the Federal Reserve’s Open Market Committee (FOMC), where important aspects of U.S. monetary policy are decided on. Surprisingly, we find that words appearing in the context of hedging, which is usually considered a way to express uncertainty, are more likely to be repeated in subsequent meetings, while strong emphasis indicated by *superlatives* has a slightly negative effect on word recurrence in subsequent meetings. We also observe interesting patterns in how these effects vary depending on social factors such as status and gender of the speaker. For instance, the positive effects of hedging are more pronounced for female speakers than for male speakers.

## 1 Introduction

Meetings play a crucial role in a wide range of settings, including collaboration, negotiation and policy decisions (Jarzabkowski and Seidl, 2008). For example, the U.S. Federal Open Market Committee (FOMC), “the monetary policymaking body of the Federal Reserve System”,<sup>1</sup> “holds eight regularly scheduled [six-hour] meetings per year [where it] reviews economic and financial conditions, determines the appropriate stance of monetary policy, and assesses the risks to its long-run goals of price stability and sustainable economic growth”; its decisions can “ultimately [affect] a range of economic variables, including employment, output,

---

\*Work done while the author was at Cornell University.

<sup>1</sup> [https://www.federalreserve.gov/faqs/about\\_12844.htm](https://www.federalreserve.gov/faqs/about_12844.htm)

and prices of goods and services”.<sup>2</sup> Studies of the FOMC’s meetings or using FOMC meeting transcripts as data include Meade (2005); Meade and Stasavage; Schonhardt-Bailey (2013); Guo, Blundell, Wallach, and Heller (2015); Zirn, Meusel, and Stuckenschmidt (2015); Hansen, McMahon, and Prat (2015).

## 1.1 Hedging “versus” superlatives

A central question for each meeting participant is how to make his or her arguments noted and valued by other participants, and thus ultimately influence the outcome of the meeting. Our interest in this paper is in the effectiveness of certain subtle presentational or rhetorical options in this regard — specifically, whether a speaker attempts to make a point using certain vs. uncertain language. Here is an example taken from the March 22, 2005 FOMC meeting. The speaker is identified in the transcript as Ms. Minehan,<sup>3</sup> President of the Federal Reserve Bank of Boston, and she is discussing an alternative wording:<sup>4</sup>

I’m also concerned in alternative B about the rise in energy prices not notably feeding through to core consumer prices. Core consumer prices are up a full percentage point on a year-over-year basis, and there has been some feed-through. We think it’s going to slacken, and *maybe you want to put that reference in the future, but **I’m not sure** that this is what we want to say in this statement.* I think we’d be better off leaving that sentence out and just going with “pressures on inflation have picked up in recent months and pricing power is more evident.”

The italicized sentence contains the highlighted hedges “maybe” and “I’m not sure”. Notice that Minehan could have uttered a more invested or committed version of this sentence that omits the expressions of uncertainty:

... and you could put that reference in the future, but this is not what we want ... (1)

Also, she could have made the point using superlative language for emphasis:

... this is the **worst** wording we could possibly pick. (2)

Would one of these choices have been more effective than the others in causing the committee members to seriously consider Minehan’s proposals?

## 1.2 Why hedging?

It may at first seem strange to choose the “emphasis” aspect of wording as a focal point. One objection runs as follows: besides wording, there are many other, perhaps more salient factors at play, such as status, social relationships, shared history, charisma, timing, and so on (Cialdini, 2009), not to mention the validity or “correctness” of the content of an argument itself (Petty and Cacioppo, 2012). However, the “omnipresence” of the idea of *framing* “across the social sciences and humanities” means that there is a great deal of scholarly interest in how speakers and authors can, *often through language*, “select some aspects of a perceived reality and make them more salient” for persuasive ends (Entman, 1993). Moreover, we have argued elsewhere that *how someone says something is one of the few factors that a speaker has some control over when he or she seeks to convey a fixed piece of content*:

---

<sup>2</sup> <https://www.federalreserve.gov/monetarypolicy/fomc.htm>

<sup>3</sup> The presence of “Ms.” and “Mr.” notations in the transcripts mean that we can easily extract gender information, a fact we take advantage of in our experiments.

<sup>4</sup> We acknowledge the meta-ness of including as an example in a paper about choices of wording a case where people are discussing choices of wording.

For example, consider a speaker at the ACL [a scientific organization's] business meeting who has been tasked with proposing that Paris be the next ACL location. This person cannot on the spot become ACL president, change the shape of his/her social network, wait until the next morning to speak, or campaign for Rome instead; but he/she can craft [their] message ... (Tan, Lee, and Pang, 2014)

We thus assert that it is both an interesting scientific question and an interesting pragmatic question to ascertain whether language aspects of delivery have an effect on the degree of influence one has, independent of non-linguistic factors.

Our particular interest in this paper in looking at employment of expressions of uncertainty arises from how fascinating the phenomenon is in its own right (see, for example, Schröder and Zimmer (1997) for a listing of perhaps hundreds of papers on the topic up to 1997, and Farkas, Vincze, Móra, Csirik, and Szarvas (2010) for a description of entrants to a shared task/competition among NLP systems for identifying uncertainty). After all, the fact that hedging exists is seemingly odd: one might first think that if people want communication to be direct and efficient, shouldn't they just cut out the extra verbiage that hedging entails? And, don't hedges make a speaker or a speaker's position seem weak? Public-speaking advice on the Internet cautions people to avoid them, and indeed, Strunk and White themselves state: "Avoid tame, colorless, hesitant, non-committal language."

But in fact hedging can be a tool for a speaker to achieve his or her aims. Consider the following excerpt from the March 22, 2005 FOMC meeting, where Kos hedges much more than Greenspan, the chair:

*Greenspan: **I assume** iron ore is in [the CRB]?* (3)

*Kos: I don't know if iron ore is in there but copper is: copper scrap is in there, **I think**.*

*Greenspan: That couldn't have done that much. Steel, for example, is actually down.*

*Kos: **I don't think** steel is in the CRB.*

Importantly, Kos's corrections of Greenspan are accurate: according to Thomson Reuters<sup>5</sup>, the CRB index contains copper but not iron ore or steel. Furthermore, Kos is presumably not actually uncertain of these facts. Rather, it would seem that Kos is softening his language to either (1) make his assertions more palatable or acceptable, or (2) trying to signal respect while contradicting the higher-status Chair.

### 1.2.1 Why the FOMC?

So far, we have not mentioned anything about language usage that seems particularly specific to economic policy discussion. But the FOMC meetings are a particularly nice domain for our empirical work because we might expect language effects to be minimized:

- The stakes are very high, since the decisions made by the FOMC are extremely consequential. Thus, one might argue that the participants would be highly motivated to focus on the content, not the wording, of the discussions.
- The participants are high-status experts in the field, and hopefully respect each other to at least some extent. One might therefore suppose that they would be less inclined to either require expressions of social deference to each other or be impressed by undeserved emphaticism, especially as the meetings wear on over multiple hours.

---

<sup>5</sup> <http://financial.thomsonreuters.com/content/dam/openweb/documents/pdf/financial/core-commodity-crb-index.pdf>

- At least some of the participants have interacted a great deal with each other, which might reduce the influence that language choices would have on how people’s suggestions are received by each other.

Hence, since the situation reduces the possibility for language choices to have an effect, any effects that we *do* see deserve consideration.

Moreover, some other experimentally convenient features are (a) the positions (job descriptions) and genders of the participants are indicated in the transcripts; and (b) pre-1993, the FOMC members were not aware that the transcripts would be made public — this fact dampens the possibility that the participants were speaking unnaturally or trying to direct their comments towards a broader audience. Other characteristics we have not exploited in this work but could be useful for other research include (c) many speakers participate in many meetings, providing relatively plentiful user-specific data; (d) there is a great deal of public documentation laying out the basis on which decisions are made and what is being decided upon, such as the Bluebooks, Greenbooks, and so on;<sup>6</sup> (e) Meade (2005) provides manually-assigned disagreement labels which indicate who argued against — not just cast a dissenting vote against — the final decision for each meeting, which may be interesting for future studies.

We intend to make our processed versions of the transcripts publicly available.

### 1.3 A repetition framework for investigation

While we would like to study whether hedging and other forms of (de-)emphasis have detrimental or positive effect on the reception of a speaker’s ideas, it can be difficult to ascertain (computationally or otherwise) whether the listeners give those ideas serious consideration or not.

We therefore employ the following computationally convenient proxy for idea uptake: *repetition* or *echoing*, inspired by Niederhoffer and Pennebaker (2002); Danescu-Niculescu-Mizil, Gamon, and Dumais (2011); Danescu-Niculescu-Mizil, Lee, Pang, and Kleinberg (2012). (See also the definition of “discussion points” of Zhang, Kumar, Ravi, and Danescu-Niculescu-Mizil (2016).)

Fig. 1 demonstrates the main idea of how we construct the specific data-points for our study. For a given context, such as “expressions of uncertainty” or “superlatives”, we find instances of the occurrence of the context in individual speeches.<sup>7</sup> Then, we pair a word from the speech appearing outside the context with a word from the speech appearing in the context word or phrase, taking care that the “in”-word has the same frequency prior to the speech as the “out”-word.

We then ask, how frequently does the out-of-context word occur *after the speech*, in comparison to the in-context word — we thus use “a word is used by other people” as a rough proxy for “other people are paying attention to the underlying concept”. The null hypothesis, given that both words are of equal prior probability and uttered by the same speaker at almost the same time, is that context will have no effect and the words will continue to have roughly equal frequency in the future. By definition this framework controls for important factors other than the phrasing, such as who the speaker is and when in the meeting does the speech happen.

Note that our framework also allows for flexibility in measuring how well statements are received: using repetition as the indicator of influence is not central to our setup.

<sup>6</sup>[https://www.federalreserve.gov/monetarypolicy/fomc\\_historical.htm](https://www.federalreserve.gov/monetarypolicy/fomc_historical.htm)

<sup>7</sup>By “speech”, we mean an uninterrupted span of speech by a single speaker.

Mr. Moskow: ... Auto and light truck sales appear to be coming in at about the 14-1/2 million units level so far in May, which is approximately 3/4 million units above the April pace but still well below the *expectations* earlier this year. [...] On the employment front, labor markets remain tight, with the District’s *unemployment* rate at its lowest level in over 15 years. ...

Figure 1: Example of a matching word pair (*unemployment*, *expectations*) for the context of *superlatives* (“lowest” highlighted and underlined in red). They have been uttered by the same speaker, and are of similar frequency in our dataset before this speech; and “unemployment” occurs in the context of “lowest”, while “expectations” does not.

### 1.4 Highlights reel

Our first contribution is the repetition-based speaker- and time-controlled framework we introduced in the previous subsection.

We look for hedging/emphasis-mediated repetition effects in the FOMC meeting transcripts both within the same meeting (intra-meeting) and in subsequent meetings (inter-meeting). One surprising finding is that although hedges have very little effect within the same meeting, words in the context of hedges are more likely to occur in subsequent meetings.

Furthermore, we investigate how these effects may vary depending on other factors such as status and gender of the speaker. One interesting finding is that the effect of context is more pronounced for female speakers. This echoes existing work that suggests that female speakers are more persuasive in an indirect manner (Burgoon, Jones, and Stewart, 1975).

## 2 Analysis framework

In order to understand the effect of (de-)emphasis on the reception — where here we approximate “reception” by “repetition” — of a speaker’s ideas, we develop a framework that controls for important confounding factors.

Throughout, we use the term *context* to refer generically to a class of (de-)emphasis techniques.

The intuition is shown in Fig. 1: for a given context, within the same speech — so that both the speaker and the meeting state are naturally controlled — we identify sentences containing an instance of that context versus sentences that do not contain any instances of that context. We then extract “similar” (in-context, out-of-context) word pairs based on their frequency in the past, and compare their frequency of occurrence later in the meeting or in subsequent meetings.

**Formal definition.** A context  $C$  is defined as a set of words or phrases. Within the same speech  $S$ , we define sentences that contain any  $c \in C$  as  $Sent_{in}$ , and the other sentences as  $Sent_{out}$ . We match content words of similar past frequency in  $Sent_{in}$  and  $Sent_{out}$ ,<sup>8</sup> and define the set of matched pairs for speech  $S$  as  $MP_C(S)$ :

$$MP_C(S) \stackrel{def}{=} \{(w_{in}, w_{out}, S) \mid w_{in} \in Sent_{in} - Sent_{out}, w_{out} \in Sent_{out} - Sent_{in}, PF_S(w_{in}) \sim PF_S(w_{out})\}, \tag{4}$$

where  $PF_S$  gives the past frequency of a word. Some further details on the finer points of refining the definition of  $MP_C$  are given in the appendix, §A.

<sup>8</sup>We exclude stopwords and words in all matching pairs, although the results are robust even if we include stopwords in the pairing process.

Next, we define the *effect* of a context by measuring the difference between the probability that words in the context are echoed *more* in the future than out-of-context words and a default prediction of 0.5, since the words in our same-speech pairs have similar past frequency:

$$E(C) = \widehat{P}_C - 0.5, \quad (5)$$

where  $\widehat{P}_C$  computes the probability that words in the context are repeated more in the future than words out of the context. Specifically, we compute the average winning rate of  $w_{in}$  in  $MP_C(S)$  for each speech  $S$  and then average over all speeches:  $\widehat{P}_C = \frac{1}{|\mathcal{S}|} \sum_{S \in \mathcal{S}} \frac{1}{|MP_C(S)|} \sum_{w_{in}, w_{out} \in MP_C(S)} I(FF_S(w_{in}) > FF_S(w_{out}))$ . Here  $FF_S$  gives the frequency of a word in *other meeting participants' speeches* after  $S$  either within the same meeting or in subsequent meetings. The definition of  $FF_S$  can vary depending on research hypotheses that we are interested in. We will present two classes of  $FF_S$  in §4.

*In our experiments, we will be concerned with whether the effect defined by Equation (5) is different from zero.* A positive effect suggests that the context is associated with more future echoing, while a negative effect suggests less.

### 3 Hypotheses

Our main interest in this work is to examine the effect of (de-)emphasis on the reception of a speaker's ideas. In addition to *hedges* and *superlatives*, we also investigate two other common contexts that can be associated with emphasis: *contrastive conjunctions* and *second person pronouns*. In the following, we develop our hypotheses based on existing studies and our intuitions.

H1: *Hedges*. We have already discussed some intuitions and prior work regarding hedging in the Introduction. Moreover, Durik, Britt, Reynolds, and Storey (2008) show that “hedges can, but do not always, undermine persuasive attempts” and Erickson, Lind, Johnson, and O’Barr (1978) show that powerless language results in lower perceived credibility of the witness in court trials. In light of these studies, we expect a *negative* effect within a meeting. We merge and manually curate several data sources to get a list of hedges (Farkas et al., 2010; Hanauer, Liu, Mei, Manion, Balis, and Zheng, 2012; Hyland).<sup>9</sup>

<sup>10</sup>

H2: *Superlatives*. As these are the strongest form to describe a fact or an action and can place an emphasis on the statement<sup>11</sup>, we expect a *positive* effect.

H3: *Contrastive conjunctions*. A contrastive conjunction like “but” places an emphasis on the text after its occurrence, so we expect a *positive* effect.

H4: *Second person pronouns*. Although using *second person pronouns* (“you”) is not a form of emphasis, it can likely attract the attention of the addressed speaker. We expect a *positive* effect shortly after the speech as these are direct mentions of other meeting participants.

<sup>9</sup>We focus on a subset of hedges where the speaker may try to shield the responsibility of a statement. For example, “to be raised” and “or” from Farkas et al. (2010) are not included.

<sup>10</sup>It should be pointed out that the automatic identification of hedging and expressions of uncertainty is not a solved problem (Farkas et al., 2010). Items that seem like hedge cues may turn out always be so in real-life usage (compare “I **\*think\*** it’s going to rain” with “**\*I\*** think it’s going to rain”); and, hedging can occur without well-recognized hedge cues (“I’m no Albert Einstein, but I say the answer is 1234.”).

<sup>11</sup>This sentence itself contains a superlative: the word “strongest”.

Num. speakers	Num. speeches	Num. tokens
26.93	404.87	37652.40

Table 1: Per-meeting averages in our dataset.

context	example pairs
<i>hedges</i>	( <i>adjustment</i> , <i>capacity</i> ), ( <i>mortgage-related</i> , <i>high-yielding</i> )
<i>superlatives</i>	( <i>sector</i> , <i>employment</i> ), ( <i>information</i> , <i>judgment</i> )
<i>contrastive conjunctions</i>	( <i>money</i> , <i>point</i> ), ( <i>export</i> , <i>signal</i> )
<i>second person pronouns</i>	( <i>disturbance</i> , <i>shipments</i> ), ( <i>spreads</i> , <i>reductions</i> )

Table 2: Example pairs for different contexts. The first element in each pair is the in-context word; the second is the outside-context word. Recall that these are words spoken by the *same* speaker at *about the same time*.

H5: No lasting effects. We expect that so much time passes between meetings and the word choices we are looking at are sufficiently subtle (for instance, the addition of the phrase “I think”) ’that there should be no effects lasting from one meeting to the next.

## 4 Dataset

Our dataset is drawn from the transcripts of all FOMC meetings from 1977 to 2008. Table 1 presents basic statistics.

In order to apply our framework, we define the *past frequency* of a word ( $PF_S$ ) with respect to its appearance in a speech  $S$  as the log probability of the word in the previous meetings, and employ two classes of functions to measure the *future frequency* of a word:

- Intra-meeting frequency. We split the speeches after  $S$  into windows of five speeches, and then compute the log probability of a word within each window after  $S$  for 20 windows (100 speeches after  $S$ ). We expect the effect of a context to fade away as the meeting moves forward, whether that effect is positive or negative.
- Inter-meeting frequency. In order to assess the effect of (de-)emphasis in subsequent meetings, we compute per-meeting log probability of the word for each of the five subsequent meetings after  $S$ .

Recall that we compare the future frequency of prior-frequency-controlled (in-context, outside-context) pairs. Table 2 presents two matching pairs of words randomly chosen from our pairs data for each of the four contexts that form the foci of our hypotheses. Indeed, it is non-trivial to guess which word, if any, will be echoed significantly more in the future a priori.

As a preliminary experiment, we used the method of Monroe, Colaresi, and Quinn (2008) to compare the words tending to appear within each type of context with the words tending to appear outside each type of context. We omit detailed results here, but in general, the differences match our intuitions. For example, *hedges* tend to occur with evaluative statements (“ought”, “risks”, “important”), while “thank” and “chairman” typically occur out of context, because a typical phrase to start a speech in these meetings is “Thank you, Mr. Chairman.” But, recall that we purposely constructed pairs to have equal prior probability, which should help mitigate any effects stemming merely from what words tend to occur in a given context.




Figure 2: The first row presents the effect of contexts with the same meeting over 100 speeches (20 windows) after speech  $S$ ; the second row presents the effects over five subsequent meetings. Dotted red line: the best linear fit of the effect over the x-axis. Dotted gray line: the null hypothesis where the context has no effect. x-axis: the number of speeches after the speech of interest in intra-meeting plots, the subsequent x-th meeting in inter-meeting plots. Bars: standard error. We use the same x-axis, y-axis, line styles, and bars in all intra-meeting and inter-meeting figures.

## 5 Effects of contexts

We first examine the overall effects of contexts, and then explore how the effects vary across different factors, including status, gender, and speech length.

### 5.1 Overall effects (Fig. 2)

**Intra-meeting effects (Hypotheses H1-H4).** We expected *hedges* (Hypothesis H1) to have a negative effect within the same meeting, given that material that people express uncertainty about might tend to receive less attention from the other participants. However, *hedges* seem to only have a small negative effect right after the speech and the effect quickly returns to 0.<sup>12</sup>

We also expected that *second person pronouns*, *contrastive conjunctions* and *superlatives* (Hypotheses H2, H3, and H4) would have a positive effect shortly after the speech of interest. However, *superlatives* seem to have a slightly negative effect, while *contrastive conjunctions* do not have much effect. In contrast, perhaps because *second person pronouns* directly mention other participants, they demonstrate a strong positive effect, although, not surprisingly, the effect diminishes over the course of the meeting.

**Inter-meeting effects (Hypothesis H5)** In contrast with intra-meeting results, surprisingly, *hedges* correlate with a positive effect in the subsequent meetings. This suggests that expressing uncertainty correlates with a better reception of ideas in the long run indicated by repetition.

<sup>12</sup>Those of us who find ourselves tending to hedge may view this as an unexpectedly positive finding.



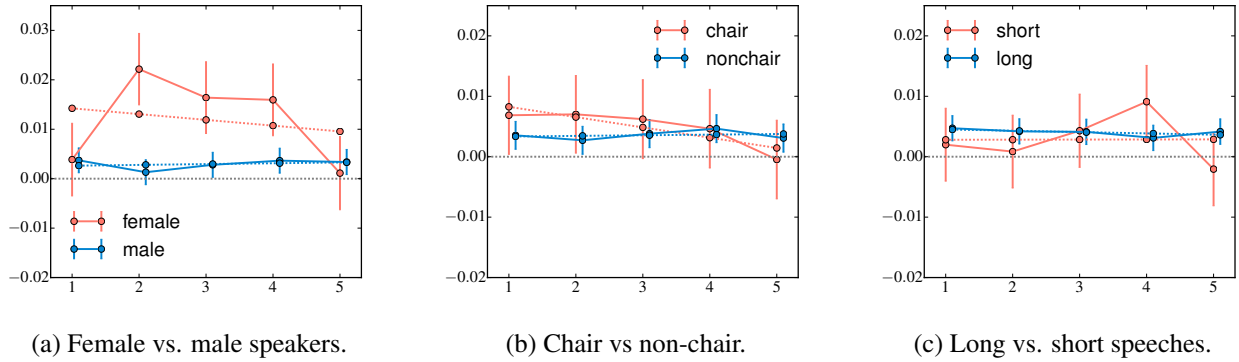


Figure 3: Inter-meeting comparisons of the effect of hedging across speaker status, gender and speech length. Note that the y-axis scale for gender comparison is different from the other two.

Consistent with intra-meeting results, *superlatives* present a negative effect on whether words are going to be repeated in subsequent meetings. Although *contrastive conjunctions* present zero effect in the next several subsequent meetings, they lead to slightly more pronounced negative effect in later subsequent meetings. Finally, the effect of *second person pronouns* mostly overlaps with the zero line (in fact, it is very similar to the random case shown in §A).

## 5.2 Impact of different factors

Despite the above aggregate results, the effect of a context may not be homogeneous conditioned on other factors, such as status (whether the speaker is chair or not), gender (whether the speaker is male or female), and speech length (whether the speech is long or short). We explore these variations in the inter-meeting effect of hedging and in the intra-meeting effect of *second person pronouns*.

### 5.2.1 Hedges (inter-meeting, Fig. 3)

**Stronger positive effect in subsequent meetings for female speakers. (Fig. 3a)** The gender of each participant can be obtained by the prefix in the speaker name. We omitted all speeches made by the chairs to avoid the influence of status.<sup>13</sup> There is a clear positive effect in subsequent meetings for female speakers, while there is not much effect for male speakers. This echoes the findings in Burgoon et al. (1975) and Carli (2001) that female speakers are more considered persuasive when employing an indirect manner. (For an interesting critique of advice that women should speak “more like men”, see Cameron (1995).)

**Similar positive effect in subsequent meetings for speakers with different statuses. (Fig. 3b)** We use whether the speaker was the chair of FOMC as a proxy of status. As a result, the number of samples is much smaller for the chair group than the non-chair group and we thus observe a larger variance for the chairs. The effect of *hedges* for the chairs and non-chairs mostly overlap with each other, although the effect for the chairs seems to be slightly above that for non-chairs in the first several subsequent meetings.

**The positive effect in long speeches is more consistent. (Fig. 3c)** The final aspect that we examine is speech length. One may expect that for long speeches, it is more important to emphasize certain parts so that others can pick up. To distinguish long speeches from short speeches, we simply split the speeches where there are matching word pairs into two groups using the median as a boundary. The positive effect of *hedges* in subsequent meetings is consistent for long speeches, while it fluctuates more for short speeches.

<sup>13</sup>We tried to exclude Yellen and the same observation holds that there is stronger positive effect for female speakers.

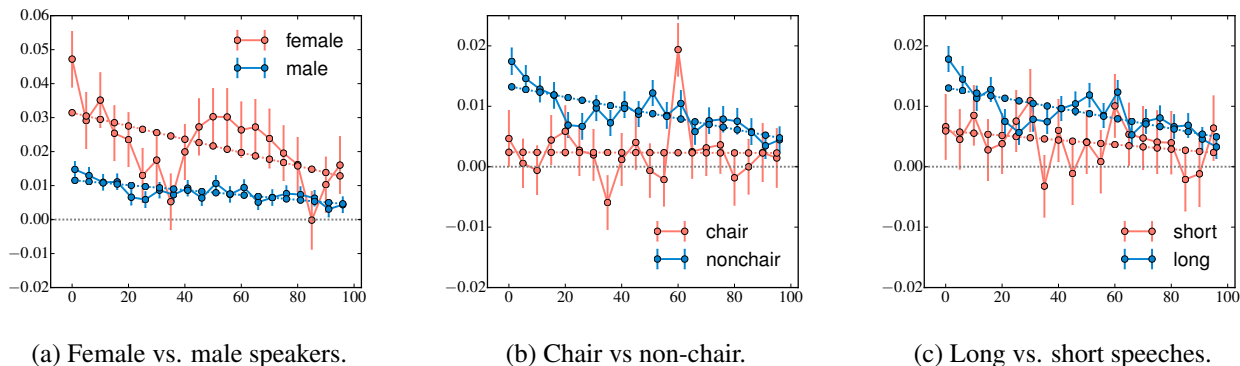


Figure 4: Intra-meeting comparisons of the effect of second-person pronouns across speaker status, gender and speech length. Note that the y-axis scale for gender comparison is different from the other two.

### 5.2.2 *Second person pronouns (intra-meeting, Fig. 4)*

We examine how the positive effect of *second person pronouns* within a meeting changes conditioned on status, gender and speech length. We follow the same procedures as above to extract status, gender and speech length information.

**Stronger positive effect for female speakers than male speakers. (Fig. 4a)** Surprisingly, the effect of *second person pronouns* is smaller for male speakers, in other words, *second person pronouns* spoken by female speakers present a stronger immediate effect on other participants after the speech. This may suggest that it is more important for female speakers to “ask” other participants to pay attention using certain contexts. This observation is consistent with the result in inter-meeting results for *hedges*: the positive effect of a context is more pronounced for female speakers.

**Stronger positive effect for speakers with lower status than speakers with higher status. (Fig. 4b)** The effect of *second person pronouns* is mitigated in the chairs’ speeches (y-values fluctuate around 0). One way to interpret this observation is that meeting participants pay similar levels of attention to the chairs’ statements regardless of the second-person-pronoun context.

**Stronger positive effect in long speeches. (Fig. 4c)** The difference between long speeches and short speeches is clearer than in §5.2.1. The effect of *second person pronouns* is stronger in long speeches.

## 5.3 Further caveats and disclaimers

We do not claim that correlation implies causation. In particular, these findings should not be viewed as positive advice on how to influence discussion.

There are some aspects of the data that we do not directly take into account in the experiments reported in this paper. There are changes over time in the style and leadership of the meetings. For instance, the number of speeches is decreasing over time. Also, after 1993, the meeting participants were aware of the fact that the transcripts would be made publicly available.

## 6 Related work

FOMC meetings have attracted significant research interests. Rosa (2013) shows that the release of FOMC minutes significantly affects the volatility of U.S. asset prices and trading volume. See Schonhardt-Bailey




Figure 5: Validation of the proposed metric using a random context ( $p = 0.05$ ).

(2013) for a comprehensive textual analysis. The focus of our study is on the effect of subtle rhetorical correlates within the meetings.

Another related line of work is accommodation and linguistic style matching (Danescu-Niculescu-Mizil and Lee, 2011; Niederhoffer and Pennebaker, 2002), which study the phenomenon of people matching each other in conversations. Here we attempt to study how subtle presentational and rhetorical (de-)emphasis may influence the reception of a speaker’s ideas and evaluate based on content words, in contrast with functional words to capture style.

Additionally, there have been other studies in the natural-language processing and computational literature of correlations between language and persuasiveness (Guerini, Strapparava, and Stock, 2008; Mitra and Gilbert, 2014; Guerini, Ozbal, and Strapparava, 2015; Tan, Niculae, Danescu-Niculescu-Mizil, and Lee, 2016; Cano-Basave and He, 2016). Hedging was one of the features examined by Tan et al. (2016).

## 7 Conclusion

In this paper, we took advantage of “natural experiments” in the same speech within meetings and proposed a computational framework for measuring the effects of subtle presentational and rhetorical (de-)emphasis. We applied our framework in FOMC meetings and found surprising patterns, including a positive effect of de-emphasis indicated by hedging. Furthermore, we demonstrated how the effect of hedging is more pronounced for speakers female speakers. This work is one step towards quantitatively understanding the effect of wording on social dynamics in meetings. This general idea of looking at words or phrases in the same speech can spur new computational frameworks to measure the influence of language.

## 8 Acknowledgments

We first learned of the availability of FOMC meeting transcripts from Cheryl Schonhardt-Bailey at the 2010 Text as Data meeting at Northwestern! We thank Bitsy Perlman, Cheryl Schonhardt-Bailey, and the 2016 Text as Data attendees for helpful comments. This work was supported by a Facebook fellowship and in part by a University of Washington Innovation Award.

## A Appendix: notes on pairing in-context vs. out-of-context words ( $MP_C$ from section 2)

Our framework considers “natural experiments” using word pairs drawn from the same speech of the same speaker. However, there can be many intricate design choices in defining  $MP_C(S)$  (Equation 4) that may affect the measurement of  $E(C)$ . These choices include whether to control the frequency of paired words within the speech, the part-of-speech tag of paired words, etc. Therefore, we use a “random” context to validate these choices.

In order to generate a “random” context, we toss a coin for each word position in the speech with probability  $p$  to decide whether this word position is a context cue.<sup>14</sup> *Since the context is randomly selected, we expect our metric to be around 0.* If the observed effects are different from 0, it suggests that there exists some systematic bias in the design choices.

Figure 5 presents the results regarding whether to control the frequency of paired words in the same speech. Surprisingly, given that we have already controlled for past frequency of paired words, it remains important to control for the number of times a word occurs in the speech. Without controlling in-speech frequency, the effect is biased towards the negative side, which could have led to the wrong conclusion that a random context has negative effects on future re-occurrences of words.

We also explore other design choices that can potentially influence the metric: 1) where the speech happens in the meeting (meetings may have different stages and contexts may provide different effects in the middle of a meeting compared to in the beginning of a meeting); 2) part-of-speech tags of paired words (contexts may have different effects on words of different part-of-speech tags). These two factors did not affect our metric. Therefore, in the following results, we enforce that paired words have the same frequency within the same speech.

## References

- Michael Burgoon, Stephen B. Jones, and Diane Stewart. Toward a Message-centered Theory of Persuasion: Three Empirical Investigations of Language Intensity. *Human Communication Research*, 1(3):240–256, 1975.
- Deborah Cameron. The new Pygmalion: Verbal hygiene for women. In *Verbal Hygiene*, pages 166–211. Routledge, 1995. URL <http://site.ebrary.com/id/10100241>.
- Amparo Elizabeth Cano-Basave and Yulan He. A study of the impact of persuasive argumentation in political debates. In *Proceedings of NAACL*, pages 1405–1413, June 2016.
- Linda L. Carli. Gender and social influence. *Journal of Social Issues*, 57(4):725–741, 2001.
- Robert B. Cialdini. *Influence: Science and Practice*. HarperCollins, 2009.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In *Proceedings of the 2Nd Workshop on Cognitive Modeling and Computational Linguistics*, 2011.
- Cristian Danescu-Niculescu-Mizil, Michael Gamon, and Susan Dumais. Mark my words! Linguistic style accommodation in social media. In *Proceedings of WWW*, 2011.

---

<sup>14</sup>We considered  $p = 0.05$  and  $p = 0.5$ . The trends are similar.

- Cristian Danescu-Niculescu-Mizil, Lillian Lee, Bo Pang, and Jon Kleinberg. Echoes of power: Language effects and power differences in social interaction. In *Proceedings of WWW*, pages 699–708, 2012.
- Amanda M. Durik, M. Anne Britt, Rebecca Reynolds, and Jennifer Storey. The effects of hedges in persuasive arguments: A nuanced analysis of language. *Journal of Language and Social Psychology*, 27(3): 217–234, 2008.
- Robert M. Entman. Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, 43(4):51–58, 1993.
- Bonnie Erickson, E. Allan Lind, Bruce C. Johnson, and William M. O’Barr. Speech style and impression formation in a court setting: The effects of “powerful” and “powerless” speech. *Journal of Experimental Social Psychology*, 14(3):266 – 279, 1978.
- Richárd Farkas, Veronika Vincze, György Móra, János Csirik, and György Szarvas. The CoNLL-2010 shared task: Learning to detect hedges and their scope in natural language text. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning—Shared Task*, pages 1–12, 2010.
- Marco Guerini, Carlo Strapparava, and Oliviero Stock. Trusting politicians’ words (for persuasive NLP). *Proceedings of CICLing*, pages 263–274, 2008.
- Marco Guerini, Gözde Ozbal, and Carlo Strapparava. Echoes of persuasion: The effect of euphony in persuasive communication. In *Proceedings of NAACL*, pages 1483–1493, 2015.
- Fangjian Guo, Charles Blundell, Hanna Wallach, and Katherine Heller. The Bayesian echo chamber: Modeling social influence via linguistic accommodation. In *Proceedings of AISTATS*, pages 315–323, 2015.
- David A. Hanauer, Yang Liu, Qiaozhu Mei, Frank J. Manion, Ulysses J. Balis, and Kai Zheng. Hedging their bets: the use of uncertainty terms in clinical documents and its potential implications when sharing the documents with patients. In *AMIA Annual Symposium Proceedings*, 2012.
- Stephen Hansen, Michael McMahon, and Andrea Prat. Transparency and deliberation within the FOMC: A computational linguistics approach. [https://www2.warwick.ac.uk/fac/soc/economics/staff/mfmcmahon/research/fomc\\_submission.pdf](https://www2.warwick.ac.uk/fac/soc/economics/staff/mfmcmahon/research/fomc_submission.pdf), 2015.
- Ken Hyland. *Hedging in scientific research articles*. Pragmatics and Beyond New Series.
- Paula Jarzabkowski and David Seidl. The role of meetings in the social practice of strategy. *Organization Studies*, 29(11):1391–1426, 2008.
- Ellen E. Meade. The FOMC: Preferences, voting, and consensus. *Federal Reserve Bank of St. Louis Review*, 87(2):93–101, 2005.
- Ellen E. Meade and David Stasavage. Publicity of debate and the incentive to dissent: Evidence from the US Federal Reserve, journal = *The Economic Journal*, year = 2008, volume = 118, number = 528, pages = 695–717.
- Tanushree Mitra and Eric Gilbert. The language that gets people to give: Phrases that predict success on kickstarter. In *Proceedings of CSCW*, 2014.

- Burt L. Monroe, Michael P. Colaresi, and Kevin M. Quinn. Fightin' words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, 16(4):372–403, 2008.
- Kate G. Niederhoffer and James W. Pennebaker. Linguistic style matching in social interaction. *Journal of Language and Social Psychology*, 21(4):337–360, 2002.
- Richard E. Petty and John T. Cacioppo. *Communication and Persuasion: Central and Peripheral Routes to Attitude Change*. Springer Science & Business Media, 2012.
- Carlo Rosa. The financial market effect of FOMC minutes. *Economic Policy Review*, 19(2), 2013.
- Cheryl Schonhardt-Bailey. *Deliberating American Monetary Policy: A Textual Analysis*. MIT Press, illustrated edition, 2013.
- Hartmut Schröder and Dagmar Zimmer. Hedging research in pragmatics: A bibliographical research guide to hedging. In *Hedging in Discourse: Approaches to the analysis of a pragmatic phenomenon in academic texts*, Research in Text Theory, pages 249–271. De Gruyter, 1997.
- Chenhao Tan, Lillian Lee, and Bo Pang. The effect of wording on message propagation: Topic-and author-controlled natural experiments on twitter. In *Proceedings of the ACL*, pages 175–185, 2014.
- Chenhao Tan, Vlad Niculae, Cristian Danescu-Niculescu-Mizil, and Lillian Lee. Winning arguments: Interaction dynamics and persuasion strategies in good-faith online discussions. In *Proceedings of WWW*, pages 613–624, 2016.
- Justine Zhang, Ravi Kumar, Sujith Ravi, and Cristian Danescu-Niculescu-Mizil. Conversational flow in Oxford-style debates. In *Proceedings of NAACL (short papers)*, 2016.
- Ccilia Zirn, Robert Meusel, and Heiner Stuckenschmidt. Lost in discussion? Tracking opinion groups in complex political discussions by the example of the fomc meeting transcriptions. In *Proceedings of RANLP*, pages 747–753, 2015.