

Echoes of power: Language effects and power differences in social interaction

Cristian Danescu-Niculescu-Mizil
Cornell University

Lillian Lee
Cornell University

Bo Pang
Yahoo!

Jon Kleinberg
Cornell University

ABSTRACT

Understanding social interaction within groups is key to analyzing online communities. Most current work focuses on structural properties: who talks to whom, and how such interactions form larger network structures. The interactions themselves, however, generally take place in the form of natural language — either spoken or written — and one could reasonably suppose that signals manifested in language might also provide information about roles, status, and other aspects of the group’s dynamics. To date, however, finding such domain-independent language-based signals has been a challenge.

Here, we show that in group discussions, power differentials between participants are subtly revealed by how much one individual immediately echoes the linguistic style of the person they are responding to. Starting from this observation, we propose an analysis framework based on linguistic coordination that can be used to shed light on power relationships and that works consistently across multiple types of power — including a more “static” form of power based on status differences, and a more “situational” form of power in which one individual experiences a type of dependence on another. Using this framework, we study how conversational behavior can reveal power relationships in two very different settings: discussions among Wikipedians and arguments before the U. S. Supreme Court.

1. INTRODUCTION

With the arrival of detailed data on the interactions within social groups — generally coming from the on-line domain — an active line of research has developed around the phenomena taking place in these groups. To date, these analyses have mainly used structural features of the interactions, including who talks to whom, how frequently, and how these patterns of interaction form larger network structures.

But the interactions themselves are generally taking place in natural language — both spoken and written — and the language content of these interactions has been a long-acknowledged missing

ingredient in this style of investigation. The reason for this is clear: while it is reasonable to suppose that signals within the language could provide insight into the social structure of the group, it has been challenging to extract useful language-level signals. A small but growing line of work has begun to use textual content for uncovering structural properties of on-line networks, either through the analysis of its flow [49, 31, 11, 9] or by enabling finer grained link-type detection [16], and it is an important direction to extend the range of social properties that can be analyzed via text.

Power and linguistic style. In this paper, we show how variations in linguistic style can provide information about power differences within social groups. Our focus is on domains in which groups engage in goal-oriented discussions — situations where people interact, not necessarily collaboratively, in order to accomplish tasks or settle on choices. An important characteristic of such discussions is that the participants are invested in the issues at hand, so that their dialogs are not simply “idle chat”, but consequential: the outcome matters. Examples include conversations among wiki editors or open-source teams regarding modifications; debates within conference program committees on which papers to accept; and discussions in legal hearings, where opposing sides compete to persuade a judge or jury.

Power differences among the participants constitute a crucial force in all these settings. Sometimes these power differences are embodied in formal roles, such as that of a judge or a program chair. Sometimes they are based on more informal differences in the respect or authority commanded by individuals within the group. And sometimes they are more situational: x may have power over y in a given situation because y needs something that x can choose to provide or not.

It is natural to ask how we might try inferring these power differences when they are not explicitly marked, simply by observing the dynamics within a group. This is particularly challenging if we are seeking methods that generalize across domains, and are not tied to specific choices of content. By way of analogy, imagine that you walk into a meeting among people you’ve never met, and on a topic that you know nothing about; what could you do to identify who are the most powerful members of the group? If you were actually able to observe the people and hear them speaking to each other, then cues such as posture and vocal pitch can provide such information [17, 24]. But if we only have the written text content of their interactions — as is the typical case in on-line discussions — how do we identify evidence of power differences?

Language coordination. What we propose is that *language coordination* in text content alone can serve as a rich source of information about power differences within a group. Language coordina-

tion is a phenomenon in which people tend to unconsciously mimic the choices of function-word classes made by the people they are communicating with [35]; roughly speaking, if you are communicating with someone who uses a lot of articles — or prepositions, or personal pronouns — then you will tend to increase your usage of these types of words as well, even if you don’t consciously realize it.¹

We measure language coordination in two datasets of goal-oriented text that arise in very different settings: discussions among Wikipedia editors, containing over 240,000 conversational exchanges; and oral arguments before the U.S. Supreme Court, as processed by Hawes et al. [21, 22] and containing 50,389 conversational exchanges among Justices and lawyers. By focusing on function word classes, rather than domain-specific substantive content, we are able to evaluate the domain-independence of our techniques and their ability to generalize across different contexts; methods that rely on subject-specific cues to determine levels of power (such as the use of “Your honor” in a legal setting) are not positioned to generalize as readily.

To be able to speak in a principled way about power differences, we draw on the framework of *exchange theory* from sociology [46]. Exchange theory and its generalizations [44] have distinguished between two forms of power, which naturally parallel the types of power in our discussion above. First, a power difference between x and y can be based on the fact that x has higher *status* than y , either through a formal designation of status, or through more informal notions of status based on reputation within the group. Second, a power difference can also arise through *dependence*: if y needs something from x , and hence is dependent on x , this can give x a form of power over y that may be temporary while the underlying dependence lasts.

Power differences from language coordination. We find that differences in the level of language coordination consistently reveal both of these types of power differences, in both of our datasets. Specifically, we will present the following results.

1. In general people with low power exhibit greater language coordination than people with high power.
2. Conversely, people coordinate more with communication partners who have higher power than with those who have lower power.
3. When a person undergoes a change in status, their coordination behavior changes, and so does the coordination behavior of people talking to them.
4. When an individual is trying to convince someone who holds an opposing view, this creates a form of dependence and hence a power deficit in the sense of exchange theory; we find increased levels of language coordination in such cases.
5. The relation between status level and the extent of language coordination transfers across domains, and is a reliable cross-domain feature for status prediction.

These results suggest clear potential applications to the analysis of on-line social groups. In particular, they could provide methods for identifying power differences and levels of status in on-line settings where one has only the text content of social interactions, rather than explicit markers of status or explicitly annotated links.

¹We note that language coordination is just one form of coordination where such phenomena occur [18] (another is posture coordination, for example); we focus on language coordination because it can be measured in textual interaction.

Similarly, they could also provide a means of analyzing conversations between two users of a social media platform so as to determine the power balance or levels of relative status in their relationship. In all such uses, the methods do not require domain-specific knowledge of the on-line application being analyzed. We also note that the role of features internal to the content can be crucial in some of these settings, since it has been observed that message frequency and message volume do not necessarily suffice to determine relative status. As Rowe et al. state [37], “As we move down the corporate ladder, the conversational flows of dissimilar employees can in fact be quite similar.” Indeed, it is easy to think of contexts where dominant individuals consume a lot of the conversational bandwidth, and others where low-status individual take up most of the airtime with their advocacy toward higher-status participants.

There is something striking about the fact that the content features being employed are properties of language that tend to escape conscious attention. The phenomena we find in the text content are consistent and significant, but they are not effects one notices in reading or listening to the interactions; in essence, they operate on levels that only show up when you use computational methods to explicitly tune in to them. Moreover, since our methods are based on function words, it means one can apply them to language samples from which the content words have been redacted, raising intriguing implications for compact representations and user privacy.

Summary: Novel contributions of present work. Our use of language coordination as a key source of information draws on a history of coordination studies originating in social psychology; we discuss this background in the section on Related Work. These psychological studies of coordination focused on small-scale settings where participant behaviors could be individually observed; the identification of language coordination phenomena in large-scale on-line text was done very recently by Danescu-Niculescu-Mizil et al. [10] using data from Twitter. Their study, however, did not use coordination to analyze underlying social structure, as we do here. Thus, to our knowledge, our work is the first to identify connections between language coordination and social power relations at large scales, and across a diverse set of individuals and domains.

In addition, our work here provides the following further novel contributions.

Multiple domains with large amounts of data. By using large amounts of data, we can pick up subtle effects and explicitly vary some of the underlying conditions for coordination across different subsets of the data. Moreover, working with two such different corpora allows us to test the domain independence of our linguistic-coordination approach.

Status change. Wikipedians can be promoted to administrator status through a public election, and almost always after extensive prior involvement in the community. Since we track the communications of editors over time, we can examine how linguistic coordination behavior changes when a Wikipedian becomes an “admin”. To our knowledge, our study is the first to analyze the effects of status change on specific forms of language use.

Situation-dependent forms of power. By generalizing from status to broader notions of power, our study is, to our knowledge, also the first to show how multiple types of power relationships — and in particular situation-dependent power — can be exposed through domain-independent textual features.

2. COORDINATION AND POWER

We can apply communication accommodation theory [34, 18, 40, 17], an influential line of research in sociolinguistics, to our investigations because the theory implies the following principle:

Principle \mathcal{P} . Linguistic coordination is a function of the power differential between the speaker and the target: the lower the power of the speaker relative to that of the target, the more she coordinates (and vice versa, the higher the relative power of the speaker, the less she coordinates).

where *speaker* refers to the person producing the reply in an exchange, and *target* refers to the person initiating the exchange (and thus the target of the speaker’s reply).² In the context of group conversations, which is the focus of the present work, this principle leads to the following two concrete hypotheses, the first one based on the power of the target and the second one based on the power of the speaker:

\mathcal{P}_{target} . People in general coordinate more towards *high-powered* people than towards *low-powered* people.

$\mathcal{P}_{speaker}$. *High-powered* people coordinate less than *low-powered* people towards their targets.

(Note that the two hypotheses are distinct; because of the asymmetry of coordination, neither hypothesis implies the other.)

In addition to power imbalance, we hypothesize that the personal traits of the participants could also determine how much they coordinate:

\mathcal{B} . People have a baseline coordination level, which is determined by personal characteristics (such as their sociability and level of social engagement).

It is worth noting that, although previously hypothesized in the literature, the relation between power and (not necessarily linguistic) coordination was empirically supported exclusively through very small scale studies. For example, [19] showed that Larry King, the host of a popular talk-show in the U. S., coordinated more in his vocal pitch to his high-status guests (such as then-President Clinton) than to low-status guests. As for *linguistic style* coordination, [35] looked at 15 Watergate transcripts involving only four people altogether — Richard Nixon and three of his aides — and found mixed effects of power, depending on the stylistic dimension. Furthermore, there are competing theories which postulate that the relation between power and coordination is the reverse of \mathcal{P} [2]. Therefore, it is not at all clear a priori whether \mathcal{P}_{target} and $\mathcal{P}_{speaker}$ hold at large.

In what follows, we present the first in-depth and large-scale examination of the power-based hypotheses ($\mathcal{P}_{speaker}$ and \mathcal{P}_{target}), as well as our additional hypothesis based on personal characteristics (\mathcal{B}) in two very different datasets of textual exchanges.

While power might correlate with certain personal traits in a given community, making the distinction between \mathcal{P} and \mathcal{B} difficult, they differ in one important aspect which we will exploit in our study: power can change abruptly — such as when an individual is assigned a new role — while personal traits, in comparison, are more stable over time. As a result, examining the temporal change in coordination level of people who have undergone changes in power can help us isolate the effect of \mathcal{P} from that of \mathcal{B} . In particular, this will help us address the following question: if we do find evidence supporting hypothesis \mathcal{B} , would it be sufficient to explain the data, or will we see power playing a role on top of the baseline coordination level?

² We use “initiate” and “reply” loosely: in our terminology, the conversation (x : “Hi.” y : “Tired?” x : “No.”) has two exchanges, one initiated by x ’s “Hi”, the other by y ’s “Tired?”.

3. POWER RELATIONS IN WIKIPEDIA AND SUPREME COURT DATA

In this section, we describe the two corpora of consequential discussions we used in our studies. The first consists of discussions between editors on Wikipedia; the second consists of transcripts of oral arguments before the United States Supreme Court. Both settings involve power differentials, both through status and dependence, as we will see below. Our Wikipedia corpus is much larger, potentially more representative of online discussions, and allows us to study the effects of changes in power; but the Supreme Court represents a less collaborative situation than Wikipedia (in the Supreme Court data, there are always explicit opposing sides) and is an instance of an off-line setting. The differences in the two corpora help us focus on general, domain-independent relationships between relative power and linguistic coordination.

We begin by briefly describing the roles and text content of our two domains, and then discuss how we formalize the different kinds of power imbalances within the domains.

3.1 Discussions among Wikipedia editors

Roles and role changes. Wikipedia editors form a close community with salient markers of status. Administrators, commonly known as *admins*, are Wikipedia editors “trusted with access to restricted technical features” such as “protect, delete and restore pages”, as well as “block other editors”³. In effect, admins have a higher status than other users (*non-admins*) in the Wikipedia community, and editors seem to be well aware of the status and activity history of other editors. Users are promoted to admins through a transparent election process known as requests for adminship⁴, or *RfAs*, where the community decides who will become admins. Since *RfAs* are well documented and timestamped, not only do we have the current status of editors, we can also extract the exact time when editors underwent role changes from non-admins to admins.

Textual exchanges. Editors on Wikipedia interact on *talk* pages⁵ to discuss changes to article or project pages. We gathered 240,436 conversational exchanges carried out on the talk pages, where the participants of these (asynchronous) discussions were associated with rich status and social interaction information: status, timestamp of status change if there is one, as well as activity level on talk pages, which can serve as a proxy of their sociability, or how socially inclined they are. In addition, there is a discussion phase during *RfAs*, where users “give their opinions, ask questions, and make comments” over an open nomination. Candidates can reply to existing posts during this time. We also extracted conversations that occurred in *RfA* discussions, and obtained a total of 32,000 conversational exchanges. Most of our experiments were carried out on the larger dataset extracted from talk pages, unless otherwise noted. (The dataset will be distributed publicly.)

3.2 Supreme Court oral arguments

While Wikipedia discussions provide a large-scale dataset with rich meta-information, overall, high-status people and low-status people are collaborating to accomplish a task. In other social hierarchies, power can be more markedly concentrated to the high-status people. Oral arguments before the Supreme Court provide such a setting.

Roles. A full court consists of nine Justices, although occasionally some recuse themselves. In the oral arguments for a case, lawyers

³<http://en.wikipedia.org/wiki/Wikipedia:Administrators>

⁴http://en.wikipedia.org/wiki/Wikipedia:Requests_for_adminship

⁵http://en.wikipedia.org/wiki/Wikipedia:Talk_page_guidelines

	Wikipedia	
	higher power	lower power
Status	admins admins	non-admins admins-to-be (before RfAs)
Dependence	diff. vote	same vote
	Supreme Court	
	higher power	lower power
Status	Justices Chief Justices	lawyers Associate Justices
Dependence	unfavorable Justice	favorable Justice

Table 1: Power differentials exhibited in our data.

for each party have thirty minutes to present their side to the Justices. The Justices may interrupt these presentations with comments or questions, leading to interactions between the lawyers and Justices. After the oral arguments and subsequent deliberations, cases are decided by majority vote of the Justices. This provides an interesting additional test ground: instead of asynchronous textual exchanges in a social hierarchy working collaboratively, here we have verbal exchanges in a social hierarchy where *Justices* decide the final outcome. In addition, conversations here are over topics in a completely different domain.

Transcripts of verbal exchanges. Transcripts of oral arguments in Supreme Court are publicly available⁶. We used a pre-processed version of this dataset described in [21]. We enhanced this dataset with the final votes from the Spaeth Supreme Court database⁷. In total, we have 50,389 verbal exchanges for 204 cases. 11 justices (two of which have little conversational data: Thomas⁸ and Alito) and 311 lawyers are represented in the dataset. 73% of the lawyers only appear in one case, and the maximum number of cases in our data where one lawyer appears is 15. As such, trends identified on this dataset should not be due to idiosyncratic behavior of a few over-represented lawyers.

3.3 Power Relations in the Data

Having now surveyed the nature of the two domains, we discuss the different kinds of power relations that they contain. An overview of the following discussion is summarized in Table 1.⁹

In our discussion of roles earlier in this section, we have already indicated some of the basic status differences: the distinction between *admins* and *non-admins* on Wikipedia, and the distinction between *Justices* and *lawyers* (as well as *amici curiae*) in the context of the Supreme Court. We can also identify certain finer-grained distinctions, including the distinction between the *Chief Justice* of the Supreme Court and the *Associate Justices*.¹⁰ And on Wikipedia, we can also study the behavior over time of users who were promoted to the position of admin — in effect, comparing their behavior as *admins* to their earlier behavior as *admins-to-be*.

⁶http://www.supremecourt.gov/oral_arguments/

⁷<http://scdb.wustl.edu/>

⁸In 2011, Justice Thomas marked five terms without speaking in any oral arguments. [30]

⁹Throughout the paper we use color coding to indicate the relative power relations relevant for the respective discussion. These colors are simply intended as a helpful mnemonic and can be ignored without any loss of meaning.

¹⁰During the time period covered by our data, there was a succession in the Chief Justice position from Rehnquist to Roberts; each appears as Chief Justice for some of the cases.

Our data also makes it possible to study several instances of power differences based on *dependence*. To begin with, we note the general principle that status and dependence are almost never completely distinct [43], since a person in a high-status role frequently appears in situations where people are dependent on them.

The data, however, offers us opportunities to study forms of dependence where the level of status has been largely controlled for. Key among these are forms of dependence created by the need to convince someone who disagrees with you. If you are advocating a position in a debate with opposing sides leading to an eventual decision (for example, a Supreme Court case, or a policy discussion on Wikipedia prior to a vote), then your audience can be roughly divided into two groups: people who would naturally tend to vote in favor of your position, and people who would naturally tend to vote against your position. Principles of exchange theory indicate that in such situations, you are more dependent on the people who would naturally vote against you, and less dependent on the people who would naturally vote for you, since in order to accomplish your goal, you need to effect a more substantial behavior change in the former group [13, 26, 47]. An important further point here is that we are studying settings where participants can readily anticipate, either through dialogue or advance knowledge, who is “on their side” and who is “on the other side,” and so it makes sense to suppose that they are aware of these dependence relations during the interaction.

Motivated by this, in the Supreme Court data we will compare levels of coordination of lawyers toward *unfavorable Justices* who (eventually) vote against their side and toward *favorable Justices* who (eventually) vote for their side; there is more dependence and hence more of a power difference in the former case. In the Wikipedia data, we will compare levels of coordination of editors with others who *vote the opposite way* and with others who *vote the same way*; here too, there is more dependence and hence more of a power difference in the former case. We should also note the exchange-theoretic principle that a dependence relation affects both sides: x ’s dependence on y is expected not just to affect x ’s behavior in their interaction, but y ’s as well.

4. LINGUISTIC STYLE COORDINATION

As discussed earlier, we use *linguistic style coordination* to quantify the degree to which one individual *immediately* echoes the linguistic style of the person they are responding to. Here, linguistic style is quantified by a person’s usage of certain linguistic style markers (corresponding to categories of function words). We start with a description of these markers, followed by a more formal definition of coordination.

Linguistic style markers. We measure the linguistic style of a person by their usage of function words that have little lexical meaning, thereby marking style rather than content.

For consistency with prior work, we employed the nine LIWC-derived categories [36] deemed to be processed by humans in a generally non-conscious fashion [25]. The nine categories are: articles, auxiliary verbs, conjunctions, high-frequency adverbs, impersonal pronouns, negations, personal pronouns, prepositions, and quantifiers (451 lexemes total). Given its sparseness in the data and the relatively high semantic meaning it carries, we ignored the negation category¹¹. Each of the remaining eight categories is considered a linguistic style *marker* (or feature) in this study.

It is important to note that in general coordination is multimodal: it does not necessarily occur simultaneously for all markers [15], and speakers may coordinate on some features but diverge on oth-

¹¹Negations were also partly discarded by [25].

ers [42]. Thus, we measure both coordination for each individual marker and coordination aggregated across markers.

4.1 Coordination measure

Here we present a variation of the measure introduced in [10], adapted to the setting of group conversations.

Coordination with respect to a marker. We start by defining the coordination of one person b towards another person a with respect to a specific linguistic style marker class m . We want to quantify how much the use of marker m in a 's utterances *immediately triggers* the occurrence of m in b 's respective replies to those utterances. To put it another way, we want to measure how much a 's use of m in an utterance u_1 increases the probability that b will use m in his reply u_2 , where the increase is relative to b 's normal usage of m in conversations with a .

Recall from §2 that we call b the *speaker* and a the *target* of a conversational exchange ($a : u_1, b : u_2$), since a is the target of b 's reply when b acts as the speaker. We say an utterance *exhibits* m if it contains a word from the corresponding category. Let $X_{u_1}^m$ be the indicator variable for an utterance u_1 exhibiting m ; similarly, let $X_{u_2 \leftrightarrow u_1}^m$ be the indicator for reply u_2 exhibiting m .

Given a set $S_{a,b}$ of exchanges ($a : u_1, b : u_2$), we define the coordination of b towards a as:

$$C^m(b, a) = P(X_{u_2 \leftrightarrow u_1}^m | X_{u_1}^m) - P(X_{u_2 \leftrightarrow u_1}^m), \quad (1)$$

where the probabilities are estimated over $S_{a,b}$. In other words, $C^m(b, a)$ measures the increase (or decrease) of the probability that b will use m in her reply as an effect of a using m in her initial utterance.

Coordination towards a group. In the context of group conversations, we can extend this definition to coordination of a particular speaker b towards a *group of targets* A by simply modifying the set of exchanges on which the probabilities in (1) are estimated. Specifically, given a set of $S_{A,b}$ of exchanges ($a : u_1, b : u_2$) involving initial utterances u_1 of various targets $a \in A$ and replies u_2 of b , the coordination of b to the group A is:

$$C^m(b, A) = P(X_{u_2 \leftrightarrow u_1}^m | X_{u_1}^m) - P(X_{u_2 \leftrightarrow u_1}^m), \quad (2)$$

where this time the probabilities are estimated over $S_{A,b}$.

We then define the coordination of one group of people B towards another group A as the average coordination of speakers in B to targets in A :

$$C^m(B, A) = \langle C^m(b, A) \rangle_{b \in B} \quad (3)$$

By taking the macro (unweighted) average, our measure will not be dominated by a few active speakers in a dataset.

Note that our measure is not symmetric, which fits the purpose of this study well, since the power relations we want to investigate are also asymmetric.

Aggregated measures. Next we define aggregated measures of coordination of B to A , which are not specific to any individual markers, and can be used to provide an overall picture of the level of coordination between the groups. Ideally we want to simply compute $C(b, A)$ as the macro-average of $C^m(b, A)$ across different markers m , and then compute $C(B, A)$ same way as in (3). Recall, however, that $C^m(b, A)$ can only be computed if $S_{A,b}$ contains enough exchanges exhibiting m to reliably estimate both probabilities in (3), which is not always the case for all people with respect to all markers. For instance, some persons rarely use quantifiers, leaving C^{quant} undefined in those instances.

We accounted for such “missing values” in three different ways, resulting in three aggregated measures:

Aggregate 1 Compute the “ideal” macro-average $C(b, A)$ only for the persons b for whom $C^m(b, A)$ can be computed for all markers (ignore all the others). This reduces the set of persons considered by the aggregated measure, but provides the most direct measure (in the sense that it does not rely on any particular “smoothing” assumptions as the other two aggregated measures do).

Aggregate 2 For each person b , if $C^m(b, A)$ is undefined, we “smooth” it by using the group average $C^m(B, A)$ instead; this measure considers everybody for which we can compute coordination for at least one marker, but assumes the people in a given group share similar coordination behavior, which arguably can be the very point we need to verify.

Aggregate 3 For each person b , we take the average only over the markers for which $C^m(b, A)$ is defined; this is equivalent to assuming that b would have exhibited the same level of coordination for the missing markers as they did with other markers. This aggregation also considers everybody for which we can compute coordination for at least one marker.

Formalization of the power hypotheses. Now that we have introduced a more formal definition of coordination between two groups of people, we present a way of formalizing the hypotheses introduced in Section 2 in terms of this definition. If people in a group G^{high} have more power than people in a group G^{low} , and U is a set of arbitrary people, the power hypotheses can be rewritten as:

$$\mathcal{P}_{target}. C(U, G^{high}) > C(U, G^{low})$$

$$\mathcal{P}_{speaker}. C(G^{high}, U) < C(G^{low}, U)$$

5. EMPIRICAL INVESTIGATION

Using the concepts and formalism introduced in the previous sections, we can now investigate the relation between linguistic coordination and power differentials in concrete conversational settings. Specifically, we test whether the principle \mathcal{P} and the hypotheses \mathcal{P}_{target} and $\mathcal{P}_{speaker}$ introduced in Section 2 can be empirically confirmed in the two datasets described in Section 3. We begin by discussing power differences arising from status in Wikipedia (where our primary status distinction will be admins vs. non-admins) and in the Supreme Court (where our primary status distinction will be Justices vs. lawyers). After this, we consider power differences arising from dependence.

5.1 Power from status: Wikipedia

First, communication behavior on Wikipedia provides evidence for hypothesis \mathcal{P}_{target} : users coordinate more toward the (*higher-powered*) **admins** than toward the **non-admins** (Figure 1(a))¹².

In the other direction, however, when comparing **admins** and **non-admins** as speakers, the data provides evidence that is initially at odds with $\mathcal{P}_{speaker}$: as illustrated in Figure 1(b), **admins** coordinate to other people *more* than **non-admins** do (while the hypothesis predicted that they would coordinate *less*).¹³ We now explore some of the subtleties underlying this result, showing how it arises as a superposition of two effects.

¹²Throughout the figures we will indicate statistical significance (according to an independent t-test) using the notation: *="p<0.05", **="p<0.01", ***="p<0.001". We also indicate next to each label the number of users considered by the Aggregated 1 measure (i.e.,

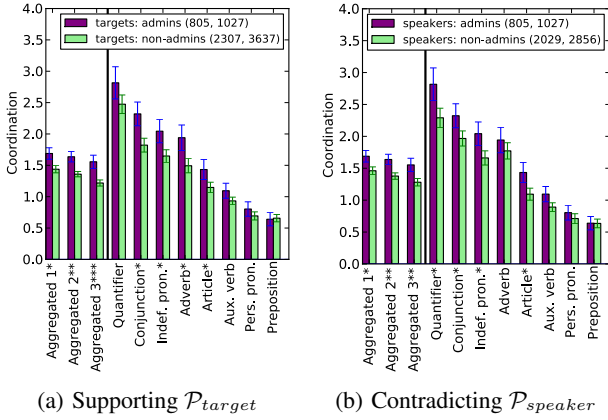


Figure 1: Status and linguistic coordination: (a) Users coordinate more towards **admins** (*high-powered*) than towards **non-admins** (*low-powered*) supporting hypothesis \mathcal{P}_{target} ; (b) on the other hand, **admins** (*high-powered*) coordinate more than **non-admins** (*low-powered*) when replying to other people, contradicting hypothesis $\mathcal{P}_{speaker}$.




Figure 2: Same-status comparisons (supporting a “winner” personality hypothesis): (a) **admins-to-be** coordinate more than those who remain **non-admins** throughout; (b) during **adminship elections** (RfAs), **admins-to-be** coordinate more than **failed-to-be**.

Personal characteristics: Hypothesis \mathcal{B} . One possible explanation for the inconsistency of our observations with $\mathcal{P}_{speaker}$ is the effect of personal characteristics suggested in Hypothesis \mathcal{B} from Section 2. Specifically, admin status was not conferred arbitrarily on a set of users; rather, admins are those people who sought out this higher status and succeeded in achieving it. It is thus natural to suppose that, as a group, they may have distinguishing individual traits that are reflected in their level of language coordination.

the users for which we can compute coordination for all markers) and the total number of users considered by Aggregated 2 and 3 (i.e., the users for which we can compute coordination for at least one marker).

¹³Note that the observations shown in Figure 1(a) do not imply those in Figure 1(b), nor vice-versa. For example, the trend in Figure 1(b) does not change if we restrict the speakers to be only non-admins (or only admins).

Fortunately we can extract rich enough data from Wikipedia that it becomes possible, to a significant extent, to separate the effect of status from these individual traits, establishing that both effects play a role. Our separation of these effects is based on the fact that status can change abruptly, while personal characteristics, though mutable, are more stable over time. On Wikipedia, status changes are well documented as they can occur only through requests for adminship (RfAs). To investigate whether the effects observed in Figure 1(b) are purely tied to status, we look at communication differences between these same two populations over time periods when there was no status difference between them: we compare the set of **admins-to-be**— future admins before they were promoted via their RfA — with **non-admins**. Figure 2(a) shows that the same differences in language coordination were already present in these two populations — hence, they are not an effect of status alone, since they were visible before the former population ever achieved its increase in status.

Above we noted that the set of admins corresponds to users who both had the ambition to achieve admin status, *and* succeeded in achieving it. In Figure 2(a) we saw a comparison in which both of these distinctions were present — the ambition and the success. One way to separate the second issue from the first is to look at differences in coordination between users who were promoted (**admins-to-be**), and those who went through the RfA process but were denied admin status (**failed-to-be**). Both **admins-to-be** and **failed-to-be** had the ambition to become admins, but only members of the former group were successful. We investigate coordination differences between these two groups during a period when their adminship ambitions are arguably most salient: during the discussions in each user’s own RfA process. Figure 2(b) shows that even in the conversations they had on their RfA pages, the **admins-to-be** were coordinating more to the others than the **failed-to-be**, providing evidence for a strong form of Hypothesis \mathcal{B} .

Revisiting status: Hypothesis $\mathcal{P}'_{speaker}$. We now return to the issue of status, and describe a method of partially controlling for personal characteristics so as to evaluate the following modification of Hypothesis $\mathcal{P}_{speaker}$:

$\mathcal{P}'_{speaker}$. When controlling for personal characteristics, *high-powered* people coordinate less than *low-powered* people.

To study $\mathcal{P}'_{speaker}$, we create two populations for comparison: the interactions of each **admin** before his or her promotion via RfA (i.e., when they were **admins-to-be**), and the interactions of each **admin** after his or her respective promotion. Figure 3(a) shows how the resulting comparison confirms $\mathcal{P}'_{speaker}$: **admins-to-be** decrease their level of coordination once they gain power.¹⁴ Interestingly, the reverse seems to be true for **failed-to-be**: after failing in their RfAs — an event that arguably reinforces their failure to achieve high status in the community — they coordinate more (p-value 0.05; we omit the figure due to space limitations.)

In addition, we can employ status-change to reinforce \mathcal{P}_{target} in a setting that controls for personal characteristics: we find that users coordinate more to **admins** after promotion than when they were **admins-to-be** (p-value<0.05).

Finally, in Figure 3(b), we investigate how quickly the change in status is reflected in the communication behavior of the users involved. In addition to the monotonic changes in coordination

¹⁴Note that the trend shown in Figure 3(a) is maintained when considering the exact same users in both groups (i.e., excluding the users which did not have enough conversations both before and after adminship). Also note that we allow a time buffer a month after the RfAs between the two sets of conversations we compare.

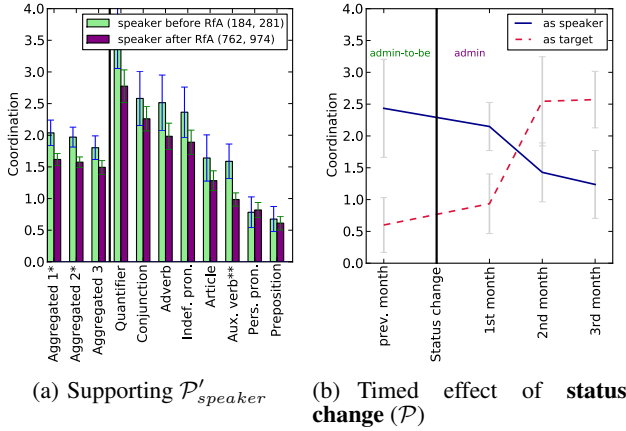


Figure 3: Effect of status-change: (a) admins-to-be coordinate less after they become admins; (b) Aggregated 1 coordination of the user (as speaker) and, respectively, towards the user (as target) in the months before and after status change occurs through RfA.




Figure 4: Justices coordinate less to lawyers than conversely.

levels over time, and in the hypothesized directions, it is interesting to note that the most dramatic change in coordination is visible in the second month after the change in status occurred. This suggests a period of accommodation to the newly gained status, both for the person that undergoes the change and for those witnessing it.

5.2 Power from status: Supreme Court

In the setting of the Supreme Court, status differences are extremely salient and do not suffer from the correlations that added complexity to the study of $\mathcal{P}_{speaker}$ in its original form. Also, conversations during the oral arguments (almost) always are between a Justice and a lawyer. Thus, our basic finding can be expressed succinctly in Figure 4, which shows significantly more coordination from lawyers to Justices than vice versa.¹⁵

In the Supreme Court setting we can also study finer-grained

¹⁵Throughout we consider each appearance of a Justice (or lawyer)

status distinctions, to see if these too are manifested in language coordination differences. In our data, we find that they are: in concordance with \mathcal{P}_{target} , we observe that lawyers coordinate significantly more toward the Chief Justice than toward the Associate Justices (p-value<0.01).

5.3 Power from dependence

As noted in Section 3, we can study power differences based on dependence — even for fixed levels of status difference — using the exchange-theoretic principle that the need to convince someone who disagrees with you creates a form of dependence [13, 26, 47]. Moreover, this power difference is predicted to be felt by both sides — the side with lower power and the side with higher power.

In the case of a lawyer-Justice interaction in the Supreme Court data, let us define the Justice to be favorable to the lawyer if he or she ends up voting on the lawyer’s side, and unfavorable if he or she ends up voting against the lawyer’s side. It is well-understood that the Justices often come into the case with a general leaning toward one side or the other based on their judicial philosophy — this has been acknowledged for example in interviews with members of the Court [38] — and lawyers through their preparation for the case will come in with knowledge of these leanings. Hence it is reasonable to suppose that the favorable-unfavorable distinction will be salient to the interaction during oral arguments.

And indeed, Figures 5(a) and 5(b) shows that the power differences created by this form of dependence are reflected in the amount of coordination, in both directions. First, lawyers coordinate more toward unfavorable Justices (on whom they are more dependent) than toward favorable Justices, in keeping with \mathcal{P}_{target} . Second, unfavorable Justices coordinate less toward lawyers than favorable Justices do, in keeping with $\mathcal{P}_{speaker}$. Given the formal framework of impartiality that characterizes the Justices’s behavior at the Supreme Court, it is intriguing to see the undercurrent of language coordination differences nevertheless hinting at their eventual decision.

We see a similar effect of dependence on coordination in the context of discussions with opposing sides on Wikipedia. During RFAs, one voter may try to change the opinion of voters on the other side who have already cast their vote. (Changing your vote during the RfA process is allowed, and hence there is an incentive to convince voters to consider this.) Users coordinated more when engaging with users on the opposite side than with those who voted the same way (p-value<0.05; for space reasons we omit the figure). This finding too, via the arguments about opposing sides and dependence, supports the general power-coordination principle \mathcal{P} .

6. CROSS-DOMAIN ANALYSIS AND INTERACTION AMONG HYPOTHESES

6.1 Coordination as a Cross-Domain Feature

Part of the motivation for studying the relation between coordination and power is that the principles that govern this relation appear to be domain-independent. Here we perform a set of analyses to show that coordination features do generalize across our two domains more effectively than other text-based features for the problem of inferring power. We find that indeed, compared to the other features we consider, they are the only set of features to display any non-trivial generalization.

in a case as a separate entity, by way of which we avoid case idiosyncrasies and data sparseness.




Figure 5: Dependence and linguistic coordination: (a) lawyers adjust their coordination level according to whether the Justice is unfavorable or favorable, supporting \mathcal{P}_{target} ; (b) favorable Justices coordinate more than unfavorable Justices ($\mathcal{P}_{speaker}$).

Our analysis is based on the following prediction task: for a given pair of discourse participants x and y (of different status) who have engaged in conversations with each other, we predict whether x has the higher status. In this setting, a random guess baseline would achieve 50% accuracy. We stress, however, that this prediction task is primarily a means to assess cross-domain generalization i.e., not as a free-standing task in itself. Indeed, the best achievable performance on this status-prediction task appears to be quite domain-dependent. In some domains such as the Supreme Court, idiosyncratic cues in text usage (e.g., lawyers begin their sentences with stylized lead-ins, such as “Your honor”, that clearly mark them as lawyers, not Justices) enable almost perfect performance when these cues are available as features. In other domains, such as Wikipedia, an informal evaluation using two human annotators familiar with the domain produced only 70% accuracy (and an inter-annotator agreement of only 80%). Thus, our interest is not in whether coordination features achieve the best within-domain performance, but in whether they are particularly effective at generalizing (as we indeed find them to be).

Experimental setup. Let R_x be x ’s replies to y , and R_y be y ’s replies to x , and $Len(S)$ be the average length of all utterances in the set S . Let \mathcal{F}_{style} be the set of 8 stylistic markers introduced in Section 4. We define the following sets of features used as input to an SVM classifier:

- coordination features: binary features indicating, for each $m \in \mathcal{F}_{style}$ as well as for Aggregated 1¹⁶, whether x coordinates more to y than y to x on m
- stylistic features: $Len(R_x)$, $Len(R_y)$ and the frequency of each marker $m \in \mathcal{F}_{style}$ in R_x and, respectively, in R_y
- bag of words: frequency of each word in R_x , frequency of each word in R_y , L_2 -normalized

For experiments on the Wikipedia data, which we denote as *wiki*, we considered (*admins*, *non-admins*) pairs and their conversations

¹⁶We only considered pairs of participants for which enough data was available to compute coordination on all stylistic features.

Training corpus	in-domain		cross-domain	
	<i>wiki</i>	<i>court</i>	<i>court</i>	<i>wiki</i>
Test corpus	<i>wiki</i>	<i>court</i>	<i>wiki</i>	<i>court</i>
coordination features (9)	57.7	70.4	57.1	55.0
stylistic features (18)	59.2	51.4	50.0	51.9
bag of words (20k)	51.4	99.5	45.2	40.1

Table 2: Accuracy for prediction task using different feature sets (size of feature sets indicated in parentheses). The cross-domain performance is in the right-hand pair of columns.

(that took place after the status difference was effective). For the Supreme Court dataset (*court*), we considered (*Justices*, *lawyer*)¹⁷.

For *in-domain* experiments, we report average accuracy over cross-validation within the same domain (i.e., training and test corpora are both *wiki* or *court*); for *cross-domain* experiments, we train on one domain and test on the other.

Results. Table 2 summarizes the results. In brief, we find first of all that coordination features are the only ones to perform statistically significantly better than random guessing in the *cross-domain* settings — the other classifiers simply learn cues that are idiosyncratic to their training data, and fail to generalize. (Note for example that the bag-of-words method picks up on the near-perfect textual cues marking lawyers in the Supreme Court data, but this method performs worse than random guessing when applied to the other domain.)

Even looking at the *in-domain* tasks — which were not our primary focus here — we find that coordination features are the only ones that perform statistically significantly better than random guessing on *both* datasets.

6.2 Interactions among Hypotheses \mathcal{P} and \mathcal{B}

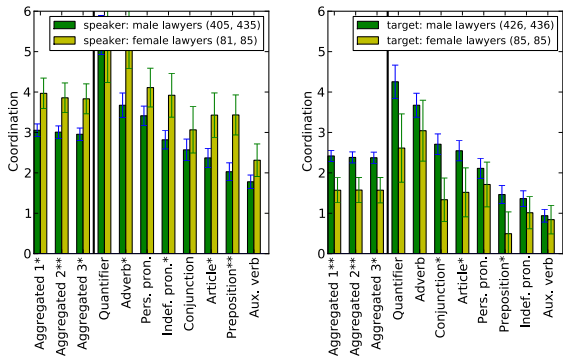
In Section 5 we saw that the interaction between personal characteristics (which form the basis for Hypothesis \mathcal{B}) and power differentials (which form the basis for Hypothesis \mathcal{P}) can lead to complex effects. Here we consider two cases where this interaction raises interesting issues, and point to open questions in the analysis of coordination.

An individual’s level of social engagement is one type of personal characteristic that interacts with coordination and power. As a simple proxy for social engagement, for purposes of discussion here, we consider the volume of communication the individual engages in. As we noted in Section 1, simple volume measures such as this do not seem to readily yield domain-independent information about power, since they vary considerably across domains — in some domains the powerful people talk a lot, and in others they talk relatively little. For example, when people are promoted to *admin* status, their volume of communication goes up while (as we have seen) their level of coordination goes down. On the other hand, *lawyers* talk more than *Justices* in the Supreme Court data, and (again as we have seen) they also coordinate more in the *lawyer-Justice* interactions.

However, if we restrict attention to a fixed sub-population within a given domain, there are interesting connections between coordination and volume that suggest further questions. In particular, on Wikipedia we consider the number of replies posted by a user on talk-pages as a measure of communication volume, and hence a

¹⁷In order to focus on the conversational exchanges and avoid exchanges in which the lawyers formally introduce their case, we considered only cases where the length difference between the two utterances were fewer than 20 words.

proxy for their level of social engagement on the site. We compared users in the top 1/3 of the sorted order by communication volume with users in the bottom 1/3, finding that users with higher numbers of replies are more likely to coordinate to others (p -value <0.05) We observed the same effect when we compared the communication volumes of users with the same status: among admins, users with more communication are also more likely to coordinate, and the same trend holds among non-admins. Similar effects also hold for other measures of communication volume. Again, we note that for other domains (such as the Supreme Court) show an inverse relation between volume and coordination in the communication transcripts, and so it is an interesting question to identify the principles that determine how this relationship plays out in different settings.



(a) gender difference in speakers (b) gender difference in targets

Figure 6: Gender differences.

We also consider a second basic example that raises an interesting challenge for distinguishing between Hypotheses \mathcal{P} and \mathcal{B} : the effect of gender on coordination, using the fact that gender information is available for participants in the Supreme Court dataset. The main finding here, in Figure 6, is that overall female lawyers coordinate more than male lawyers when talking to Justices, and correspondingly, Justices coordinate more towards female lawyers than towards male lawyers. Given the extensive history of work exploring culturally-embedded status and power differences based on gender [3, 43], one interpretation of this finding is directly in terms of Hypothesis \mathcal{P} . However, since it is also potentially related to theories of gender-based communication differences [23], the question of separating Hypotheses \mathcal{P} and \mathcal{B} becomes challenging here. We think it is a promising possibility that language coordination effects may be able to serve as a lens through which to measure many similar kinds of distinctions in both on-line and off-line conversational settings.

7. FURTHER RELATED WORK

In the opening sections, we have discussed some of the ways in which earlier work used text content to analyze on-line networks [49, 31, 11, 9, 16], as well as background on language coordination and the exchange-theoretic notions of power from status and dependence. Here we discuss some further work that is related to the general issues we consider here.

Power and structural features. There has been extensive work on using structural features, rather than language, to infer notions of “importance” in networks, both in the literature on social networks [14] and on the Web [8]. Recent work has also studied the inference of status from on-line social network features [20, 29].

Power and language. The relation between linguistic coordination¹⁸ and status has been studied in an interesting but small-scale study [35] which looked at 15 Watergate transcripts involving only four people altogether (as such, it did not explore the bulk of the phenomena we identify here). A recent large-scale study of language coordination in the on-line domain [10] used data from Twitter, where markers of status and power are not as readily inferred; they identified a weak correlation between language coordination and Twitter follower counts, suggesting a potential connection to status measures. Additionally, researchers have used text features other than linguistic coordination to identify status differences [12, 33, 39, 5]; in contrast with our work these methods picked up situation-specific cues, such as the word “termination” for the Enron corporate-email corpus [12], which are unlikely to generalize across contexts.

Collaborative communities. An extensive amount of work has contributed to a better understanding of the rules that govern collaborative behavior in online communities. Wikipedia was used as a testbed for studying user interaction at large [4, 32, 45, 41, 27] and the promotion process in such communities [7, 28]. Reviewer behavior and incentives to participate in the collaborative process were studied in the context of commercial review sites [16, 48, 6, 32].

8. CONCLUSIONS

We have shown that *language coordination* can provide information about the power differences between individuals taking part in the social interactions within a group. While traditional methods for inferring social structure at large scales have used network properties, our findings thus suggest methods for taking information from the textual content of social interactions and bringing it to bear on these types of problems

There are many additional directions suggested by the current work. One set of questions is to further understand the types of social distinctions that are manifested by differences in coordination; while we have seen that multiple forms of power can be observed this way, there could well be other properties that can be exposed as well. More broadly, our work makes clear that language use contains subtle features that reveal latent social information, and identifying new classes of such features promises to extend our understanding of group interaction.

Acknowledgments. We thank Timothy Hawes for sharing the processed Supreme Court transcripts from [22], and Bruno Abrahao, Chris Diehl, Susan Dumais, Shimon Edelman, Jacob Eisenstein, Michael Gamon, Susan Herring, Molly Ireland, Diana Minculescu, Alex Niculescu-Mizil, Daniel Romero, and the participants in the Cornell NLP Seminar for valuable discussions and suggestions.

References

- [1] K. Aronsson, L. Jönsson, and P. Linell. The courtroom hearing as a middle ground: Speech accommodation by lawyers and defendants. *Journal of Language and Social Psychology*, 6(2):99–115, 1987.
- [2] A. Bell. Language style as audience design. *Language in Society*, 13(2):145–204, 1984.

¹⁸For brevity, we exclude examinations of the effects of status on other types of coordination, such as pitch and vocal features, which are absent from textually-manifested discussions (see [17] for a survey) or on related phenomena such as content matching [1].

- [3] J. Berger, B. P. Cohen, and M. Zelditch Jr. Status Characteristics and Social Interaction. *American Sociological Review*, 37(3):241–255, 1972.
- [4] M. Billings and L. A. Watts. Understanding dispute resolution online: Using text to reflect personal and substantive issues in conflict. In *Proceedings of CHI*, pages 1447–1456, 2010.
- [5] P. Bramsen, M. Escobar-Molana, A. Patel, and R. Alonso. Extracting social power relationships from natural language. In *Proceedings of ACL HLT*, 2011.
- [6] S. L. Bryant, A. Forte, and A. Bruckman. Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. In *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*, pages 1–10, 2005.
- [7] M. Burke and R. Kraut. Mopping up: Modeling Wikipedia promotion decisions. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, pages 27–36, 2008.
- [8] S. Chakrabarti. *Mining the Web: discovering knowledge from hypertext data*. Morgan Kaufmann Pub, 2003.
- [9] M. D. Conover, J. Ratkiewicz, M. Francisco, B. Goncalves, A. Flammini, and F. Menczer. Political polarization on twitter. In *Proc. 5th Intl. Conference on Weblogs and Social Media*, 2011.
- [10] C. Danescu-Niculescu-Mizil, M. Gamon, and S. Dumais. Mark my words! Linguistic style accommodation in social media. In *Proceedings of WWW*, 2011.
- [11] M. De Choudhury, W. A. Mason, J. M. Hofman, and D. J. Watts. Inferring relevant social networks from interpersonal communication. In *Proceedings of the 19th international conference on World wide web*, pages 301–310, 2010.
- [12] C. P. Diehl, G. Namata, and L. Getoor. Relationship identification for social network discovery. In *Proceedings of the AAAI Workshop on Enhanced Messaging*, volume 22, pages 546–552, 2007.
- [13] R. M. Emerson. Power-dependence relations. *American Sociological Review*, 27:31–40, 1962.
- [14] K. Faust and S. Wasserman. *Social network analysis: Methods and applications*. Cambridge University Press New York, 1994.
- [15] K. Ferrara. Accommodation in therapy. In *Accommodation theory: Communication, context, and consequences*. Cambridge University Press, 1991.
- [16] E. Gilbert and K. Karahalios. Predicting tie strength with social media. In *Proceedings of CHI*, pages 211–220, 2009.
- [17] H. Giles. Communication Accommodation Theory. In *Engaging theories in interpersonal communication: Multiple perspectives*, pages 161–173. Sage Publications, Inc, 2008.
- [18] H. Giles, J. Coupland, and N. Coupland. Accommodation theory: Communication, context, and consequence. In *Contexts of accommodation: Developments in applied sociolinguistics*, page 1. Cambridge University Press, 1991.
- [19] J. Gregory, Stanford W. and S. Webster. A nonverbal signal in voices of interview partners effectively predicts communication accommodation and social status perceptions. *Journal of Personality and Social Psychology*, 70(6):1231, 1996.
- [20] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *Proceedings of the 13th international conference on World Wide Web*, pages 403–412, 2004.
- [21] T. Hawes. Computational analysis of the conversational dynamics of the United States Supreme Court. Master’s thesis, University of Maryland, 2009.
- [22] T. Hawes, J. Lin, and P. Resnik. Elements of a computational model for multi-party discourse: The turn-taking behavior of Supreme Court justices. *Journal of the American Society for Information Science and Technology*, 60(8):1607–1615, 2009.
- [23] S. C. Herring. Gender and power in on-line communication. In J. Holmes and M. Meyerhoff, editors, *The handbook of language and gender*, pages 202–228. 2003.
- [24] L. Huang, A. D. Galinsky, D. H. Gruenfeld, and L. E. Guillery. Powerful Postures Versus Powerful Roles. *Psychological Science*, 22(1):95–102, 2011.
- [25] M. E. Ireland, R. B. Slatcher, P. W. Eastwick, L. E. Scissors, E. J. Finkel, and J. W. Pennebaker. Language style matching predicts relationship initiation and stability. *Psychological Science*, 22:39–44, 2011.
- [26] J. P. Kotter. Power, dependence, and effective management. *Harvard Business Review*, pages 125–136, July 1977.
- [27] D. Laniado, R. Tasso, Y. Volkovich, and A. Kaltenbrunner. When the Wikipedians talk: Network and tree structure of Wikipedia discussion pages. In *Proceedings of ICWSM*, 2011.
- [28] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Governance in Social Media: A case study of the Wikipedia promotion process. In *Proceedings of the International Conference on Weblogs and Social Media (ICWSM A&S10)*, 2010.
- [29] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Predicting positive and negative links in online social networks. In *Proceedings of the 19th international conference on World wide web*, pages 641–650, 2010.
- [30] A. Liptak. Clarence Thomas Keeps 5-Year Supreme Court Silence. *The New York Times*, February 2011.
- [31] A. Livne, M. P. Simmons, E. Adar, and L. A. Adamic. The Party is Over Here: Structure and Content in the 2010 Election. In *ICWSM*, 2011.
- [32] Y. Lu, P. Tsaparas, A. Ntoulas, and L. Polanyi. Exploiting social context for review quality prediction. *Proceedings of the 19th international conference on World wide web*, pages 691–700, 2010.
- [33] A. McCallum, X. Wang, and A. Corrada-Emmanuel. Topic and role discovery in social networks with experiments on Enron and academic email. *Journal of Artificial Intelligence Research*, 30(1):249–272, 2007.
- [34] M. Natale. Convergence of mean vocal intensity in dyadic communication as a function of social desirability. *Journal of Personality and Social Psychology*, 32(5):790 – 804, 1975.

- [35] K. G. Niederhoffer and J. W. Pennebaker. Linguistic style matching in social interaction. *Journal of Language and Social Psychology*, 21(4):337–360, 2002.
- [36] J. W. Pennebaker, R. J. Booth, and M. E. Francis. Linguistic inquiry and word count (LIWC): A computerized text analysis program. <http://www.liwc.net/>, 2007.
- [37] R. Rowe, G. Creamer, S. Hershkop, and S. J. Stolfo. Automated social hierarchy detection through email network analysis. In *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*, pages 109–117, 2007.
- [38] "interview with justice antonin scalia (transcript)". C-SPAN (2009-06-19), 2009.
- [39] A. J. Scholand, Y. R. Tausczik, and J. W. Pennebaker. Social language network analysis. In *Proceedings of CSCW*, pages 23–26, 2010.
- [40] R. L. Street and H. Giles. Speech accommodation theory. In *Social cognition and communication*. Sage Publications, 1982.
- [41] D. Taraborelli and G. L. Ciampaglia. Beyond notability. [sic] Collective deliberation on content inclusion in Wikipedia. *Second international workshop on quality in techno-social systems*, 2010.
- [42] J. Thakerar, H. Giles, and J. Cheshire. Psychological and linguistic parameters of speech accommodation theory. In C. Fraser and K. Scherer, editors, *Advances in the Social Psychology of Language*. Cambridge, 1982.
- [43] S. Thye. A status value theory of power in exchange relations. *American Sociological Review*, 65:407–432, June 2000.
- [44] S. Thye, D. Willer, and B. Markovsky. From status to power: New models at the intersection of two theories. *Social Forces*, 84:1471–1495, 2006.
- [45] F. B. Viégas, M. Wattenberg, J. Kriss, and F. v. Ham. Talk Before You Type: Coordination in Wikipedia. In *HICSS: Annual Hawaii International Conference on System Sciences*, pages 78–78, 2007.
- [46] D. Willer. *Network Exchange Theory*. Praeger, 1999.
- [47] R. J. Wolfe and K. L. McGinn. Perceived relative power and its influence on negotiations. *Group Decision and Negotiation*, 14:3–20, 2005.
- [48] F. Wu and B. A. Huberman. Opinion formation under costly expression. *ACM Transactions on Intelligent Systems and Technology*, 1(1):1–13, 2010.
- [49] F. Wu, B. A. Huberman, L. A. Adamic, and J. R. Tyler. Information flow in social groups. *Physica A: Statistical and Theoretical Physics*, 337(1-2):327–335, 2004.