Linguistic overfitting in empirical psychology

Travis Riddle¹

¹ Department of Psychology, Columbia University

Author Note

This work was supported by a grant from the National Science Foundation, NSF #1420446. Data, code, and supplementary material are available at <https://osf.io/5xwe6/>

Correspondence concerning this article should be addressed to Travis Riddle, 406 Schermerhorn Hall, 1190 Amsterdam Avenue MC 5501, New York, NY 10027. E-mail: tar2119@columbia.edu

Abstract

Recent controversies in psychology have highlighted disagreements about how science works and how science proceeds. Despite the existence of tools and data that can help to address some of these disagreements, psychologists have been slow to provide empirical evidence that speaks to these issues. Indeed, there is essentially no work on how scientists communicate their findings. In this paper, I use language from scientific literature to answer questions about scientific behavior. Specifically, I find that features of data analysis practices to be associated with two narrative features of these manuscripts – hedging (diminishing claims) and boosting (amplifying claims). The use of hedging and boosting were found to be associated with statistical errors. Hedging and boosting were also found to be associated with and statistical overfitting, though this relationship is more uncertain. Additionally male first-authors are less likely to hedge. No other characteristics of the author were found to be associated with these linguistics features. Though the work is exploratory, the results are largely consistent with a linguistic overfitting hypothesis, in which the authors of a paper use language to amplify and diminish findings in a paper in order to create a more compelling narrative around theory. The flexibility of language is a powerful tool used to communicate research, but there is a fine line between effective interpretation of data and convenient overinterpretation, and it is worthwhile to be cognizant of when that line may be crossed.

Keywords: Meta-science, computational science, text analysis, bayesian statistics, power

Word count: 6,066

Linguistic overfitting in empirical psychology

Introduction

Recently, psychologists, economists, biologists and others have begun using scientific research to better understand how science itself operates (Schooler, 2014). Although much effort has been expended attempting to document the prevalence of various methodological or quantitative phenomena present in the literature (M. B. Nuijten, Hartgerink, Assen, Epskamp, & Wicherts, 2015; Open Science Collaboration and others, 2015; Simonsohn, Nelson, & Simmons, 2014), there has been relatively little attention paid to the social and behavioral phenomena of science (though see Smaldino and McElreath (2016) for a take on institutional phenomena). In particular, there is little work examining the psychological properties of the patterns in scientific communication. The lack of attention is particularly notable in light of the centrality that communicating research plays in scientific behavior.

Scientific results are communicated through talks, blog posts, press releases, and preprints published to digital archives, but peer-reviewed papers published in major journals are still the bread and butter of scientific correspondence. These papers represent the fundamental units through which science is communicated, scientific beliefs are established, and scientists are judged. The language used within peer-reviewed papers are traces of scientific behavior, operating within the constraints, standards and norms of academic journals. Connecting these traces of behavior with characteristics of the science featured in these papers can highlight inefficiencies in the machinery of science, or provide clues to the social and behavioral structures that may make it difficult for some groups of people to achieve success and visibility in the field (Knobloch-Westerwick, Glynn, & Huge, 2013; Moss-Racusin, Dovidio, Brescoll, Graham, & Handelsman, 2012). In this paper, I report efforts to understand how some typical features of scientific discourse—hedging and boosting—vary with respect to characteristics of a paper and its lead author.

Hedging and Boosting

One linguistic phenomenon in scientific discourse that has received some attention in the linguistics (Hyland, 1996, 1998) and computer science communities (Farkas, Vincze, Móra, Csirik, & Szarvas, 2010) is the way in which a claim is contextually amplified or diminished. Formally, these linguistic patterns are known as *boosting* (amplifying) and *hedging* (diminishing) respectively. From the perspective of knowledge accumulation in scientific literature, these linguistic devices are of interest because they may indicate statements that have more or less empirical support than statements without such devices. For instance, if the author of the paper concludes with “... *the intergroup contact hypothesis is supported by our data*” that suggests a much different strength of evidence than “... *the intergroup contact hypothesis may be supported by our data*”. From a psychological perspective of the scientific process, hedging and boosting hold interest because of the way these linguistic devices can be used to drive opinion and influence others. For instance, some work shows that the use of hedges is negatively associated with perceptions of power and persuasion (Blankenship & Holtgraves, 2005; Burrell & Koper, 1998; Hosman & Siltanen, 2011). Additional work suggests that hedges are associated with decreases in the persuasiveness of the message itself under certain conditions (Gibbons, Busch, & Bradac, 1991). Boosting is a somewhat rarer occurrence in scientific discourse (Hyland, 1998; Peacock, 2006), and has received less empirical attention, but qualitative work in linguistics suggests that it serves a complementary persuasive function (Vázquez Orta & Giner, 2009). Additionally, both of these linguistic features may be used to more carefully craft a narrative around the research, emphasizing and de-emphasizing for rhetorical or persuasive effect

Hedging and boosting are not only interesting because they influence how others interpret the evidence; they are also interesting as markers of linguistic *power*. There is a rich body of research on the antecedents and consequences of powerful versus powerless language use (Blankenship & Holtgraves, 2005; Haleta, 1996). Some features that mark powerless language include increases in qualifiers, tag questions (e.g. “*don’t you think?*”) and hedges,

while the absence of these mark powerful language. Powerful and powerless language are differentially used based on circumstances and individual differences. For instance, supreme court justices have been shown to use relatively powerful language when speaking to a lawyer arguing a case before the court, but when addressing the chief justice of the court, these same judges revert to a relatively powerless form of language (Danescu-Niculescu-Mizil, Lee, Pang, & Kleinberg, 2012). Additionally, women have been shown to be less likely than men to use powerful language (Carli, 1990), although this could reflect a general bias for women to occupy lower status positions than a gender effect per-se (Simkins-Bullock & Wildman, 1991).

Informed by this literature, there are two classes of paper meta-features that I propose may be related to hedging and boosting in interesting ways. First are markers of scientific behavior itself, namely the presence of statistical errors, and statistical overfitting (described below). The second are markers of power and status - the gender of the author, the prestige of the institution, and the number of citations an article has received. To the extent that the language in a paper is indicative of how science is approached and conducted, or the relative status of the scientist, one would expect that hedging and boosting would be prime candidates to reflect some of these phenomena. In the next section, I describe how the markers of scientific behavior would be expected to be associated with these linguistic features according to two hypotheses. In the subsequent section, I describe how power and status are expected to be associated with hedging and boosting according to the same hypotheses.

Errors and Overfitting. A statistical error, for the present purpose, is defined as an incongruous set of numbers from a statistical test. For instance, when the reported p-value does not match with the p-value derived based on the reported degrees of freedom and t or f statistic. Overfitting (more typically called *p-hacking*, though after Yarkoni and Westfall (In Press) I prefer overfitting as a less loaded term) reflects a set of behaviors where the analyst has a model fitting procedure that is more complicated than can be supported by

the data. Any behavior or technique that yields additional flexibility (e.g. including/excluding covariates or participants, selective reporting of variables or conditions, pursuing further analyses of a dataset contingent upon the results of previous analyses of the same dataset, etc.) is a potential contributor to overfitting. How might hedging and boosting be associated with the scientific behaviors of making statistical errors and overfitting the data?

One possibility, which I will refer to as the *linguistic power* hypothesis, is that hedging and boosting—as markers of power—reflect a relationship between the training one has received and the likelihood of success in academia. This basic idea has been thoroughly explored at a structural level by Smaldino and McElreath (2016). Their work roughly indicates that the structure of rewards in science favors those who find statistically significant results and publish many papers, with relatively few rewards for accuracy. Additionally, their work suggests that punishments for publishing irreproducible research have only minimal effects on success in science. Consequently, the system of promotion and reward in science is one that leads to a preponderance of labs that are good at publishing a lot of papers, but have very little incentive to conduct research that faithfully captures the true state of the world.

Methodological training in psychology has historically matched this environment. Researchers have been explicitly encouraged to conduct exploratory, low-powered studies with ill-defined analysis plans. Researchers were then taught to write up these studies for publication as if the research was confirmatory and the analysis was planned that way all along. This type of training encourages statistical overfitting. That this type of training should flourish is not particularly surprising - universities, departments, and individual academic mentors would like their graduates to go on to obtain good jobs, which may only happen if the researcher obtains statistical significance and publishes. Whatever circumstances drive the presence of overfitting are also likely to contribute to the presence of statistical errors insofar as the errors tend to lead to similar outcomes - improved prospects for publication and promotion. Since errors are known to be biased in such a way that favors

statistically significant effects (i.e., errors in mistakenly rejecting the null are more likely than errors in mistakenly failing to reject the null (M. B. Nuijten et al., 2015)), this suggests that the same reasoning should apply to them. To be clear, it is not that individuals intentionally make errors to improve their prospects, but rather that errors occur, and those that are consistent with achieving success are less likely to be scrutinized as closely as those that are barriers to achieving success. In this view, statistical reporting errors and statistical overfitting should be more likely from researchers who have attained some measure of power by obtaining jobs and being promoted¹. Given this, one would expect that their language use reflects their relatively high-status positions. If this is the case, then we would expect to find that hedging is negatively associated and boosting is positively associated with these statistical features of a paper.

To summarise the linguistic power hypothesis, individuals have been more likely to achieve success and power if they practice science in such a way that maximizes the number of publications and significant findings. In doing this, these scholars have unknowingly statistically overfit their data and been more likely to overlook errors that are favorable to publication than errors that are not favorable. Because hedging and boosting are thought to mark linguistic powerlessness and powerfulness, respectively, this hypothesis predicts that statistical errors and statistical overfitting are negatively associated with hedging and positively associated with boosting.

A second possibility is what I will call the *linguistic overfitting* hypothesis. The central idea behind this hypothesis is that hedging and boosting are used to amplify and diminish

¹One might suggest that this also reflects a cohort effect based on when people received their PhD, but broad appreciation of these problems did not exist until relatively recently, leaving very little time for those trained under more recent philosophies to have matured into positions of power. Instead, this hypothesis suggests that variation in power (and consequently, language), is more likely to be tied to variation in the degree to which individuals tended to favor maximization of publication and statistical significance, as this tendency will better represent the slice of academic practice captured by the papers examined in the data here (published in 2014).

claims in the paper in a way that better suits a theoretical story or leads to a more compelling narrative. To the extent that editors and reviewers act in ways that select for publication of exciting, surprising, and significant findings, then it is logical that individual researchers will do their best to write papers that fit this expectation. In this view, a researcher is interested in their paper communicating an interesting and surprising story about the world and human behavior. In order to craft such a story, the write-up of the research will involve an additional degree of emphasizing some ideas or results and de-emphasizing others. At the extreme, this rhetorical strategy is similar to the quantitative outcomes of overfitting - a given set of data can be used to craft many different stories of what the research means. If the linguistic overfitting hypothesis is an accurate description of behavior, then one would expect hedging and boosting to be positively associated with statistical errors and statistical overfitting, as they are all means of improving the likelihood of a paper's acceptance and publication.

It should be emphasized that neither of these hypotheses require any nefarious or ill-intentioned behavior on the part of individuals. Instead, they are both expectations of a system optimized to distribute rewards based on statistical significance and quantity of publications.

Power and Status in Science. As in many fields, men occupy a higher status level in science. Social and institutional biases ensure that men receive preferential treatment in a number of ways. For instance, the same work is perceived as higher quality when authored by a man (Knobloch-Westerwick et al., 2013), male students are perceived as more competent (Moss-Racusin et al., 2012), and male scientists are rated more highly in grant applications (R. van der Lee & Ellemers, 2015). As such, maleness confers a status and power advantage in science. Additionally, and more obviously, institutional ranking confers a high degree of status, advantage, and privilege broadly within academia (Clauset, Arbesman, & Larremore, 2015, Pinheiro, Melkers, and Newton (2017)). Finally, the number of citations a paper receives is, for better or worse, one of the primary metrics through which its impact

185 is judged.

186 If the linguistic power hypothesis is an apt description for academic communication,
187 then one would expect authors from higher-ranked institutions, male authors, and highly
188 cited papers to use greater degrees of boosting and reduced degrees of hedging. On the other
189 hand, if hedging and boosting are instead used as tools to craft stories around research in the
190 fashion described by the linguistic overfitting hypothesis, then the expectation is that
191 hedging and boosting, as tools used in the pursuit of publication and status would be more
192 prevalent in papers authored by individuals from higher-ranked institutions, and papers with
193 more citations. This expectation is on account of the fact that the practice of overfitting
194 should lead individuals to obtain more success in the field, as measured by institutional
195 prestige and citation counts. The prediction for male authors in comparison to female
196 authors is less clear under the linguistic overfitting hypothesis. One might predict that males
197 have achieved their position of privilege in part by succumbing to the institutional pressures
198 to present research in a particular way. The expectation here would then be that male
199 authors would be more likely to use hedging and boosting. On the other hand, one can just
200 as easily predict that males are *less* susceptible to the institutional pressures to present
201 research in a particular way, and are therefore expected to hedge and boost less than female
202 authors. As such, findings with respect to gender will not be able to distinguish between
203 these hypotheses.

204 Here, I use computational methods to investigate a collection of psychology journal
205 papers, extracting the relevant portions of the manuscripts, and explore the degree to which
206 the linguistic indicators of hedging and boosting are associated with meta-features of the
207 text and of the author as described above. While this work is largely exploratory, the
208 expectation going into the project was that the evidence would be congruent with the
209 linguistic power hypothesis. Although the pattern of results is not completely consistent
210 with either hypothesis, the linguistic overfitting hypothesis, developed after the data was
211 analyzed, appears to be better supported by the data, especially with respect to overfitting

and statistical errors. However, it deserves mention that the hypotheses are not mutually exclusive. It is possible, and even likely that both hypotheses are operating, and the patterns seen in these data reflect variation in the strength of these effects across different factors.

Disclosures

The work described here was not preregistered. Data and code can be found at <https://osf.io/5xwe6/>. I report how sample size was determined, as well as all data exclusions and measures used. I have no conflict of interest with respect to this work. I wrote the code to gather and analyze data and wrote the manuscript. I acknowledge Mark Conley, Jonathan Cook, Maneeza Dawood, Hale Forster, Rebecca Mohr, Maya Rossignac-Millon, Kate Turetsky, Matti Vuorre, and Katherine Zee for comments on previous versions of this manuscript. Acknowledgements also to Vanessa Anyanso, Julia Bottesini, Kyle Johnson, and Yeji Park for help with organizing, improving code efficiencies, and providing labels for training the classifiers. This manuscript was first deposited to PsyArXiv (<https://osf.io/preprints/psyarxiv/qasde/>).

Methods

Data Collection

I contacted the American Psychological Association and asked permission to automatically download articles from their website. After receiving this permission, I downloaded all articles from the year 2014 in 18 of their journals (see Figure 1) using the scrapy framework². These journals were chosen because they covered a range of research areas, were variable in their level of impact, and tended to publish empirical articles. In all, I obtained 1,583 articles from this process.

²<https://scrapy.org>

Figure 1. Number of articles from each journal

Data Filtering

Having done this, I next went through two filtering steps to limit the investigation to text that was directly relevant to the question of how the narrative portions of empirical work used hedging and boosting as linguistic devices. That is, I excluded non-empirical papers and focused on the introduction and discussion portions of the papers, as this is where hedging and boosting occurred most frequently³ (see Figures 2 and 3). To accomplish this, I used a machine learning approach. Details of these procedures can be found in the supplementary material. At the conclusion of this procedure, I was left with 1467 empirical articles.

³Filtering was only performed for the measurement of hedging and boosting, and not for evaluating overfitting or whether there were errors

Figure 2. Distribution of hedged sentences by position in document (sentence_number/ $n_{sentences}$) and journal. Journals are ordered left-to-right, top-to-bottom based on the count of articles.

Measurement

To measure prevalency of hedging and boosting, I used lists of words and phrases from previous work. The list of hedges comes from Prokofieva and Hirschberg (2014), who in turn sourced it from Prince, Frader, Bosk, and others (1982). Some minor changes have been made to reflect the fact that scientific papers often include references to first-person plural pronouns (e.g. I added *in our understanding* in addition to the already present *in my understanding*). The list of boosts were taken from Hyland (1998) without change as they were developed specifically for academic articles. Because the lists contain phrases as well as individual words, I measured the presence of hedging and boosting at the sentence level. A

Figure 3. Distribution of boosted sentences by position in document (sentence_number/ $n_{sentences}$) and journal. Journals are ordered left-to-right, top-to-bottom based on the count of articles.

sentence was marked as hedged (or boosted) if it contained a word or phrase from the relevant list. The hedge and boost lists may be viewed in the supplementary material.

I measured two types of markers related to the statistical results of the paper. I ran each paper through Statcheck (S. Epskamp & Nuijten, 2015), and marked each paper for whether it contained a statistical mistake or not. I also extracted all the reported p-values from each paper and sorted them into two bins: between .03 and .04 inclusive and above .04 up to .05. Following an approach described by Head, Holman, Lanfear, Kahn, and Jennions (2015), who took advantage of the distribution of p-values described by Simonsohn et al. (2014), I marked a paper as likely to have been overfit if the number of p-values in the upper bin exceeded the number of p-values in the lower bin. Simonsohn et al. (2014) describe this

logic in detail, but the computation hinges on the fact that a “true” effect will have a right-skewed distribution of p-values (i.e. small values are more likely). If there are more p-values just below the .05 cut-off than there are further below the cutoff, then this is a pattern that is consistent with the data having been quantitatively overfit.

I measured three other properties of the papers. I obtained the number of citations each paper had received by querying the Web of Science’s API. I established the gender of the first author by extracting the name and running the names through a multicultural name-gender association database (Elmas, 2013), yielding a count of 624 female and 677 male first authors. An additional 150 papers had first authors whose gender was identified as androgynous or unknown. I then matched the listed affiliation of the first author to the 2014 Shanghai Academic Rankings of World Universities. Details of this matching methodology can be found in the supplementary material. Of all listed affiliations, the matching procedure left 427 institutions unmatched, all of which were then designated as unranked universities.

Modeling

To evaluate the degree to which these meta-characteristics are related to prevalence of hedging and boosting, I fit a bayesian multilevel model separately for each linguistic marker. Each model predicted proportion of sentences hedged/boosted for each paper from the number of citations (standardized within journal), the presence of a statistical error (0=no error; 1=error), the presence of overfitting (0=not overfit; 1=overfit), the gender of the first author (dummy coded, with female as the baseline), and the institutional rank of the first author. Because the Shanghai rankings data is composed of inconsistent groupings through the range of ranks (e.g. ranks 1-100 are enumerated, but ranks 101-150, 151-200, 201-300, 301-400, and 401-500 are not differentiated within group), I created bins of institutional rankings that were consistent with the rankings approach published in the Shanghai data. I treated these bins (1-25; 26-50; 51-75; 76-100; 101-150; 151-200; 201-300; 301-400; 401-500; 500+) as a continuous numeric variable (i.e. 0-9), with the top four bins of my own creation,

and the remaining bins coming directly from the Shanghai rankings. All universities and other institutions not captured by the matching approach described above were designated as ranked 500+. As this latter category is difficult to interpret, the statistical model included an interaction term to allow this final group to vary from the linear fit for the other ranking groups.⁴

The nature of this model is largely similar to a standard multilevel model, though because of the bayesian nature, the types of statements we can logically make are somewhat different. In particular, rather than basing all conclusions on an assumption of a null hypothesis being true, it is instead possible to ask the question of interest more directly. Specifically, the posterior distribution of a bayesian model describes what parameter values are most likely to be true, in light of the data. For each result, I will present three metrics. The mean of the posterior distribution is presented as a summary of the center of the distribution and gives a sense of the estimated magnitude of the effect. The 95% uncertainty intervals represent the lower and upper points of the posterior distribution that define where 95% of the mass exists. Heuristically, this defines the most probable values for the parameter. For key results, I'll also present the probability of a negative or positive effect (depending on the direction of the estimate). This is computed as the proportion of the posterior distribution that lies below or above zero, respectively. Heuristically, we can interpret this as our confidence that the effect is negative or positive, given the data.

⁴Given that the work is largely exploratory, I had no specific hypotheses for interactions between these variables. It seems important to at least allow for such possibilities, but including all possible interactions would yield a model with 107 parameters, excluding the intercept. Such a model is difficult to interpret in a setting where little is known about the underlying phenomena. To strike a balance between the extremes of a simple model with just main effect predictors and one with an excess of parameters, I fit a model that placed horseshoe priors on all possible 2-way interaction terms ($p = 18$) (Piiironen & Vehtari, 2015). This model is the bayesian analogue of a strongly-regularizing LASSO regression (Tibshirani, 1996). Though a number of posteriors were strongly skewed away from zero, in no case did more than 69% (hedging) or 68% (boosting) of the posterior distribution lie on one side of zero. Given this, interactions were not explored any further in these analyses.

The multilevel structure of the model is defined such that papers are nested within journals, and all estimated parameters were allowed to vary by journal. These models were fit using Stan (vers. 2.15.0; Stan Development Team, 2016) and the rstanarm interface (vers ; Gabry & Goodrich, 2016). All plots were generated with ggplot (vers. ; Wickham, 2009) and bayesplot (vers.); Gabry, 2016). Data collection and processing was conducted using the R programming language (vers. 3.3.2; R Core Team, 2015) and Python 2.7. Classification models were fit using Scikit-learn (vers. 0.18.1; Pedregosa et al., 2011). All reported results are transformed to percentages

Results

There was a generally higher level of hedging ($M_{hedge} = 32.90, SD = 9.10$) than boosting ($M_{boost} = 17.60, SD = 6.70$) in the data. However, caution should be taken when interpreting this finding, as the dictionary for hedges is slightly longer (83 words and phrases) than that for boosting (64 words and phrases).

Hedging

Figure 4 displays the posterior distributions for the hedging model. As is apparent from the figure, the association between hedging and citations was estimated to be close to zero (posterior mean = -0.20) with a relatively high amount of uncertainty (95% uncertainty interval = $[-0.80, 0.40]$). Indeed, the posterior distribution tells us that the probability of a negative effect is $Pr_{negative_effect} = 0.74$. Similarly small or highly uncertain effects are present for the difference between female first authors and those whose gender is unknown ($M = 0.20, [-1.40, 1.80]$), institutional ranking (mean effect of decreasing a level in rankings = $0, [-0.30, 0.20]$), and the unranked interaction term ($M = -0.10, [-0.30, 0.10]$).

On the other hand, estimates for some other parameters suggest stronger associations. The presence of a statistical error, as detected by Statcheck is positively associated with hedging (mean = $0.90, [-0.20, 1.90]$, $Pr_{positive_effect} = 0.95$). The presence of statistical

Figure 4. Posterior distribution estimates for hedging model. Dark vertical line is the mean of each distribution. Shaded area is 95% uncertainty interval

overfitting shows an estimate that is in the same direction, though with slightly greater uncertainty (mean = 0.60, $[-0.70, 1.90]$, $Pr_{positive_effect} = 0.82$)).

To understand these effects in real terms, the average paper has approximately 383 sentences. The number of sentences in the introduction and discussion sections averages 166. Given this, a paper written by a woman with the average number of citations, without an error, not overfit, and from a top-ranked institution (i.e. the baseline paper in the statistical model) would be expected to have about 55 hedged sentences in these sections, A paper with an error would be expected to have about 56 hedged sentences. A paper that has been overfit would also be expected to have 56 hedged sentences.

Having a male first author is also associated with a decrease in the amount of hedging. Papers with a male first author are estimated to have a hedge rate that is 0.20 percent below

female first author papers $[-2.40, -0.30]$, $Pr_{negative_effect}.41$).

Converting these results to real language, in comparison to the baseline paper that contains 55 hedged sentences in the introduction and discussion sections, a paper whose first author is male, would be expected to have 52 hedged sentences in these sections.

When plainly interpreted as above, these effects may seem underwhelming. However, they are realistic in terms of effect sizes expected with vocabulary differences in large linguistic corpora (see supplementary material for a comparison with gender effect sizes seen in Newman, Groom, Handelman, and Pennebaker (2008)).

Boosting

Figure 5 displays the posterior distributions for the boosting model.

Figure 5. Posterior distribution estimates for boosting model. Dark line is the mean of each distribution. Shaded area is 95% uncertainty interval

The associations between boosting and presence of an error (mean = 0.90, [0.20, 1.70], $Pr_{positive_effect}$.99 and boosting and presence of overfitting (mean = 0.40, [-0.50, 1.40], $Pr_{positive_effect}$ = 0.82 are both estimated to be in the same direction as observed in the hedging model.

More concretely, a paper written by a woman with the average number of citations, without an error, not overfit, and from a top-ranked institution (i.e. the baseline paper in the statistical model) would be expected to have about 29 boosted sentences in the introduction and discussion sections. A paper with an error would be expected to have about 31 boosted sentences in these sections. A paper that has been overfit would be expected to have 30 boosted sentences.

In contrast, the estimates for the effect of citation (mean = -0.10, [-0.40, 0.30], the first author having unknown gender (mean = -0.30, [-1.50, 0.90], the first author being male (mean = 0, [-0.80, 0.70], institutional ranking -0.10, [-0.30, 0.10], and the ranking interaction term (mean = -0.10, [-0.30, 0.10] are estimated to be very close to zero, or with sufficiently high certainty to make any directional statements unwarranted.

Discussion

The findings described suggest that the presence of statistical errors and, though more weakly, that statistical overfitting share a similar relationship with the use of hedging and boosting. Papers with errors and with overfitting have more hedged sentences and boosted sentences than papers without errors or overfitting. This systematic difference suggests that the occurrence of these phenomena is a feature of scientific communication in the psychology literature sampled here and that some underlying approach to conducting and communicating science drives both of these occurrences. This observation is consistent with the linguistic overfitting hypothesis.

The linguistic overfitting hypothesis suggests that hedging and boosting are used to emphasize and de-emphasize themes in the narrative used to describe and interpret the

research. The writing in research papers is influenced by the same types of pressures that lead researchers to engage in statistical overfitting. In the case of linguistic overfitting, the writing is characterized by more language used to emphasize and de-emphasize themes in the narrative of the research. The same structural pressures that lead researchers to publish papers with statistical overfitting and statistical errors also encourage them to have results that are cleanly interpretable with respect to theory.

The finding that the occurrence of hedging language varies between genders is in line with previous work showing that men tend to use less tentative language (Carli, 1990), and is broadly consistent with the linguistic power hypothesis. However, in contrast to the predictions of this hypothesis, and to previous work (Miller, Cooke, Tsang, & Morgan, 1992; Rudman, 1998), I did not observe a corresponding increase in boosting. The estimated relationships between citations and the linguistic features are relatively uncertain. Without more data, it's difficult to know what direction this relationship truly takes. Additionally, it may be the accumulated citations are yet a noisy measurement of each paper's influence. Perhaps examining papers that have had a longer opportunity for impact would yield more robust estimates of this relationship.

Given these data, strong claims are also not possible with other markers of power and status. Although there were reasons to expect individuals from higher ranked universities, or papers with more citations to yield differing rates of hedging and boosting, that was not observed in these data. This tells us that any systematic relationship between these variables is likely to be relatively small, and we would need more data to detect it reliably. Although the findings here support the idea that the language in a scientific paper is associated with approaches to how science is conducted, the evidence that it is associated with the relative status of the scientist is weaker.

Holistically, the evidence is more strongly in favor of the linguistic overfitting hypothesis, though I emphasize again that these hypotheses are not mutually exclusive. Rather, they may both be operating and yielding effects differentially depending on feature

and circumstances. Knowledge of scientific communication patterns may benefit from additional work with planned analyses between these features at longer timescales, with more data, and examining specific interactive effects or relationships between the covariates used here.

As with any research endeavor, this work has some limitations. The first is that dictionary methods for text analysis, as were used here, are very crude. The words and phrases found in the hedging and boosting dictionaries may not always be a hedge or a boost. Measurement of hedging and boosting is not a straightforward problem for the computational linguist, but more sophisticated methods such as employing shallow parsing or taking advantage of distributed word representations (Garten et al., 2017; Mikolov, Sutskever, Chen, Corrado, & Dean, 2013) could yield benefits over those employed here. However, this concern is mitigated somewhat by the fact that the boosting dictionary was developed specifically for scientific articles.

Additionally, the measurement of overfitting employed here has been previously criticized as being insensitive to certain varieties of overfitting (Bishop & Thompson, 2016). Though I agree that a p-curve is a coarse tool, to my knowledge it is the only tool that can be applied to this question in this circumstance. Additionally, the nature of my investigation is not so much one of determining the prevalence of these behaviors, as past researchers have done. Instead, I'm connecting this behavior with other characteristics of the paper, and even if a particular *type* of overfitting is not accounted for here, those that are show a consistent relationship with the language investigated. If this pattern were to change with a more precise tool, then this would be suggestive of some relationship between the linguistic phenomena explored here and the type of overfitting behavior the researchers engaged in.

Constraints on Generality

The units of analysis for this research was a sample of empirical psychology articles published in 18 APA journals in 2014. I expect that the general findings will hold for

empirical psychology articles published in 2014, and perhaps to articles of other social and behavioral scientific disciplines as well. Given that publishing practices are presently undergoing rapid change, I suspect that the patterns observed here may be changing as time progresses. Thus, I am hesitant to declare that these results continue to hold with work to be published in the future. For similar reasons, these relationships may be different for articles published in the distant past.

To my knowledge, this is the first work examining how linguistic contextualizers are associated with meta-features of a paper. Additionally, I know of no other work exploring how statistical reporting errors or overfitting are related to scientific communication. This issue is worthy of further exploration. Though there is a good deal of work examining the methodological problems researchers used due to the pressures to tell a story consistent with a pre-conceived or newsworthy hypothesis (Gelman & Loken, 2014), there is less work examining how that story is actually told. The current results suggest that linguistic features connected with emphasis and de-emphasis systematically vary with meta-features of the paper, such as the gender of the author, and whether the paper is likely to have been overfit.

An important unanswered question is whether hedging and boosting are desirable qualities. Like many real life phenomena, the answer is complicated. If hedging and boosting are calibrated to the claims of the paper, few would argue that they are being misused. On the other hand, hedging and boosting to warp the details of a paper around theory or to warp theory around quantitative findings could be less desirable. There is an appropriate balance to strike between a completely transparent yet wooden reporting of the results, and a write up that selectively amplifies or downplays ideas and findings in the service of making research and theory work in harmony. The work presented here cannot identify where that balance is, but does provide empirical evidence that this continuum exists and that it appears to be associated with markers of scientific reliability.

Although the strength of the data should be the most important property of any scientific investigation, the written description of the research is also a powerful contributor

to how the research is perceived. Most consumers of science appreciate that variation in writing style is responsible for manuscripts that range from terrifically readable to those that are difficult to comprehend. However, this variation in writing style is also one additional, perhaps unavoidable avenue through which inefficiencies may be introduced into the truth-finding mechanisms of science.

References

- Bishop, D. V., & Thompson, P. A. (2016). Problems in using p-curve analysis and text-mining to detect rate of p-hacking and evidential value. *PeerJ*, 4, e1715.
- Blankenship, K. L., & Holtgraves, T. (2005). The role of different markers of linguistic powerlessness in persuasion. *Journal of Language and Social Psychology*, 24(1), 3–24.
- Burrell, N. A., & Koper, R. J. (1998). The efficacy of powerful/powerless language on attitudes and source credibility. In *Persuasion: Advances through meta-analysis* (pp. 203–215). Hampton Press Cresskill, NJ.
- Carli, L. L. (1990). Gender, language, and influence. *Journal of Personality and Social Psychology*, 59(5), 941.
- Clauset, A., Arbesman, S., & Larremore, D. B. (2015). Systematic inequality and hierarchy in faculty hiring networks. *Science Advances*, 1(1), e1400005.
- Danescu-Niculescu-Mizil, C., Lee, L., Pang, B., & Kleinberg, J. (2012). Echoes of power: Language effects and power differences in social interaction. In *Proceedings of the 21st international conference on world wide web* (pp. 699–708). ACM.
- Elmas, F. (2013). SexMachine. Retrieved from <https://github.com/ferhatelmas/sexmachine/>
- Epskamp, S., & Nuijten, M. (2015). Statcheck: Extract statistics from articles and recompute p values (r package version 1.0. 1.).
- Farkas, R., Vincze, V., Móra, G., Csirik, J., & Szarvas, G. (2010). The conll-2010 shared task: Learning to detect hedges and their scope in natural language text. In

484 *Proceedings of the fourteenth conference on computational natural language learning—Shared*
485 *task* (pp. 1–12). Association for Computational Linguistics.

486 Gabry, J. (2016). *Bayesplot: Plotting for bayesian models*. Retrieved from
487 <https://CRAN.R-project.org/package=bayesplot>

488 Gabry, J., & Goodrich, B. (2016). *Rstanarm: Bayesian applied regression modeling via*
489 *stan*. Retrieved from <https://CRAN.R-project.org/package=rstanarm>

490 Garten, J., Hoover, J., Johnson, K. M., Boghrati, R., Iskiwitch, C., & Dehghani, M.
491 (2017). Dictionaries and distributions: Combining expert knowledge and large scale textual
492 data content analysis. *Behavior Research Methods*, 1–18.

493 Gelman, A., & Loken, E. (2014). The statistical crisis in science data-dependent
494 analysis—a “garden of forking paths”—explains why many statistically significant
495 comparisons don’t hold up. *Am Sci*, 102(6), 460.

496 Gibbons, P., Busch, J., & Bradac, J. J. (1991). Powerful versus powerless language:
497 Consequences for persuasion, impression formation, and cognitive response. *Journal of*
498 *Language and Social Psychology*, 10(2), 115–133.

499 Haleta, L. L. (1996). Student perceptions of teachers’ use of language: The effects of
500 powerful and powerless language on impression formation and uncertainty. *Communication*
501 *Education*, 45(1), 16–28.

502 Head, M. L., Holman, L., Lanfear, R., Kahn, A. T., & Jennions, M. D. (2015). The
503 extent and consequences of p-hacking in science. *PLoS Biol*, 13(3), e1002106.

504 Hosman, L. A., & Siltanen, S. A. (2011). Hedges, tag questions, message processing,
505 and persuasion. *Journal of Language and Social Psychology*, 341–349.

506 Hyland, K. (1996). Writing without conviction? Hedging in science research articles.
507 *Applied Linguistics*, 17(4), 433–454.

508 Hyland, K. (1998). Boosting, hedging and the negotiation of academic knowledge.
509 *TEXT*, 18(3), 349–382.

510 Knobloch-Westerwick, S., Glynn, C. J., & Huge, M. (2013). The matilda effect in

science communication: An experiment on gender bias in publication quality perceptions and collaboration interest. *Science Communication*, 35(5), 603–625.

Lee, R. van der, & Ellemers, N. (2015). Gender contributes to personal research funding success in the netherlands. *Proceedings of the National Academy of Sciences*, 112(40), 12349–12353.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems* (pp. 3111–3119).

Miller, L. C., Cooke, L., Tsang, J., & Morgan, F. (1992). Should i brag? Nature and impact of positive and boastful disclosures for women and men. *Human Communication Research*, 18(3), 364–399.

Moss-Racusin, C. A., Dovidio, J. F., Brescoll, V. L., Graham, M. J., & Handelsman, J. (2012). Science faculty’s subtle gender biases favor male students. *Proceedings of the National Academy of Sciences*, 109(41), 16474–16479.

Newman, M. L., Groom, C. J., Handelman, L. D., & Pennebaker, J. W. (2008). Gender differences in language use: An analysis of 14,000 text samples. *Discourse Processes*, 45(3), 211–236.

Nuijten, M. B., Hartgerink, C. H., Assen, M. A., Epskamp, S., & Wicherts, J. M. (2015). The prevalence of statistical reporting errors in psychology (1985–2013). *Behavior Research Methods*, 1–22.

Open Science Collaboration and others. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716.

Peacock, M. (2006). A cross-disciplinary comparison of boosting in research articles. *Corpora*, 1(1), 61–84.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., . . . Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine*

Learning Research, 12, 2825–2830.

Piironen, J., & Vehtari, A. (2015). Projection predictive variable selection using stan+
r. *arXiv Preprint arXiv:1508.02502*.

Pinheiro, D. L., Melkers, J., & Newton, S. (2017). Take me where i want to go:
Institutional prestige, advisor sponsorship, and academic career placement preferences. *PloS
One*, 12(5), e0176977.

Prince, E. F., Frader, J., Bosk, C., & others. (1982). On hedging in
physician-physician discourse. *Linguistics and the Professions*, 8, 83–97.

Prokofieva, A., & Hirschberg, J. (2014). Hedging and speaker commitment. In *5th intl.
workshop on emotion, social signals, sentiment & linked open data, reykjavik, iceland*.

R Core Team. (2015). *R: A language and environment for statistical computing*.
Vienna, Austria: R Foundation for Statistical Computing. Retrieved from
<https://www.R-project.org/>

Rudman, L. A. (1998). Self-promotion as a risk factor for women: The costs and
benefits of counterstereotypical impression management. *Journal of Personality and Social
Psychology*, 74(3), 629.

Schooler, J. W. (2014). Metascience could rescue the replication crisis. *Nature*, 515, 9.

Simkins-Bullock, J. A., & Wildman, B. G. (1991). An investigation into the
relationships between gender and language. *Sex Roles*, 24(3), 149–160.

Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). P-curve: A key to the
file-drawer. *Journal of Experimental Psychology: General*, 143(2), 534.

Smaldino, P. E., & McElreath, R. (2016). The natural selection of bad science. *Royal
Society Open Science*, 3(9). doi:10.1098/rsos.160384

Stan Development Team. (2016). Stan modeling language users guide and reference
manual. Retrieved from <http://mc-stan.org>

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the*

563 *Royal Statistical Society. Series B (Methodological)*, 267–288.

564 Vázquez Orta, I., & Giner, D. (2009). Writing with conviction: The use of boosters in
565 modelling persuasion in academic discourses. *Revista Alicantina de Estudios Ingleses*.

566 Wickham, H. (2009). *Ggplot2: Elegant graphics for data analysis*. Springer-Verlag New
567 York. Retrieved from <http://ggplot2.org>

568 Yarkoni, T., & Westfall, J. (In Press). Choosing prediction over explanation in
569 psychology: Lessons from machine learning. *Perspectives in Psychological Science*.