

# Statistical Inferences for Polarity Identification in Natural Language

**Journal Article****Author(s):**

Pröllochs, Nicolas; Feuerriegel, Stefan; Neumann, Dirk

**Publication date:**

2018-12-21

**Permanent link:**

<https://doi.org/10.3929/ethz-b-000309052>

**Rights / license:**

[Creative Commons Attribution 4.0 International](#)

**Originally published in:**

PLoS ONE 13(12), <https://doi.org/10.1371/journal.pone.0209323>

RESEARCH ARTICLE

# Statistical inferences for polarity identification in natural language

Nicolas Pröllochs<sup>1\*</sup>, Stefan Feuerriegel<sup>2</sup>, Dirk Neumann<sup>3</sup>

**1** Department of Engineering Science, University of Oxford, Oxford, United Kingdom, **2** Chair of Management Information Systems, ETH Zurich, Zurich, Switzerland, **3** Chair for Information Systems Research, University of Freiburg, Freiburg, Germany

\* [nicolas.prollochs@eng.ox.ac.uk](mailto:nicolas.prollochs@eng.ox.ac.uk)



## Abstract

Information forms the basis for all human behavior, including the ubiquitous decision-making that people constantly perform in their every day lives. It is thus the mission of researchers to understand how humans process information to reach decisions. In order to facilitate this task, this work proposes LASSO regularization as a statistical tool to extract decisive words from textual content in order to study the reception of granular expressions in natural language. This differs from the usual use of the LASSO as a predictive model and, instead, yields highly interpretable statistical inferences between the occurrences of words and an outcome variable. Accordingly, the method suggests direct implications for the social sciences: it serves as a statistical procedure for generating domain-specific dictionaries as opposed to frequently employed heuristics. In addition, researchers can now identify text segments and word choices that are statistically decisive to authors or readers and, based on this knowledge, test hypotheses from behavioral research.

## OPEN ACCESS

**Citation:** Pröllochs N, Feuerriegel S, Neumann D (2018) Statistical inferences for polarity identification in natural language. PLoS ONE 13 (12): e0209323. <https://doi.org/10.1371/journal.pone.0209323>

**Editor:** Søren Wichmann, Leiden University, NETHERLANDS

**Received:** April 5, 2018

**Accepted:** December 4, 2018

**Published:** December 21, 2018

**Copyright:** © 2018 Pröllochs et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** All relevant data is available from GitHub (<https://github.com/nproellochs/SentimentDictionaries>).

**Funding:** The authors received no specific funding for this work.

**Competing interests:** The authors have declared that no competing interests exist.

## Introduction

The power of word choice and linguistic style is undisputed in the social sciences. For instance, linguistic style provides a means for deception [1, 2]. Likewise, marketing professionals have long understood the value of choosing the right terms when advertising products. For example, the use of technical terms facilitates the success of print advertisements [3]. Similarly, the valence of messages helps to explain consumer behavior. Here, the use of more positive expressions in user- and marketer-generated content in social media has a clear impact on purchase decisions [4]. The subtleties of language also receive increasing attention in the financial domain. In a recent study, [5] manipulate the tone of corporate news in a randomized controlled experiment and find that subjects expect a higher future return from a given firm when reading an article skewed towards positive language.

Psychological research has found that negative terms, especially, are vital in forming impressions, perceptions and attributions [6]. For instance, subjects use more positive-emotion words in self-disclosures; yet negative-emotion words have a significantly greater impact on formed impressions [7]. Further works by Pennebaker and his colleagues shed light on the

use of linguistic terms as psychological markers of personality and personal states [8]. As an illustrative example, linguistic style serves as a predictor of age, gender, mood, emotion and mental health. This and other findings stem from the calculation of the occurrences of certain, aggregated word categories (e. g. cognitive words, past tense, pronouns). However, though not all of the words are likely to be relevant, there is a scarcity of resources that identify the decisive entries within these categories.

While the above applications demonstrate the great importance and need for profound language understanding, the reception of individual words and their effects on human behavior remain subject to research [8, 9, 10]. When studying the reception of natural language, researchers commonly utilize a document-level score that measures the overall perception of natural language, including negative wording, tone, sentiment, moods and emotions (e. g. [11, 12, 13]). However, this does not allow for a granular understanding of how individual pieces of information are perceived within narratives. In fact, understanding word choice, the perception of wording and the corresponding human responses present open questions for research, and especially in terms of rigorous statistical inferences [14].

Related research in the area of social sciences commonly relies on manually selected dictionaries (e. g. [15, 16, 17]). Prevalent examples are the Harvard IV psychological dictionary from the General Inquirer software or Linguistic Inquiry and Word Count (LIWC). These contain lists of words that classify terms as either positive or negative based on human judgments, which makes them prone to severe drawbacks. Foremost, the word lists are selected *ex ante* based on subjective opinions of their authors. They thus can be neither comprehensive nor as precise as statistical rigor. Furthermore among these is the fact that the labor-intensive process of their construction prevents dictionaries from being tailored to arbitrary domain-specific applications. Moreover, dictionaries rarely discriminate between different levels of positivity (or negativity), since the underlying words are merely grouped into two classes of positive and negative expressions without further weighting.

To overcome the previous shortcomings, this paper proposes a novel approach that utilizes LASSO regularization to extract words that are statistically decisive based on an outcome variable. Examples include ordinal ratings on review portals, which summarize the connotation of user-generated comments, or the stock market reaction, which assesses investors' perceptions of financial materials. Our approach specifically builds upon these response variables as they mirror narrative content in an accurate and objective manner. Here we extend our previous work [18, 19] and introduce statistical inferences to identify cues that convey a positive or negative polarity. At the same time, the analysis can be replicated for the prose of arbitrary applications in order to adapt to the domain-specific particularities.

This work immediately reveals manifold implications for social sciences and behavioral research: first, our approach offers a tailored means by which to study the perception of language and word choice through the eyes of readers and authors with statistical rigor. The results are highly interpretable and serve as an input to further hypothesis tests. After all, this contributes to behavioral research by addressing the crucial question of how textual information impacts individual behavior and decision-making.

The remainder of this paper is organized as follows. Section 2 provides background literature concerning the reception of natural language, which necessitates a statistical approach to measure the response to word choice. As a remedy, we present our LASSO-based methodology in Section 3. Subsequently, Section 4 demonstrates the value of this approach with examples from recommender systems and finance. Section 5 then provides thorough comparisons, followed by Section 6 with implications for behavioral hypothesis testing. Section 7 discusses the advantages and limitations of our method and provides detailed implications for both theory and practice.

## Backgrounds

This section posits that extracting statistically relevant terms based on a decision variable is both an innovative and relevant research question to the social sciences. Therefore, we review previous works and methods concerned with measuring the reaction to word choice. We also outline how our approach differs from opinion mining, which gives a lever to measure subjective information in narrative content.

### Relationship to opinion mining

Drawing inferences regarding how wording relates to a decision variable is closely related to the concept known as sentiment analysis or opinion mining. It refers to the use of natural language processing as a way to extract subjective information from narrative content. The underlying methods aim at measuring the semantic orientation (i. e. the positivity and negativity) of the overall text, or with respect to a particular topic or aspect [20]. The result is then either a continuous sentiment score or else a classification as positive or negative. The surveys in [12] and [21] provide a comprehensive, domain-independent overview of common methodological choices. These techniques can primarily be grouped into two categories, namely, approaches utilizing pre-defined dictionaries or machine learning.

The former, dictionary-based approaches, mainly serve explanatory purposes, especially when a response variable is not present. They extract subjective information from the occurrences of pre-defined polarity words, which are selected *ex ante* based on the intuition of experts. This creates an approach that is not only straightforward, but also produces reliable and interpretable results in various applications (e. g. [22, 23]). Previous research has devised several variants with different scopes and objectives (cf. next section for an overview). These dictionaries can be combined with linguistic rules that specifically account for linguistic modifiers that signal, for instance, uncertainty or activation [10].

Machine learning methodologies utilize a baseline variable to train a predictive model, which is later applied to unseen documents where it should predict the semantic orientation. Previous research has tested various models, including support vector machines and artificial neural networks, that typically take (transformed) word frequencies as input (e. g. [21, 24]). As a result, machine learning often achieves a high predictive accuracy but might suffer from overfitting. In addition, it remains a black-box with low interpretability and hardly any insights into its reasoning.

The above approaches target applications in which whole texts are classified according to their semantic orientation. Thereby, sentiment analysis either serves explanatory or predictive purposes, which have both become prevalent in behavioral research. These methods work at document level (or aspect level); however, they cannot draw statistical inferences at word level, which is the goal of research aimed at understanding the reception of word choice at a granular level.

### Overview of common dictionaries

Gaining insights into the subtle differences between word choice requires methods that analyze narrative content at a granular level. Therefore, a common strategy is to build upon manually selected dictionaries from previous research. In this vein, humans label terms as either positive and negative or, alternatively, according to other semantic, syntactic or psychological categories. Table 1 provides an overview of prevalent dictionaries in behavioral research. For example, the Harvard IV dictionary from the General Inquirer software comprises various psychological categories beyond positive and negative valence: e. g. emotions, strength, or overstatement. LIWC was designed to identify emotion-laden writing but also measures linguistic style based on expressions that were individually assigned to over 70 linguistic dimensions by

**Table 1. Common dictionaries in behavioral research.**

Dictionary	Size	Categories	Domain	Selection Process	Polarity Levels	Notes
Diction	10,000	35 linguistic categories (e. g. optimism, satisfaction, praise, blame, denial)	Politics	Expert judgment	Binary	Accessible for purchase via the Diction software for text analysis
Harvard IV	4206	15 linguistic categories (e. g. polarity, motivation, pleasure, pain, cognitive orientation)	Psychology	Expert judgment	Binary	Shipped in General Inquirer
LIWC	4500	64 linguistic dimensions (e. g. polarity, part-of-speech, cognitive and psychological words)	Psychology	Independent judges	Binary	Accessible for purchase from the LIWC text analysis software
Loughran-McDonald	2709	Polarity (positive, negative)	Finance	Manual selection procedure	Binary	Based on 2of12inf dictionary
QDAP	6789	Polarity (positive, negative)	General	Heuristic based on co-occurrences to positive/negative seed words	Binary	Synset of WordNet
SentiStrength	763	Positivity, negativity	Social media	Human judgment	Continuous rating	Derived from LIWC
SentiWordNet 3.0	28,431	Positivity, negativity, neutrality	General	Heuristic based on co-occurrences to positive/negative seed words	Continuous rating	Based on 86,994 terms from WordNet

<https://doi.org/10.1371/journal.pone.0209323.t001>

independent judges. Other dictionaries are devoted to domain-specific applications, such as the Loughran-McDonald dictionary, which consists of polarity terms found in earnings reports. With the exception of SentiStrength, SentiWordNet and QDAP, the dictionaries usually cannot differentiate between different degrees of polarity among words.

In order to computerize the construction of dictionaries, researchers have devised various rule-based approaches and heuristics, which are frequently referred to as dictionary generation or lexicon creation. On the one hand, several algorithms follow a semi-supervised approach that considers word embeddings, similarity or co-occurrences between terms. For instance, SentiWordNet, as well as QDAP, starts with a small set of seed words labeled as positive or negative, based on which neighboring terms are classified [25]. On the other hand, some algorithms base their classifications on a response variable (the gold standard). This sounds similar to our statistical procedure, but they then propose the use of heuristics to label words depending on their appearances in documents rated with a high or low gold standard. The underlying heuristics adapt concepts from information retrieval, such as information gain, pointwise mutual information and  $\chi^2$ -based selection (e. g. [26]). The heuristics aim at differentiating varying degrees of sentiment strength; however, they lack statistical justification, which impairs the possibility of drawing any reliable inferences.

In addition to the above shortcomings, only a small portion of the content of the dictionaries in Table 1 overlaps and some even contain contradictory entries. As a result, choosing the most suitable dictionary to facilitate an understanding of written information is challenging and any choice is likely to be imperfect. This is particularly relevant, since words often feature a highly domain-specific meaning. The above elaborations immediately reveal that there is no one “right” dictionary, and authors of [27] argue that the state-of-the-art methods for polarity scoring are subpar, which affects sentiment-related analysis and conclusions drawn from it.

### Statistical approaches for dictionary generation

The objective of this work is to come up with a statistical procedure that deduces the true perception of explicit and implicit polarity terms. The few existing approaches entail several

statistical deficiencies. [28] count frequencies (tf-idf) of selected words and then insert them into an ordinary least squares estimation with a gold standard. However, this approach is subject to multicollinearity and, hence, the authors decided to restrict their analysis specifically to words that appear in the Loughran-McDonald finance-specific dictionary. [29] and [30] develop variants of multinomial regressions that can handle high-dimensional count data. However, both are limited to categorical outcome variables, which makes them infeasible in our setting. Furthermore, the multinomial regressions only work with absolute term frequencies, instead of using common weighting schemes from information retrieval (e. g. tf-idf), which are often regarded as more efficient. Lastly, the underlying dimension reductions return loadings in the reduced-space, which allows for the ranking of word polarities, but lacks direct statistical interpretation (e. g. standard errors).

We later draw upon the LASSO as a procedure for extracting decisive variables. This method has been applied to textual input, but merely in predictive settings, where it serves either as tool for weighting the salience of predictive features (e. g. [31]) or black-box forecasts (e. g. [32]). However, we adapt it for explaining outcomes *ex post*. To the best of our knowledge, it has not been combined with statistical confidence estimates or proposed as technique for measuring the reception of language. Beyond an earlier draft [18], the use of the LASSO has, in particular, been neither propagated as a tool for generating domain-specific dictionaries nor experimentally evaluated against manual dictionary annotations.

## Research gap

Altogether, we see that the above research neglects to draw rigorous statistical inferences from a comparison between word choice and the regressands. As a remedy, we develop a regularization technique to select granular polarity expressions from documents that *statistically* elicit a positive or negative response. It even extracts terms that convey valence implicitly, helps to discriminate between subtle differences in polarity strength, and adapts to domain-specific particularities—all in order to enable an in-depth analysis of the relationship between language and decisions. This ultimately contributes a better understanding of human text processing.

Previous literature has pointed out the need for understanding the reception of natural language. Related works predominantly draw on manual and labor-intensive procedures in which human judgments are assumed to reflect the ground truth. The outcome of this process usually results in a set of positive and a set of negative cues, which one refers to collectively as a dictionary. However, there is no doubt that such a setup is error-prone as perceptions of individuals and experts are eminently subjective and thus biased. These dictionaries also entail further shortcomings. First, they usually struggle to capture domain-specific characteristics. For instance, a finance-specific dictionary cannot distinguish between language describing developments in the real estate market in comparison to the linguistic style of technology firms. In addition, most dictionaries also presuppose an equal importance across all words in the same polarity group and thus do not exhaust a continuous bandwidth of sentiment levels.

## Method development

This section proposes a novel methodology by which to investigate the granular perception of natural language and to examine the textual cues that trigger decision-making. Our methodology comprises two stages, of which the first step performs several preprocessing operations to transform running text into a document-term matrix. The second step performs a variable selection to extract only the relevant terms. This essentially utilizes a LASSO regression, treating each document as an observation, while we use all words as explanatory factors explaining an exogenous response variable. We have released our method publicly in the

form of an R package. The package *SentimentAnalysis* is available for download via CRAN: <https://cran.r-project.org/package=SentimentAnalysis>.

### Preprocessing of natural language

The preprocessing phase transforms the running text into a structured format that allows for further calculations. This includes a myriad of standard routines from natural language processing (cf. online appendix for details). For instance, we remove stop words without a deeper meaning and truncate inflected words to their stems [24].

We then obtain frequencies  $x_{d,t}$  of how often term  $t$  occurs in document  $d$ . In order to focus only on the characteristic terms in a document, we transform the frequencies  $x_{d,t}$  by using a common weighting scheme from information retrieval, namely, the term frequency-inverse document frequency or tf-idf for short [24]. Thereby, the raw frequency  $x_{d,t}$  is weighted by the ratio of the total number of documents divided by the number of documents that contain the term  $t$ , i. e.

$$\hat{x}_{d,t} = tf-idf(t, d, D) = tf(t, d) idf(t, D) = x_{d,t} \log \frac{|D|}{|\{d \in D \mid t \in d\}|}, \quad (1)$$

given a corpus,  $D$ , of documents. We have also tested different variants of using the raw term frequencies as part of our robustness checks; however, these result in a slightly lower goodness-of-fit and thus yield inferior prediction performance in both datasets.

### Model specification

Let  $y$  denote the gold standard that measures our response variable of interest. We now construct a linear model where the number of occurrences of individual words explains the response variable. That is, we specify a linear model

$$y = \beta_0 + \sum_{t=1}^n \beta_t \hat{x}_t + \varepsilon \quad (2)$$

to quantify the effect of words  $\hat{x}_t = [\hat{x}_{1,t}, \dots, \hat{x}_{|D|,t}]^T$  for  $t = 1, \dots, n$  on the dependent variable  $y$  with error term  $\varepsilon$ . In addition, we standardize the word variables  $\hat{x}_t$  in order to facilitate comparison between coefficients. The estimated coefficients  $\beta_0, \dots, \beta_n$  then gauge the effect of words on that gold standard.

Estimating the above model is not trivial, since the appearance of words is likely to be highly correlated, i. e.  $|\text{cor}(\hat{x}_i, \hat{x}_j)| \gg 0$  for  $i \neq j$ . This raises serious issues of multicollinearity and, consequently, the ordinary least squares (OLS) estimator can be misleading. Moreover, it also results in low predictive power [33] and entails limited interpretability when facing a large number of variables [34].

In order to overcome these statistical challenges, we perform regularization via the least absolute shrinkage operator (LASSO). Regularization can serve as a viable alternative to OLS when the number of regressors is large and highly correlated. As our main contribution to the existing body of literature on this topic, we propagate the application of this regularization approach to word frequencies in order to infer decisive words and interpret them statistically. Alternative estimators entails disadvantages, since, for instance, ridge regression and elastic net perform no variable selection and can thus not benefit from parsimonious models.

### Reasoning behind regularization

Theory from natural language builds upon Zipf's law according to which word counts follow a power law distribution [24]. It further suggests that certain words have a potentially large

impact, while a high number elicit only a marginal response [35]. To conduct an analysis focusing on those relevant words, we need a mathematical mechanism that extracts terms deemed important. To this end, regularization is a common tool for implicit variable selection and has recently gained considerable traction in data science [36]. A common choice is the LASSO [33, 34, 37], since it identifies covariates that fit the data best, while simultaneously shrinking some coefficients to zero.

The LASSO entails several properties that makes its use beneficial for estimating our model. First of all, the LASSO automatically identifies decisive variables in a linear model. It thus chooses a subset of variables and filters out non-informative model terms. In our setting, this allows us to discard words that are statistically not relevant with respect to the exogenous variable. This property of variable selection leads to parsimonious and more interpretable models. At the same time, the LASSO mitigates the issue of multicollinearity, which is present when estimating the model via ordinary least squares. Additionally, by finding a reasonable trade-off between bias and variance, it solves the problem of overfitting, which occurs if the model complexity is too high [33, 34, 37].

The LASSO can be identically formalized both as an OLS estimator with an additional regularization parameter or as Bayesian model with a specific prior distribution. The LASSO has recently been extended by significance tests [38]. Alternatively, one can utilize standard errors from the Post-LASSO procedure [39].

On the whole, the LASSO specifically enables us to treat each distinct word from a corpus as a potential regressor. Its use, together with the standard errors, thereby introduces statistical inferences to natural language on a word-by-word level.

### Statistical inferences from word choice

The LASSO incorporates an additional regularization term that penalizes non-zero coefficients [33, 34, 37], given by a minimization problem

$$\beta_{\text{LASSO}} = \arg \min_{\beta_0, \dots, \beta_n} \sum_{i=1}^{|D|} \left[ y_i - \beta_0 - \sum_{t=1}^n \beta_t \hat{x}_{d,t} \right]^2 \text{ s. t. } \sum_{t=1}^n |\beta_t| \leq \lambda \quad (3)$$

with a suitable tuning parameter  $\lambda$ . The magnitude of the regression coefficients measures the perception of individual words statistically.

Because of the  $L_1$ -penalty, the LASSO typically produces estimates in which some of the coefficients are set exactly to zero and, thereby, performs an implicit feature selection. In practice, the parameter  $\lambda$  is selected using cross-validation to find a value that minimizes the error on the hold-out set. Afterwards, we re-fit the model with that specific  $\lambda$  using all the observations in order to determine its coefficients. Our standard errors stem from the Post-LASSO, and allow us to make statistical tests that correspond to the use of specific words.

As a result, our procedure identifies statistically relevant words, while the corresponding coefficients measure their polarity. One major benefit of our approach is that it overcomes the problem of ex ante selected words. Hence, we no longer run the risk of labeling words for subjective reasons or on the basis of erroneous knowledge, since all outcomes measure the influence of words on the dependent variable with statistical validation.

### Empirical results

This section evaluates our method with two studies from different domains: (I) we investigate the role of word choice in recommender systems by extracting opinionated terms from user-generated reviews. (II) We further study the impact on stock markets of the wording in

financial disclosures. Subsequently, we compare the resulting word lists to the manually-selected dictionaries from previous research and show how our method can be used with higher-order word combinations to incorporate context. For all websites exploited for the collection of data, we complied with the terms of service.

## Study I: Opinionated terms in user-generated reviews

**Corpus with reviews.** The first study demonstrates the added value of our approach in the domain of recommender systems, where we automatically infer terms that convey opinions. Professionals from marketing can exploit these expressions to gain insights into how people judge products or services. In related research, movie reviews represent a popular choice when it comes to studying opinion mining (e. g. [40]). Among the reasons is that movie reviews pose a particularly difficult challenge, since they often contain a mixture of feedback, critique and summaries of movie scenes. For example, positive reviews still refer to some unpleasant scenes and negative reviews to pleasant ones. We utilize a free and publicly available corpus of 5006 movie reviews from the Internet Movie Database (IMDb), each annotated with an overall rating. The scaled dataset is available from <http://www.cs.cornell.edu/people/pabo/movie-review-data/>. All reviews are written by four different authors and preprocessed, e. g. by removing explicit rating indicators [40].

**Statistical inferences for polarity word scoring.** We now extract opinionated terms from the movie reviews. The corpus contains a total number of 1195 word stems after preprocessing. Our methodology results in a final model with 549 (47.21%) statistically relevant terms. Out of these, 294 terms feature a positive and 255 a negative connotation. Unsurprisingly, the coefficients are generally small as a single word does not flip the whole meaning of the document but merely of a sentence. We report the top 15 expressions with the highest and lowest coefficients in Table 2. The table lists stems instead of complete words due to stemming being part of the preprocessing. We additionally calculate standard errors via the Post-LASSO [39].

Table 2 renders it possible to precisely discriminate different levels of positive and negative polarity strength. Many of the listed terms seem plausible and might be used independent of the context of a movie review, such as *perfect* or *bad*. These words frequently appear in sentences, such as “*the story is perfect*” or “*this is just a bad film*”. In addition, we observe a large number of words that are specific to the domain of motion pictures. This includes terms, such as *recommend* and *long*, that, for instance, occur in sentences such as “*the movie was too long*”. However, other terms, such as *war* (coefficient of 0.0041) or *crime* (coefficient of 0.0004) appear unexpected at first glance. A potential reason is that these words are often related to certain actions and scenes that appeal to the audience and are—on average—more positively perceived than other parts in the plot.

Furthermore, Table 2 states, in percentage, how often each word occurs in reviews with positive or negative ratings. For instance, the term *best* appears in 65% of all positive reviews and *brilliant* in 73% of the cases. Yet the pure number of appearances is misleading: the term *best* amounts to a much higher coefficient of 0.0571 compared to 0.0480, thereby indicating that it expresses a more positive sentiment. We note here again that both the response variables, as well as our regressors, are standardized for easier comparisons.

Our model features a relatively high explanatory power with an adjusted  $R^2$  amounting to 0.5668. We also see clear indications of multicollinearity in the model prior to performing variable selection, since 18 (1.51%) out of all the variance inflation factors exceed the critical threshold of 4, hence, making regularization a vital ingredient of our procedure.

Table 2 also compares the inferred polarity score to expert judgments. Evidently, there is a considerable number of opinionated terms that are not covered by dictionary word lists. Among

Table 2. Empirical results of top 15 opinionated terms in movie reviews.

Word Stem	Coef.	Stand. Error (Post-LASSO)	Relative Freq. (%)	Positive Doc. (%)	Negative Doc. (%)	Harvard IV
POSITIVE TERMS (TOP 15)						
great	0.0709	0.0098	31.58	66.29	33.71	⊕
perfect	0.0707	0.0096	19.18	74.17	25.83	⊕
excel	0.0572	0.0156	19.68	64.57	35.43	⊕
best	0.0571	0.0098	47.16	64.59	35.41	⊕
life	0.0551	0.0102	49.82	63.79	36.21	
delight	0.0515	0.0098	10.69	76.82	23.18	⊕
brilliant	0.0480	0.0095	7.19	73.06	26.94	⊕
intens	0.0469	0.0097	9.27	74.35	25.65	
uniqu	0.0416	0.0098	8.39	73.81	26.19	⊕
recommend	0.0393	0.0138	18.57	59.78	40.22	
marvel	0.0390	0.0096	5.19	79.62	20.38	⊕
hilari	0.0373	0.0096	6.97	75.64	24.36	⊕
easi	0.0353	0.0095	15.46	70.67	29.33	⊕
matur	0.0347	0.0102	10.49	74.67	25.33	⊕
fascin	0.0346	0.0099	10.59	77.74	22.26	⊕
NEGATIVE TERMS (BOTTOM 15)						
bad	-0.1124	0.0103	34.50	47.60	52.40	⊖
worst	-0.1011	0.0132	16.26	52.21	47.79	⊖
wast	-0.0762	0.0144	19.28	52.54	47.46	⊖
review	-0.0741	0.0169	53.28	51.48	48.52	
suppos	-0.0699	0.0097	15.66	41.07	58.93	
least	-0.0672	0.0097	22.67	47.58	52.42	
movi	-0.0671	0.0130	84.66	56.06	43.94	
cinematograph	-0.0538	0.0151	21.17	44.81	55.19	
flat	-0.0526	0.0096	6.31	35.44	64.56	
unfortun	-0.0512	0.0102	14.12	43.42	56.58	⊖
dull	-0.0483	0.0096	5.43	32.72	67.28	⊖
bore	-0.0483	0.0097	8.21	37.23	62.77	⊖
denni	-0.0468	0.0236	23.33	42.21	57.79	
lack	-0.0450	0.0097	16.48	48.61	51.39	⊖
wors	-0.0442	0.0097	7.11	38.48	61.52	⊖

Notes: This table reports the extracted terms that convey a particularly positive or negative sentiment in movie reviews. Top: the 15 most positive word stems, together with their estimated coefficient. Standard errors are calculated via the Post-LASSO [39]. Bottom: the 15 most negative word stems. In addition, we provide the relative frequency within the corpus, as well as the ratio of positive and negative documents that contain each word. The last column show the overlap with the Harvard IV psychological dictionary. The symbol “⊕” indicates terms that appear in the positive word list and “⊖” in the negative word list of this dictionary. The complete list with all 549 stems is given in the supplementary materials.

<https://doi.org/10.1371/journal.pone.0209323.t002>

the 15 most positive words, for example, only 12 have found their way in the Harvard IV psychological dictionary, whereas this is true for only 8 of the 15 most negative terms. We later detail the overlap for the complete list of terms in Section 5, finding only a minor consensus of 40.44%. This stems from the fact that authors commonly utilize implicit polarity words to express their opinions, which are not included in psychological dictionaries. This highlights the shortcomings of human dictionaries and provides strong evidence that authors convey their message by utilizing different and highly domain-specific wording to communicate their opinion.

## Study II: Impact of wording on financial markets

**Financial corpus.** Our second study demonstrates the reception of language in regulatory Form 8-K filings from the United States. These inform investors about important corporate events, such as management changes, the departure of directors, bankruptcy, layoffs, and other events deemed significant. Form 8-K filings are publicly accessible and quality-checked by the Securities and Exchange Commission (SEC) to ensure that the content meets formal requirements. These reports are of high relevance for the stock market and communicate very informative material [41]; this suggests a strong relationship between their content and market responses.

Our filings (including amendments) span the years 2004 to 2013, originating from the EDGAR website of the SEC. The complete sample consists of 901,133 filings, which then undergo several filtering steps. First, we select only filings from firms whose stocks were publicly traded on the New York Stock Exchange (NYSE). Second, in order to gain information about the stock market reaction, we remove filings for which we are not able to match the SEC CIK numbers to Thomson Reuters Datastream (from which all financial data is retrieved). Consistent with prior research, we exclude filings that contain fewer than 200 words and penny stocks below \$5 per share [42]. These filtering steps then result in a final corpus of 76,717 filings.

We measure the stock market reaction subsequent to a disclosure by the abnormal return of the corresponding company, since it corrects the nominal return for concurrent market movements. In short, we implement a market model that assumes a stable linear relation between market return and normal return. We model the market return using a stock market index, namely, the NYSE Composite Index, along with an event window of 10 trading days prior to the disclosure. The supplementary materials provide a thorough explanation of this approach.

**Statistical inferences for word reception.** We now report the cues that are relevant for the decision-making of investors when reading financial materials. Our approach selects a total of 172 statistically relevant terms, out of 1724 entries in the preprocessed corpus, i. e. 9.98%. Out of this subset, 82 entries are linked to a positive stock market response, 90 word stems to a decreasing firm valuation. Such a relatively small subset of decisive terms is in line with the suggestion from Zipf's law [24, 35]. We observe generally smaller coefficients as compared to our first study with movie reviews. This is not an unexpected result, since the average length of financial filings (3473 words) is higher than that of reviews (1066 words). Hence, the proportional influence of a single word as measured by the magnitude of its coefficient is smaller. Table 3 reports the 15 words with the highest and lowest coefficients based on our procedure, for which we again provide only stemmed words due to our preprocessing. As before, we additionally calculate standard errors via the Post-LASSO. The complete list is provided in the supplements.

Similarly to the previous corpus, we observe several terms that are specific to the given domain of financial reporting, e. g. *improv*, *strong payroll* and *lower*. These words crop up, for instance, in sentences such as “*the strong business development was sustainably confirmed*”. In contrast, we also find unexpected outcomes, which appear predominantly in the negative list. Examples include *although* (standardized coefficient of -0.0036) and *however* (standardized coefficient of -0.0015). Most likely, these cues convey uncertainty, attenuate other statements or overturn earlier expectations.

Overall, the current model features a lower explanatory power when compared to the previous model based on user-generated reviews. We expected such an outcome, since previous work has found that very few variables can predict stock returns in efficient markets [23]. In

**Table 3. Empirical results of top 15 polarity expressions in financial filings.**

Word Stem	Coef.	Stand. Error (Post-LASSO)	Relative Freq. (%)	Positive Doc. (%)	Negative Doc. (%)	Harvard IV	Loughran-McDonald
POSITIVE TERMS (TOP 15)							
improv	0.0325	0.0045	37.27	49.70	50.30	⊕	⊕
rais	0.0160	0.0038	11.06	51.17	48.83	⊖	
strong	0.0144	0.0045	28.41	50.32	49.68		⊕
increas	0.0113	0.0051	60.51	49.16	50.84		
facil	0.0106	0.0039	35.28	49.29	50.71		
waiver	0.0095	0.0039	15.33	48.33	51.67		
stronger	0.0080	0.0039	5.44	50.58	49.42		⊕
vacat	0.0076	0.0037	5.58	49.72	50.28		
repurchas	0.0074	0.0039	22.97	50.03	49.97		
favor	0.0073	0.0040	25.45	49.71	50.29	⊕	⊕
consumm	0.0067	0.0040	15.13	48.43	51.57	⊕	
annum	0.0056	0.0039	9.2707	48.00	52.00		
avoid	0.0051	0.0037	11.55	48.82	51.18	⊖	
payrol	0.0049	0.0037	6.69	49.32	50.68		
middl	0.0046	0.0037	5.16	49.15	50.85		
NEGATIVE TERMS (BOTTOM 15)							
declin	-0.0204	0.0045	23.59	48.65	51.35	⊖	⊖
negat	-0.0162	0.0040	20.03	47.85	52.15	⊖	⊖
lower	-0.0138	0.0047	27.27	48.72	51.28	⊖	
experienc	-0.0117	0.0038	12.17	47.93	52.07		
delay	-0.0091	0.0038	18.72	47.65	52.35	⊖	⊖
broad	-0.0063	0.0038	11.46	48.22	51.78		
advertis	-0.0056	0.0042	8.86	48.35	51.65		
project	-0.0055	0.0041	36.88	48.84	51.16		
pressur	-0.0055	0.0038	9.42	48.96	51.04		
now	-0.0054	0.0040	27.14	48.80	51.20		
challeng	-0.0054	0.0039	15.58	48.41	51.59	⊖	⊖
offer	-0.0052	0.0045	40.33	48.76	51.24	⊕	
depreci	-0.0051	0.0052	23.11	48.59	51.41	⊖	
impact	-0.0041	0.0046	39.79	48.62	51.38		
weak	-0.0039	0.0038	8.5841	48.35	51.65	⊖	⊖

Notes: This table reports verbal expressions that convey positive and negative information in financial disclosures (Form 8-K filings). Top: 15 most positive word stems, together with their estimated coefficient. Standard errors are calculated via the Post-LASSO [39]. Bottom: the 15 most negative word stems. In addition, we provide the relative frequency in financial filings, as well as the ratio of documents with a positive or negative market response. The last columns shows the agreement between our statistical inferences and two common dictionaries based on human annotations, namely, the Harvard IV psychological and Loughran-McDonald finance-specific dictionary. The symbol “⊕” indicates terms that appear in the respective positive word list, “⊖” in the negative one. The complete table with all 172 entries is given in the supplements.

<https://doi.org/10.1371/journal.pone.0209323.t003>

addition, we see again strong evidence of multicollinearity, since 24 (1.39%) of the variance inflation factors in the full model before variable selection exceed the critical threshold of 4. This stresses once more the need for regularization in our approach.

In addition, Table 3 compares the inferred polarity scores to the classifications from psychological and finance-specific dictionaries. As suggested by [35], we observe that dictionary labels deviate extensively from the true perception of stock market investors. From the 15 most positive words, only 5 words are also contained in the Harvard IV psychological

Table 4. Comparison of human classifications to statistical inferences.

	Size	Overlapping Terms		Consensus Classification		Correlation	Reliability
		Count	Share	Count	Share		
STUDY I: MOVIE REVIEWS							
Harvard IV	4206	222	0.4044	138	0.6216	0.3236 ***	0.2246
Henry	190	26	0.0474	20	0.7692	0.5593 **	0.5446
Loughran-McDonald	2709	73	0.1330	45	0.6164	0.4303 ***	0.2311
SentiWordNet	28431	440	0.8015	246	0.5591	0.2649 ***	0.1001
QDAP	6789	176	0.3206	114	0.6477	0.3638 ***	0.2863
STUDY II: FINANCIAL FILINGS							
Harvard IV	4206	55	0.3198	34	0.6182	0.2742 *	0.2270
Henry	190	21	0.1221	19	0.9048	0.6333 **	0.8102
Loughran-McDonald	2709	20	0.1163	18	0.9000	0.6433 **	0.8030
SentiWordNet	28431	118	0.6860	69	0.5847	0.2089 *	0.1715
QDAP	6789	40	0.2326	28	0.7000	0.4524 **	0.3939

Notes: This table compares common, human-generated word lists to extracted terms based on our statistical inferences. We omitted LIWC and Diction, since these are commercial products with proprietary dictionaries. When computing correlation coefficients and reliability scores, we exclude non-overlapping terms and count binary dictionary entries with a negative label as -1 and positive ones as 1. Statistical significance levels are \*\*\*0.001, \*\*0.01, \*0.05. Reliability (i. e. the concordance with our statistical inferences) is measured in terms of Krippendorff's alpha coefficient [43].

<https://doi.org/10.1371/journal.pone.0209323.t004>

dictionary, whereas this is true for 8 of the negative terms. Similarly, the Loughran-McDonald dictionary contains only 4 out of the 15 most positive words and 5 out of the 15 most negative words. A detailed comparison with other dictionaries is provided as part of our robustness checks in Section 5. Overall, this indicates that the human experts do not accurately judge the true reception of natural language in the financial domain.

### Comparison to dictionaries from human selection

We now compare the results of our statistical inferences to the manually selected dictionaries from previous research. For this purpose, Table 4 details the number of overlapping terms and compares to what extent classifications agree. In addition, we present the inter-rater reliability (i. e. the concordance with our statistical inferences) in terms of Krippendorff's alpha coefficient [43]. Here, a reliability value of 1 indicates a perfect overlap between the classifications in positive and negative groups, whereas a value of 0 denotes that human dictionaries and our statistical inferences are statistically unrelated.

The results demonstrate that the ex ante selected dictionaries show only a small overlap with the word lists from our statistical procedure. In the case of movie reviews, only 222 out of 549 (i. e. 40.44%) extracted words have found their way into the Harvard IV dictionary that is frequently utilized in IS and behavioral research. Out of these, only 62.16% actually exhibit the same polarity direction. This is in line with our in-depth investigations, since many negative expressions from this dictionary feature a positive connotation in the context of movie evaluations. Psychological dictionaries classify words, such as such as *crime*, *force* or *war*, in the negative list, while, in film reviews, these often refer in a positive sense to the suspense in certain scenes. Unsurprisingly, we find the highest number of overlapping terms in the dictionary that includes the most entries, i. e. the SentiWordNet. However, this dictionary shows the lowest reliability (0.10) and correlation (0.26) with our statistical inferences. In contrast, the highest reliability (0.54) and correlation (0.56) is achieved by the Henry dictionary which, however, consists of a mere 190 entries, resulting in a minor overlap of 26 words.

We observe similar results for our financial disclosures, where 55 out of 172 extracted words (i. e. 31.98%) also appear in the Harvard IV dictionary. Out of these, 61.82% feature the same direction. Overall, we find a correlation of 0.27 between the estimated coefficients and the binary Harvard IV dictionary (encoded as  $\pm 1$ ). Even the dictionaries that were specifically designed for financial reports reveal large deviations from the statistical inferences. We observe only a total number of 21 overlapping terms for the Henry dictionary, and 20 for the Loughran-McDonald dictionary. Nonetheless, compared to psychological dictionaries, we see that the finance-specific dictionaries are indeed more accurate in measuring the reception of words in financial disclosures. For example, the Loughran-McDonald dictionary shows a consensus classification of 90.48% and a correlation of 0.64 with our statistical inferences. Moreover, finance-specific dictionaries also yield the highest reliability. For example, the Henry dictionary shows a Krippendorff's alpha coefficient of 0.8102 (compared to e. g. 0.2270 for the Harvard IV).

Table 4 identifies a consistent disagreement between human classification and statistical selection. Although most ex ante dictionaries feature a large volume of words, many statistically relevant terms are not included. In addition, overlapping terms show a relatively low correlation that is, in some cases, only significant at the 5% level. As a consequence, misclassification and the erroneous exclusion of words limit the suitability of ex ante dictionaries.

The aforementioned dictionaries have frequently been utilized also in predictive settings and we thus also compare the out-of-sample performance of the above dictionaries with our method. We briefly outline the results here, while we provide further statistics and elaboration in our supplementary materials. In short, our method outperforms all of the investigated dictionaries for both movie reviews and financial disclosures. In the case of movie reviews, the best performing dictionary (Harvard IV) results in a 90.66% higher mean squared error compared to the LASSO. We observe a similar pattern for financial disclosures. These results thus reinforce our previous finding that manually selected dictionaries deviate from true perception.

## Statistical inferences with word phrases

Human-generated dictionaries commonly categorize only isolated words without incorporating any contextual information. However, the position of a word in a sentence is likely to contribute to the meaning and the overall interpretation. Consequently, related research attempts to work with higher-order word combinations, i. e. so-called  $n$ -grams. However, findings indicate mixed results regarding the extent to which their inclusion improves performance. Expert dictionaries refrain from labeling word pairs, since it requires considerable manual labor. Similarly, heuristics for dictionary creation are also rarely designed to process  $n$ -grams. This is in contrast to our statistical procedure, which works effortlessly with  $n$ -grams as the corresponding frequencies are simply inserted in the variable selection procedure. These benefits become particularly evident when considering the sheer number of input variables (2971 bigrams for financial filings and 1059 bigrams for movie reviews). Such large numbers of highly correlated predictors would imply serious overfitting issues for almost any type of statistical model without variable selection.

Table 5 compares the results from using  $n$ -grams. First of all, we observe fewer relevant bigrams than unigrams. In the case of unigrams, our method extracts 549 relevant terms from the movie reviews and 172 from the financial corpus, while using bigrams results in a total number of 442 terms for movie reviews and 51 for financial filings. We provide the complete lists of extracted phrases in the supplementary materials due to space limitations, but summarize a few intriguing insights here. For instance, the bigram with the highest

**Table 5. Summary statistics of statistical inferences with word tuples.**

	STUDY I: MOVIE REVIEWS	STUDY II: FINANCIAL FILINGS
<u>Bigrams</u>		
Regressors before regularization	1059	2971
Extracted terms	442.0000	47.0000
Ratio of extracted terms	41.7400%	1.5800%
Positive terms	234	19
Negative terms	208	28
Ratio positive terms	52.9400%	40.4300%
Ratio negative terms	47.0600%	59.5800%
Adjusted $R^2$	0.3184	0.0036
Correlation between model estimate and gold standard	0.6300	0.0800
<u>Bigrams &amp; unigrams</u>		
Regressors before regularization	2254	4695
Extracted terms	798	132
Ratio of extracted terms	35.4000%	2.8110%
Positive terms	394	62
Negative terms	404	70
Ratio positive terms	49.3700%	46.9700%
Ratio negative terms	50.6300%	53.0300%
Adjusted $R^2$	0.6126	0.0072
Correlation between model estimate and gold standard	0.8300	0.1100

*Notes:* The table compares our statistical inferences for different inputs, consisting of bigrams and the combination of unigrams and bigrams. These are evaluated in terms of goodness-of-fit and by comparing the number of selected entries. The complete lists of extracted variables and their coefficients are given in the supplements.

<https://doi.org/10.1371/journal.pone.0209323.t005>

positive coefficient in the review corpus is *best film*, while the most negative bigrams are *bad movie* and *waste time*.

According to Table 5, we also observe a drop in the adjusted  $R^2$  for both corpora. In the case of movie reviews, the adjusted  $R^2$  declines from 0.5668 for unigrams to 0.3184 for bigrams due to its penalty on the degrees-of-freedom. We observe a similar pattern for our financial corpus. Here, the adjusted  $R^2$  decreases from 0.0079 for unigrams to 0.0036 for bigrams. Finally, we also tested a configuration that incorporates both unigrams and bigrams. While this approach yields the highest fit for the review corpus, we observe a slightly inferior goodness-of-fit for the financial corpus. Altogether, this shows that our method is not limited to single terms, but also serves as an appropriate tool to study the influence of higher-order word combinations, and even phrases, on a response variable.

### Implications for hypothesis testing using natural language

Our method presents also a valuable tool for analyzing behavioral research questions. This section demonstrates two applications that allow for the testing of hypotheses with focus on word choice.

### Placement of negative information in movie reviews

We utilize our method to test where authors place negative statements in their reviews. Writers might start with negative thoughts, as suggested by the law of primacy in persuasion. On the

Table 6. Summary statistics for hypothesis testing with movie reviews.

	PANEL I: ALL REVIEWS			PANEL II: POSITIVE RATING			PANEL III: NEGATIVE RATING		
	$\mu_1$	$\mu_2$	$\mu$	$\mu_1$	$\mu_2$	$\mu$	$\mu_1$	$\mu_2$	$\mu$
Mean	0.1025	0.0578	0.1604	0.1626	0.1402	0.3027	-0.0098	-0.0930	-0.1028
Min.	-0.4159	-0.5922	-0.8903	-0.2020	-0.3977	-0.2947	-0.4159	-0.5922	-0.8903
25% Quantile	0.0159	-0.0517	-0.0173	0.0825	0.0507	0.1539	-0.0729	-0.1639	-0.2028
Median	0.0996	0.0528	0.1476	0.1557	0.1342	0.2896	-0.0021	-0.0849	-0.0947
75% Quantile	0.1853	0.1646	0.3331	0.2325	0.2251	0.4302	0.0587	-0.0178	-0.0014
Max.	0.7336	0.7655	1.3848	0.7336	0.7655	1.3848	0.3624	0.3322	0.3687
Std. Dev.	0.1340	0.1615	0.2569	0.1175	0.1357	0.2066	0.1060	0.1123	0.1604
Skewness	0.1717	0.1535	0.2202	0.4439	0.3028	0.5058	-0.2684	-0.3898	-0.3062
Kurtosis	0.6052	0.2682	0.2204	0.6407	0.5611	0.6167	0.5321	0.7152	0.6622

Notes: Panel I compares the sentiment of the first ( $\mu_1$ ) and second half ( $\mu_2$ ) of movie reviews, as well as the overall sentiment  $\mu$ . The additional panels present the same statistics for reviews with positive (Panel II) or negative (Panel III) gold standard only.

<https://doi.org/10.1371/journal.pone.0209323.t006>

other hand, one might be inclined to instead utilize the regency effect, according to which arguments presented last garner more attention. Given the overall movie rating, we can evaluate where authors place negative information when composing movie reviews, i. e. do they generally introduce negative aspects at the beginning or rather at the end?

*HYPOTHESIS: Negative information is more likely to be placed at the end than at the beginning of a review.*

In order to test this hypothesis, we compute the sentiment of the first and second half of each review by summing over products of coefficient and weighted term frequency. We refer to them as  $\mu_1$  and  $\mu_2$  respectively. Summary statistics of  $\mu_1$ ,  $\mu_2$  and the document sentiment  $\mu$  are shown in the first panel of Table 6. In addition, we present the same statistics for reviews that are filtered for a positive (Panel II) or negative (Panel III) gold standard only. We then test the null hypotheses

$$H_0 : \mu_1 < \mu_2 \quad (\text{law of primacy applied to negative content}), \text{ and}$$

$$H_0 : \mu_1 > \mu_2 \quad (\text{regency effect for negative content})$$

respectively.

According to our results, the second half of movie reviews generally conveys a more negative tone than the first half. The mean sentiment in the first half amounts to  $\mu_1 = 0.1025$ , whereas it is  $\mu_2 = 0.0578$  for the remainder part. The corresponding difference  $\mu_1 - \mu_2$  between both sentiment values is statistically significant at the 0.1% significance level when performing a two-sided Welch *t*-test (test statistic of 15.06). The results in Panel II and III follow a similar picture. For instance, Panel III shows that the first half of negatively rated movie reviews yields a negative sentiment of  $\mu_1 = -0.0098$  on average, while the second half results in an even more negative sentiment of  $\mu_2 = -0.0930$ . This difference is also significant at the 0.1% significance level with a *t*-value of 19.26. In Panel II, we observe a similar pattern for reviews with positive ratings (*t*-value of 6.70). We thus accept our hypothesis regarding the presence of a regency effect. This result also coincides with psychological research according to which senders of information are more likely to place negative content at the end [44], but, in contrast, our evidence is collected outside of an artificial laboratory setting, as it stems from actual human communication.

## Response of financial markets to non-informative wording

In our second application of hypothesis testing, we examine to what extent financial markets trade upon non-informative wording. Previous works have established a robust market response to fact-related information encoded in written materials, which is primarily measured by using the positive and negative word lists from Loughran-McDonald or Harvard IV. Yet it is unclear how the remaining words—which are not deemed as either positive or negative from an external standpoint and which we refer to as non-informative—are processed by markets. Consistent with classical economic theory, we expect that investors ignore these terms and, instead, solely focus on essential, fact-related information, i. e. clearly positive and negative cues.

*HYPOTHESIS: Financial markets are not distracted by the wording in corporate communication that falls outside the clearly delineated categories of positive and negative.*

Interestingly, we present empirical results in the following section which reject the above hypothesis and suggest the opposite. The extracted words from Table 3 list the polarity terms that are statistically relevant for the investment decisions of traders. However, most of them are not necessarily classified as positive or negative according to the Harvard IV psychological or Loughran-McDonald finance-specific dictionary. We thus test our hypothesis by grouping all words into two categories according to the previous dictionaries: one group contains all words that are labeled as either positive or negative. This group represents all terms that feature an explicit, fact-based statement. The remaining entries form a group that can be characterized as non-informative wording. For instance, the latter contains entries such as *although* and *however*. We find that the perception of investors depends on many terms that feature no explicit positive or negative statement polarity. According to the Harvard IV dictionary, only 31.97% of the extracted words can be associated with a fact-based meaning, whereas 68.03% of the extracted words are expected not to contribute to the informative content. The Loughran-McDonald dictionary presents a similar picture. Here, the fact-based group contains 11.63% of all extracted words, while the remaining 88.37% can be regarded as non-informative wording.

Finally, we perform an  $F$ -test to validate whether the subset of words that are neither labeled as positive nor negative has a combined effect on stock returns. In the case of the Harvard IV dictionary, this results in an  $F$ -statistic of 5.37, which is statistically significant at the 0.1% level. Similarly, the  $F$ -statistic for the Loughran-McDonald dictionary numbers to 5.48, which is also significant the 0.1% level. We must thus reject our hypothesis and provide evidence that expressions deemed as non-informative wording by previous research have a statistically significant effect on financial markets.

## Discussion

In the following, we discuss the implications of our research method as it not only improves understanding of natural language but also enables intriguing inferences in behavioral sciences. Furthermore, our research is highly relevant for practitioners seeking to operationalize natural language in Information Systems.

## Implications for behavioral sciences

Understanding decision-making and providing decision support both increasingly rely upon computerized natural language processing. In contrast to many black-box methods from the domain of machine learning, our methodology provides a vehicle for content analysis and opinion mining that is fully comprehensible for deep insights. Specifically, it allows one to maintain high interpretability as it explains an effect in terms of the presence of individual words. It thus allows researchers to dissect the relationship between natural language and a

given outcome variable. In addition, our approach goes beyond pre-defined dictionaries that classify words into groups of positive and negative words as we assign individual word weights to each word, thereby accounting for differences in the valence levels of words of the same polarity class.

Our results indicate that common, manually selected dictionaries from the literature, such as the Harvard IV psychological dictionary, are neither complete nor adequate for arbitrary domains. For instance, in the area of finance, they classify words as positive that are not necessarily interpreted positively by investors. To overcome these previous limitations, our methodology provides a means by which to automate the process of dictionary generation. Altogether, our study thus provides evidence that applications of dictionary-based sentiment analysis can be significantly improved when adapting the dictionaries to the corresponding domain.

## Applications

Analyzing the perceptions of word choice and understanding the response to natural language on a granular level can yield new insights in a large number of use cases. In the following points, we illustrate prominent applications in the areas of both practice and research:

- **Recommender systems.** Recommender systems support users by predicting their rating or preference towards products or services, and, similarly, product reviews on web platforms guide individuals considering purchases. Yet it is unclear which expressions actually convey an opinion, even though this would allow for a better understanding of how judgments are formed. Our statistical inferences aid enterprises in identifying success factors of products, while they present researchers new opportunities to study behavioral theories at word level.
- **Social and behavioral sciences.** In the context of social interactions, it is highly relevant to understand how humans express and perceive information in natural language. Our methodology helps to answer various questions, such as which wording drives word-of-mouth. Moreover, it enables the identification of word choices that convey information regarding personality and psychology or linguistic cues that are linked to deception in human communication.
- **Finance.** Before exercising ownership in stocks, investors often consult financial disclosures and pay especially close attention to their soft content, such as linguistic style. To analyze the language in financial materials, researchers, investors and automated traders utilize the simple categorization of terms as either positive or negative. However, working with black-box approaches or inferring the overall valence of a disclosure merely from term frequencies is prone to error, since companies often frame negative news using positive words. Our statistical procedure remedies this issue as it labels words based on their actual interpretation in financial materials. On the other hand, regulators can utilize our mechanism to put in place effective warning mechanisms for disclosures whose content can provoke critical market developments.
- **Marketing.** Practitioners in the field of marketing strive to understand how people perceive language in advertisements and press releases. Here, a granular understanding on a word-by-word basis would enable them to carefully consider phrasing in order to enhance sales. In addition, marketing teams could utilize our inference technique to make early predictions regarding the success of ad campaigns, product launches or the popularity of product attributes. In this vein, our method can identify words that influence customers in a positive or negative direction.

These examples highlight several prominent applications that benefit from a granular understanding of language at word level.

Ultimately, it is hoped that the contributions and advantages presented in this paper—such as quantifying the reception of language—will become an important tool in future research papers. Application of this method can yield novel insights into behavioral research questions regarding the information processing of natural language. This should help those in the field of social sciences to add to the growing body of knowledge on the role of behavior in individual decisions and population-wide outcomes, such as voting, consumer demand, information sharing, product evaluation and opinion aggregation. As demonstrated in this paper, our methodology has the potential to enable unprecedented opportunities in terms of validating behavioral research outside of existing laboratory setups. Yet it also fuels innovations in the theoretical advancement and formalization of theories as its high interpretive power facilitates new discoveries.

### Limitations and opportunities for future research

Our approach gauges the domain-specific effects of words or  $n$ -grams based on a decision variable. A sufficient number of labeled training data is thus a necessary prerequisite to extract statistically relevant terms from documents. However, at the same time, this presents also the strength of our approach as we explain the variation in the dependent variable through word use.

This work is targeted at the vast number of users of dictionary-based approaches. The objective behind dictionaries is that they obtain polarity scores that are context-independent. For instance, each term has the same polarity independent of whether it is a main-/sub-clause, the beginning/end of a text or used in metaphor. These are known issues in computer science, since an exact understanding of language remains a daunting undertaking. As a remedy, we suggest the following extension to our framework: if desired, one could extend the LASSO-based approach with a hierarchical formulation, such that terms are associated with context-specific polarity score; however, the resulting caveats are a larger corpora, the challenges from a context-dependent interpretation and the mismatch with the majority of dictionary-based use cases.

We implemented our statistical inferences as a L1-norm-based optimization problem with confidence intervals from the Post-LASSO. This framework comes with the following limitations, for which we also name straightforward remedies. One the one hand, dependent variables with other distributional assumptions can be easily handled by fitting the framework with a fully Bayesian approach that relies upon Markov chain Monte Carlo (MCMC) sampling. On the other hand, the Post-LASSO makes mathematical assumptions regarding the derivation of the confidence intervals. As an alternative, one can either revert to the significance test from [38] or again utilize a MCMC-based estimation. The latter is especially prone to co-occurrence patterns of words when computing confidence intervals. Notably, the two aforementioned adaptations do not change the overall model specification but merely exchange the underlying estimation technique.

### Summary

Understanding the decision-making of individuals, enterprises and organizations presents a fundamental pillar of behavioral research. However, the challenges associated with processing natural language have been largely associated with simple decision models featuring predominantly structured data. Yet an unparalleled source of information is encoded in unstructured formats, and especially textual materials. The reasons behind this are multifaceted, including

the recent advent of the big data era and the increasing availability of data through the World Wide Web, which has made a vast number of written documents—such as user-generated content and news—available to the public.

Past research has laid the groundwork for inferring the polarity of written contents, albeit in a manner that is usually limited to a few psychological dictionaries that classify single terms. Such approaches work almost out-of-the-box and thus seem promising at first, but entail inevitable and major shortcomings. The elements of these word lists are selected *ex ante* by manual inspection and subjective judgment. As such, our paper exposes the weaknesses of common dictionary methods: they only allow one to assess the overall polarity of documents and not of individual expressions, thereby leaving any deeper insights in the underlying text processing untapped. In addition, they often prove insufficient in adequately reflecting the domain-specific perception of a given audience.

As a remedy, this paper proposes the use of LASSO regularization as a form of variable selection to extract relevant words that statistically impact decisions. Social science researchers can greatly benefit from such a procedure, as it infers *ex post* relevant terms based on the outcome of a decision. It can therefore efficiently adapt to domain-specific peculiarities of narratives and discriminate between subtle polarity levels across words.

## Supporting information

**S1 Supplementary Materials. This document contains additional analyses and methodological details.**

(PDF)

## Acknowledgments

The valuable contributions of Ryan Grabowski are gratefully acknowledged.

## Author Contributions

**Conceptualization:** Nicolas Pröllochs, Stefan Feuerriegel, Dirk Neumann.

**Data curation:** Nicolas Pröllochs, Stefan Feuerriegel.

**Formal analysis:** Nicolas Pröllochs, Stefan Feuerriegel.

**Resources:** Dirk Neumann.

**Supervision:** Dirk Neumann.

**Visualization:** Nicolas Pröllochs, Stefan Feuerriegel.

**Writing – original draft:** Nicolas Pröllochs, Stefan Feuerriegel.

**Writing – review & editing:** Nicolas Pröllochs, Stefan Feuerriegel.

## References

1. Fuller CM, Biro DP, Burgoon J, Nunamaker J. An Examination and Validation of Linguistic Constructs for Studying High-Stakes Deception. *Group Decision and Negotiation*. 2013; 22(1):117–134. <https://doi.org/10.1007/s10726-012-9300-z>
2. Jensen ML, Averbek JM, Zhang Z, Wright KB. Credibility of Anonymous Online Product Reviews: A Language Expectancy Perspective. *Journal of Management Information Systems*. 2013; 30(1):293–324. <https://doi.org/10.2753/MIS0742-1222300109>
3. Anderson RE, Jolson MA. Technical Wording in Advertising: Implications for Market Segmentation. *Journal of Marketing*. 1980; 44(1):57. <https://doi.org/10.2307/1250035>

4. Goh KY, Heng CS, Lin Z. Social Media Brand Community and Consumer Behavior: Quantifying the Relative Impact of User- and Marketer-Generated Content. *Information Systems Research*. 2013; 24(1):88–107. <https://doi.org/10.1287/isre.1120.0469>
5. Bosman R, Kräussl R, Mirgorodskaya E. The “Tone Effect” of News on Investor Beliefs: An Experimental Approach. *CFS Working Paper Series*. 2015; 522(1):1–76.
6. Baumeister RF, Bratslavsky E, Finkenauer C, Vohs KD. Bad Is Stronger Than Good. *Review of General Psychology*. 2001; 5(4):323–370.
7. Berry DS, Hiller WS, Mueller JS, Pennebaker JW. Linguistic Bases of Social Perception. *Personality & Social Psychology Bulletin*. 1997; 23(5):526–548. <https://doi.org/10.1177/0146167297235008>
8. Pennebaker JW, Mehl MR, Niederhoffer KG. Psychological Aspects of Natural Language Use: Our Words, Our Selves. *Annual Review of Psychology*. 2003; 54:547–577. <https://doi.org/10.1146/annurev.psych.54.101601.145041> PMID: 12185209
9. Agarwal R, Dhar V. Big Data, Data Science, and Analytics: The Opportunity and Challenge for IS Research. *Information Systems Research*. 2014; 25(3):443–448. <https://doi.org/10.1287/isre.2014.0546>
10. Villarreal OF, Ludwig S, de Ruyter K, Grewal D, Wetzels M. Unveiling What Is Written in the Stars: Analyzing Explicit, Implicit, and Discourse Patterns of Sentiment in Social Media. *Journal of Consumer Research*. 2017; 43(6):875–894.
11. Bollen J, Mao H, Zeng X. Twitter Mood Predicts the Stock Market. *Journal of Computational Science*. 2011; 2(1):1–8. <https://doi.org/10.1016/j.jocs.2010.12.007>
12. Pang B, Lee L. Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval*. 2008; 2(1):1–135. <https://doi.org/10.1561/15000000011>
13. Tetlock PC. Giving Content to Investor Sentiment: The Role of Media in the Stock Market. *Journal of Finance*. 2007; 62(3):1139–1168. <https://doi.org/10.1111/j.1540-6261.2007.01232.x>
14. Lim EP, Chen H, Chen G. Business Intelligence and Analytics. *ACM Transactions on Management Information Systems*. 2013; 3(4):1–10. <https://doi.org/10.1145/2407740.2407741>
15. Shen W, Hu YJ, Rees J. Competing for Attention: An Empirical Study of Online Reviewers’ Strategic Behaviors. *MIS Quarterly*. 2015; 39(3):683–696. <https://doi.org/10.25300/MISQ/2015/39.3.08>
16. Stieglitz S, Dang-Xuan L. Emotions and Information Diffusion in Social Media: Sentiment of Microblogs and Sharing Behavior. *Journal of Management Information Systems*. 2013; 29(4):217–248. <https://doi.org/10.2753/MIS0742-1222290408>
17. Xu SX, Zhang XM. Impact of Wikipedia on Market Information Environment: Evidence on Management Disclosure and Investor Reaction. *MIS Quarterly*. 2013; 37(4):1043–1068. <https://doi.org/10.25300/MISQ/2013/37.4.03>
18. Pröllochs N, Feuerriegel S, Neumann D. Generating Domain-Specific Dictionaries Using Bayesian Learning. In: 23rd European Conference on Information Systems (ECIS). Münster, Germany; 2015.
19. Pröllochs N, Feuerriegel S, Neumann D. When Words Matter Most: Tailoring Domain-Specific Dictionaries with Decision Analytics. In: Conference on Information Systems and Technology (CIST 2015). Philadelphia, PA, USA; 2015.
20. Thet TT, Na JC, Khoo CSG. Aspect-Based Sentiment Analysis of Movie Reviews on Discussion Boards. *Journal of Information Science*. 2010; 36(6):823–848. <https://doi.org/10.1177/0165551510388123>
21. Ravi K, Ravi V. A Survey on Opinion Mining and Sentiment Analysis: Tasks, Approaches and Applications. *Knowledge-Based Systems*. 2015; 89:14–46. <https://doi.org/10.1016/j.knosys.2015.06.015>
22. Henry E. Are Investors Influenced by How Earnings Press Releases are Written? *Journal of Business Communication*. 2008; 45(4):363–407. <https://doi.org/10.1177/0021943608319388>
23. Tetlock PC, Saar-Tsechansky M, Macskassy S. More Than Words: Quantifying Language to Measure Firms’ Fundamentals. *Journal of Finance*. 2008; 63(3):1437–1467. <https://doi.org/10.1111/j.1540-6261.2008.01362.x>
24. Manning CD, Schütze H. *Foundations of Statistical Natural Language Processing*. Cambridge, MA: MIT Press; 1999.
25. Baccianella S, Esuli A, Sebastiani F. SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. In: *LREC 2010 Proceedings*; 2010. p. 2200–2204.
26. Oliveira N, Cortez P, Areal N. Stock Market Sentiment Lexicon Acquisition Using Microblogging Data and Statistical Measures. *Decision Support Systems*. 2016; 85(5):62–73. <https://doi.org/10.1016/j.dss.2016.02.013>
27. Abbasi A, Sarker S, Chiang RH. Big Data Research in Information Systems: Toward an Inclusive Research Agenda. *Journal of the Association for Information Systems*. 2016; 17(2):3. <https://doi.org/10.17705/1jais.00423>

28. Jegadeesh N, Wu D. Word Power: A New Approach For Content Analysis. *Journal of Financial Economics*. 2013; 110(3):712–729. <https://doi.org/10.1016/j.jfineco.2013.08.018>
29. Taddy M. Multinomial Inverse Regression for Text Analysis. *Journal of the American Statistical Association*. 2013; 108(503):755–770. <https://doi.org/10.1080/01621459.2012.734168>
30. Taddy M. Distributed Multinomial Regression. *The Annals of Applied Statistics*. 2015; 9(3):1394–1414. <https://doi.org/10.1214/15-AOAS831>
31. Joshi M, Das D, Gimpel K, Smith NA. Movie Reviews and Revenues: An Experiment in Text Regression. In: *The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*; 2010. p. 293–296.
32. Schwartz HA, Giorgi S, Sap M, Crutchley P, Ungar L, Eichstaedt J. DLATK: Differential Language Analysis ToolKit. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*; 2017. p. 55–60.
33. Tibshirani R. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society*. 1996; 58(1):267–288.
34. Hastie T, Tibshirani R, Friedman JH. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd ed. Springer Series in Statistics. New York: Springer; 2009.
35. Loughran T, McDonald B. Textual Analysis in Accounting and Finance: A Survey. *Journal of Accounting Research*. 2016; 54(4):1187–1230. <https://doi.org/10.1111/1475-679X.12123>
36. Varian HR. Big Data: New Tricks for Econometrics. *Journal of Economic Perspectives*. 2014; 28(2):3–27. <https://doi.org/10.1257/jep.28.2.3>
37. Hastie T, Tibshirani R, Wainwright M. *Statistical Learning with Sparsity: The Lasso and Generalizations*. vol. 143 of *Monographs on Statistics and Applied Probability*. Boca Raton, FL: CRC press; 2015.
38. Lockhart R, Taylor J, Tibshirani RJ, Tibshirani R. A Significance Test for the LASSO. *The Annals of Statistics*. 2014; 42(2):413–468. <https://doi.org/10.1214/13-AOS1175> PMID: 25574062
39. Belloni A, Chernozhukov V. Least Squares After Model Selection in High-Dimensional Sparse Models. *Bernoulli*. 2013; 19(2):521–547. <https://doi.org/10.3150/11-BEJ410>
40. Pang B, Lee L. Seeing Stars: Exploiting Class Relationships for Sentiment Categorization with Respect to Rating Scales. In: *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics (ACL'05)*; 2005. p. 115–124.
41. Carter ME, Soo BS. The Relevance of Form 8-K Reports. *Journal of Accounting Research*. 1999; 37(1):119–132. <https://doi.org/10.2307/2491399>
42. Loughran T, McDonald B. When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks. *Journal of Finance*. 2011; 66(1):35–65. <https://doi.org/10.1111/j.1540-6261.2010.01625.x>
43. Krippendorff K. *Content Analysis: An Introduction to Its Methodology*. 3rd ed. Thousand Oaks, CA: SAGE Publications; 2013.
44. Legg AM, Sweeny K. Do You Want the Good News or the Bad News First? The Nature and Consequences of News Order Preferences. *Personality and Social Psychology Bulletin*. 2014; 40(3):279–288. <https://doi.org/10.1177/0146167213509113> PMID: 24177520