

When Words Sweat: Identifying Signals for Loan Default in the Text of Loan Applications

Oded Netzer, Alain Lemaire, and Michal Herzenstein

Abstract

The authors present empirical evidence that borrowers, consciously or not, leave traces of their intentions, circumstances, and personality traits in the text they write when applying for a loan. This textual information has a substantial and significant ability to predict whether borrowers will pay back the loan above and beyond the financial and demographic variables commonly used in models predicting default. The authors use text-mining and machine learning tools to automatically process and analyze the raw text in over 120,000 loan requests from Prosper, an online crowdfunding platform. Including in the predictive model the textual information in the loan significantly helps predict loan default and can have substantial financial implications. The authors find that loan requests written by defaulting borrowers are more likely to include words related to their family, mentions of God, the borrower's financial and general hardship, pleading lenders for help, and short-term-focused words. The authors further observe that defaulting loan requests are written in a manner consistent with the writing styles of extroverts and liars.

Keywords

consumer finance, loan default, machine learning, text mining

Online supplement: <https://doi.org/10.1177/0022243719852959>

Imagine that you are considering lending \$2,000 to one of two borrowers on a crowdfunding website. The borrowers are identical in terms of their demographic and financial characteristics, the amount of money they wish to borrow, and the reason for borrowing the money. However, the text they provided when applying for a loan differs: Borrower #1 writes, "I am a hard working person, married for 25 years, and have two wonderful boys. Please let me explain why I need help. I would use the \$2,000 loan to fix our roof. Thank you, God bless you, and I promise to pay you back." Borrower #2 writes, "While the past year in our new place has been more than great, the roof is now leaking and I need to borrow \$2,000 to cover the cost of the repair. I pay all bills (e.g., car loans, cable, utilities) on time." Which borrower is more likely to default? This question is at the center of our research, as we investigate the power of words in predicting loan default. We claim and show that the text borrowers write at loan origination provides valuable information that cannot be otherwise extracted from the typical data lenders have on borrowers (which mostly include financial and demographic data), and that additional information is crucial to predictions of default. The idea that text written by borrowers can predict their loan default builds on recent research showing that text is indicative of people's psychological states, traits, opinions,

and situations (e.g., Humphreys and Jen-Hui Wang 2018; Matz and Netzer 2017).

In essence, the decision of whether to grant a loan depends on the lender's assessment of the borrower's ability to repay it. Yet this assessment is often difficult, because loans are repaid over a lengthy period of time, during which unforeseen circumstances may arise. For that reason, lenders (e.g., banks) and researchers collect and process as many pieces of information as possible within this tightly regulated industry.¹ These pieces of information can be classified into four categories: (1) the borrower's financial strength, as reflected by one's credit history, FICO score, income, and debt (Mayer, Pence, and Sherlund 2009), which is most telling of the borrower's ability to repay (Avery et al. 2000); (2) demographics, such as gender or

¹ Indeed, our personal discussions with data scientists at banks and other lending institutions suggest that they analyze as many as several thousand data points for each borrower.

Oded Netzer is Professor of Business, Columbia Business School, Columbia University, USA (email: onetzer@gsb.columbia.edu). Alain Lemaire is a doctoral student, Columbia Business School, Columbia University, USA (email: alemaire18@gsb.columbia.edu). Michal Herzenstein is Associate Professor of Marketing, Lerner College of Business and Economics, University of Delaware, USA (email: michalh@udel.edu).

geographic location (Rugh and Massey 2010)²; (3) information related to the loan (i.e., the requested borrowed amount and interest rate; Gross et al. 2009); and (4) everything else that can be learned from interactions between borrowers and people at loan-granting institutions. Indeed, Agarwal and Hauswald (2010) find that supplementing the loan application process with the human touch of loan officers significantly decreases default rate due to better screening and higher interpersonal commitment from borrowers. However, these interactions are often laborious and expensive.

Recently, human interactions between borrowers and lenders have been largely replaced by online lending platforms operated by banks, other lending institutions, or crowdfunding platforms. In such environments, the role of both hard and soft pieces of information becomes crucial. Accordingly, our main proposition is that the text people write when requesting an online crowdfunded loan provides additional important information on them, such as their intentions, personality, and circumstances—information that cannot be deduced from the financial and demographic data and is, in a sense, analogous to body (or unspoken) language detected by loan officers. Furthermore, because the evaluation of borrowers by loan officers has been shown to provide additional information to the financial information (Agarwal and Hauswald 2010), we contend that, in a similar vein, the text borrowers write when applying for a crowdfunded loan is predictive of loan default above and beyond all other available information. This hypothesis extends the idea that our demeanor can be a manifestation of our true intentions (DePaulo et al. 2003) in the text we write.

To answer the question we posed previously—who is more likely to default—we apply text-mining and machine learning tools to a data set of over 120,000 loan requests from the crowdfunding platform Prosper. Using an ensemble stacking approach that includes tree-based methods and regularized logistic regressions, we find that the textual information significantly improves predictions of default. A simple back-of-an-envelope analysis shows that using textual information can increase lenders' return on investment (ROI) over an approach that uses only financial and demographics information by as much as 5.75%.

To learn which words, writing styles, and general ideas conveyed by the text are more likely to be associated with defaulted loan requests, we further analyzed the data using a multimethod approach including both machine learning tools (e.g., naïve Bayes and L1 regularization binary logistic model) and standard econometric tools such as logistic regression of the topics extracted from a latent Dirichlet allocation (LDA) analysis, and the subdictionaries of the Linguistic Inquiry and Word Count dictionary (LIWC; Tausczik and Pennebaker 2010). Our results consistently show that loan requests written by defaulting borrowers are more likely to include words (or

themes) related to the borrower's family, financial and general hardship, mentions of God, mentions of the near future, pleading lenders for help, and using verbs in present and future tenses. Therefore, the text and writing style of borrower #1 in our opening example suggest this person is more likely to default. In fact, all else being equal, our analysis shows that based on the loan request text, borrower #1 is approximately eight times more likely to default relative to borrower #2. These analyses demonstrate the successful use of machine learning tools in going beyond merely predicting the outcome and into the realm of interpretation by inferring the words, topics, and writing styles that are most associated with a behavioral outcome.

Applying the LIWC dictionary to our data allows for a deeper exploration into the potential traits and states of borrowers. Our results suggest that defaulting loan requests are written in a manner consistent with the writing styles of extroverts and liars. We do not claim that defaulting borrowers were intentionally deceptive when they wrote the loan request; rather, we believe their writing style may have reflected (intentionally or not) doubts in their ability to repay the loan.

Our examination into the manifestation of consumers' personalities and states in the text they write during loan origination contributes to the fast-growing literature in consumer financial decision making. Consumer researchers have been investigating factors that affect consumer saving (Dholakia et al. 2016, Sussman and O'Brien 2016), debt acquisition, and repayment (Herzenstein, Sonenshein, and Dholakia 2011). Most of these investigations have been done on a smaller scale, such as with experimental participants or smaller data sets. We add to this literature by showing, on a large scale and with archival data, how participants in the crowdfunding industry can better assess the risk of default by interpreting and incorporating soft unverifiable data (the words borrowers write) into their analysis.

The rest of this article is organized as follows. In the next section, we discuss the limitations of structured financial and demographic data in the context of consumer financial decisions and the opportunity in leveraging textual information. Specifically, drawing on extant literature we argue that credit scores and demographics miss important information about borrowers that lenders can learn from mining the text these borrowers write. We then delineate the data, our text-mining and modeling approaches, the results, and their generalization.

What Do Hard Data Miss? The Opportunity in Text

Financial data such as credit scores have been used extensively to predict consumers' credit riskiness, which is the likelihood that they will become seriously delinquent on their credit obligations over the next few years. FICO, the dominant credit scoring model, is used by 90% of all financial institutions in the United States in their decision-making process (according to myFICO.com). In calculating consumers' scores, credit score agencies take into account a multitude of data points,

² Following the Equal Credit Opportunity Act (ECOA) and the Fair Housing Act (FHA), most demographic variables cannot be used directly in the decision to grant loans.

including past and current loans as well as credit cards (their allowance and utilization) and any delinquencies that were reported to the credit bureau by companies such as cable TV, cell phone providers, and so on. However, while the usage of credit scores and financial information clearly has its benefits in predicting consumer risk, these measures have been found to be insufficient and even biased. For example, Avery et al. (2000) argue that credit scores are snapshots of the past because they rely on available consumers' credit history and therefore miss important factors such as health status and length of employment, which are more forward looking in nature. The authors find that those with low credit scores often benefit most from supplementing credit scores with additional data. Agarwal, Skiba, and Tobacman (2009) provide further evidence for the deficiency of credit scores by showing that more flexible credit score systems, which weigh differently the various data sources that comprise the score for each person, are eight times better at predicting loan default than more rigid systems such as credit scores. In summary, credit scores might predict future financial behavior well for some but not others.

Analyzing the subprime mortgage default during 2006–2007, Palmer (2015) and Sengupta and Bhardwaj (2015) show that current contextual information such as loan characteristics (e.g., loan amount) and economic characteristics (e.g., housing prices) are often more predictive of mortgage default than traditional financial measures such as credit scores and debt-to-income ratios. Thus, it is clear that credit scores are missing important information that may be predictive of consumers' financial health and that the heavy reliance on such scores might blind lending organizations to additional sources of information.

Recognizing the limitation of financial measures, financial institutions and researchers added other variables to their predictive models, mostly related to the individuals' demographic, geographic, and psychographic characteristics (e.g., Barber and Odean 2001; Pompian and Longo 2004). However, other characteristics, such as those related to emotional and mental states, as well as personalities have been found to be closely tied to financial behaviors and outcomes (Norvilitis et al. 2006; Rustichini et al. 2016), yet are still missing from those predictive models. The problem with relying on purely structured financial information is even more severe in the growing realm of crowdfunded unsecured loans, because human interactions between lenders and borrowers is scarce—but therein lies the problem. Personalities and mental states are difficult, if not impossible, to infer from financial data alone, thus the pressing need to go beyond such traditional data when attempting to predict behavior. We suggest and demonstrate that the text borrowers write at loan origination can provide the much-needed supplemental information. Specifically, this textual information can be useful in understanding not only the past behavior of the borrower but also the present context of the loan and the borrower's future intentions.

Language Styles, Individual Characteristics, and Financial Behaviors

Our proposition that text can predict default builds on research in marketing, psychology, linguistics, and finance that establishes two connections: (1) word usage and writing styles are indicative of some stable inner traits as well as more transient states and (2) these traits and states affect people's financial behaviors. In this section, we provide evidence for these two connections.

We begin with the relationship between words and personality traits, emotional states, and demographics. The premise that the text may be indicative of deeper traits is predicated on the idea that there is a systematic relationship between the words people use and their personality traits (Hirsh and Peterson 2009), identities (McAdams 2001), and emotional states (Tausczik and Pennebaker 2010). The relationship between word usage and personality traits has been found across multiple textual media such as essays about the self (Hirsh and Peterson 2009), blogs (Yarkoni 2010), social media (Schwartz et al. 2013), and daily speeches (Mehl, Gosling, and Pennebaker 2006). This relationship stems from the human tendency to tell stories and express internal thoughts and emotions through these stories, which are essentially made possible by language. Therefore, even if the content might be similar across different individuals, the manner in which they convey that content differs.

Research over the last two decades has established the association between word usage with the text originator's personality, focusing on the "big five" personality traits: extroversion, agreeableness, conscientiousness, neuroticism, and openness (Kosinski, Stillwell, and Graepel 2013; Pennebaker and Graybeal 2001; Pennebaker and King 1999). Using a sample of over 1,000 students, Pennebaker and King (1999) collected scales for the big five personality traits, emotional states (positive and negative affect schedule), as well as two essays related to their current thoughts and feelings and thoughts about coming to college. Using LIWC and correlational analyses, the authors concluded that extroversion is associated with more positive-emotion word usage, neuroticism with more first-person singular and negative-emotion words, openness with more article usage, agreeableness with less article usage, and conscientiousness with fewer negations. These results were later corroborated and extended by other researchers (see, e.g., Yarkoni [2010] and Schwartz et al. [2013, Figure 2]). Many of the signals of personality in text come from writing style rather than content. Indeed, language style, more so than content, is at the heart of our research.

In addition to personality traits, words have been found to be associated with emotional and behavioral states (Pennebaker, Mayne, and Francis 1997), such as physical and mental health (Preotiuc-Pietro et al. 2015), impulsivity (Arya, Eckel, and Wichman 2013), and deception (Newman et al. 2003; Ott, Cardie, and Hancock 2012), among others. Text has also been used to predict demographic characteristics such as age, gender, education, and income (Pennebaker and Stone 2003;

Schwartz et al. 2013). For example, frequent use of long words (six letters or more) was found to be related to higher educational levels (Tausczik and Pennebaker 2010), and spelling mistakes in written text are correlated with income levels (Harkness 2016).

Next, we examine the relationship between personality traits, emotional states, and financial behavior. The aforementioned literature conveys the potential of text analysis to extract a wealth of information otherwise unobtainable in many settings. Such information has been shown to be a valuable predictor of financial behaviors. For example, Anderson et al. (2011) show that credit scores are correlated with the big five personality traits. They find negative correlations between extroversion and conscientiousness and credit scores. Extroverts want to have exciting life styles and thus sometimes spend beyond their means. The result for conscientiousness is a bit surprising, because many conscientious people are diligent and responsible; however, their need for achievement is high, which might induce a pattern of spending that is larger than their means. Berneth, Taylor, and Walker (2011) find that agreeableness is negatively correlated with credit scores because these people aim to please and are less likely to say no to unnecessary expenses. The extent to which such traits are predictive of financial behavior above and beyond credit scores is an empirical question, which we investigate in this research.

Other individual characteristics are also related to financial behavior. Arya, Eckel, and Wichman (2013) find that credit scores are correlated with personality, time and risk preference, trustworthiness, and impulsiveness; Nyhus and Webley (2001) show that emotional stability, autonomy, and extroversion are robust predictors of saving and borrowing behaviors; and Norvilitis et al. (2006) show that debt is related to delay of gratification. Financial behaviors such as saving, taking out loans, and credit card usage, as well as overall scores such as FICO, have also been found to be correlated with education (Nyhus and Webley 2001), financial literacy (Fernandes, Lynch, and Netemeyer 2014), number of hours people work (Berneth, Taylor, and Walker 2011), stress and physical well-being (Netemeyer et al. 2018), self-regulation (Freitas et al. 2002), and even the number of social media connections (Wei et al. 2015).

In summary, we postulate that many of the behavioral characteristics and situational factors that have been found to be related to financial behaviors, such as personality traits and future intentions, can be extracted from the text borrowers write in their loan request. Therefore, that text is an invaluable addition to the hard financial data when predicting loan default.

Settings and Data

We examine the value of text in predicting default using data from Prosper, the first online crowdfunding platform and currently the second largest in the United States, with over 2 million members and \$14 billion in funded unsecured loans. In Prosper, potential borrowers submit their request for a loan for a specific amount with a specific maximum interest rate they are willing to pay, and, during the span of our data set time

period, lenders then bid in a Dutch-like auction on the lender rate for loan. We downloaded all loan requests posted between April 2007 and October 2008, a total of 122,479 listings. In October 2008, the Securities and Exchange Commission required Prosper to register as a seller of investment, and when Prosper relaunched in July 2009, it made significant changes to the platform. We chose data from the earlier days because they are richer and more diverse, particularly with respect to the textual information in the loan requests.

When posting a loan request on Prosper, potential borrowers have to specify the loan amount they wish to borrow (between \$1,000 and \$25,000 in our data), the maximum interest rate they are willing to pay, and other personal information, such as their debt-to-income ratio and whether they are homeowners. Prosper verifies all financial information, including the potential borrower's credit score from Experian, and assigns each borrower a credit grade that reflects all of this information. The possible credit grades are AA (lowest risk for lenders), A, B, C, D, E, and HR (highest risk to lenders). Table A1 in the Web Appendix³ presents correspondence between Prosper's credit grades and FICO score.⁴ In addition, borrowers can upload as many pictures as they wish and can use an open text box to write any information they wish, with no length restriction. The words borrowers write in that text box are at the center of our research.

Borrowers could choose to become part of a group of borrowers. These groups often have a subject matter (groups can pertain to borrowers' alma maters, purpose for the loan, geographical regions, certain professions, etc.), and they must have a leader who supposedly acts as another layer of vetting. Group membership may affect the likelihood of a loan being granted and its default likelihood (Hildebrand, Puri, and Rocholl 2017). Accordingly, group affiliation is included in our models. After borrowers posted their listings, they may interact with lenders through a question-and-answer interface. Unfortunately, we do not have access to those chats, but unofficial reports suggest that use of this feature was limited (3% of the lenders participated).

Because our interest is in predicting default, we focus on those loan requests that were funded (19,446 requests). The default rate in our sample is 35%.

We automatically text-mined the raw text in each loan application using the *tm* package in R. Our textual unit is a loan application. For each loan application, we first tokenize each word, a process that breaks down each loan application into the distinct words it contains. We then use Porter's stemming algorithm to collapse variations of words into one. For example, "borrower," "borrowed," "borrowing," and "borrowers" become "borrow." In total, the loan requests in our data set

³ All tables and graphs whose numbering begins with "A" are presented in the Web Appendix.

⁴ The distribution of credit scores among Prosper borrowers in our sample is different from the distribution in the United States at the time (see Web Appendix Table A2): our sample has fewer people with high credit scores and more with medium scores.

have over 3.5 million words, corresponding to 30,920 unique words that are at least three letters long (we excluded numbers and symbols from our analysis).⁵ In addition to words/stems, we also examine two-word combinations (an approach often referred to as n-gram, in which for $n = 2$, we get bigrams).⁶ To reduce the dimensionality of the textual data and avoid over-relying on more obscure words, we focus our analyses on the most frequent stemmed words and bigrams that appeared in at least 400 loan requests, which left us with 1,052 bigrams.⁷

Textual, Financial, and Demographic Variables

Our dependent variable is loan repayment/default as reported by Prosper⁸ (binary: 1 = paid in full, 0 = defaulted). Our data horizon ends in 2008, and all Prosper loans at the time were to be repaid over three years or less; therefore, we know whether each loan in our database was fully repaid or defaulted—there are no other options. The set of independent variables we use includes textual, financial, and demographic variables. We elaborate on each group next.

Textual variables. These variables include the following:

1. *The number of characters in the title and the text box in the loan request.* The length of the text has been associated with deception; however, the evidence for this is inconclusive. Hancock et al. (2007) show that liars wrote much more when communicating through text messages than nonliars. Similarly, Ott, Cardie, and Hancock (2012) demonstrated that fake hospitality reviews are wordier though less descriptive. However, in the context of online dating websites, Toma and Hancock (2012) showed that shorter profiles indicate the person is lying, because they wished to avoid certain topics.
2. *The percentage of words with six or more letters.* This metric is commonly used to measure complex language, education level, and social status (Tausczik and Pennebaker 2010). More educated people are likely to have higher income and higher levels of financial literacy and thus are less likely to default on their loan relative to less educated people (Nyhus and Webley 2001). However, the use of complex language can also be risky if readers perceive it to be artificially or frivolously

complex, suggesting that the higher language was likely used deceptively (Oppenheimer 2006).

3. *The “simple measure of gobbledygook” (SMOG; McLaughlin, 1969).* This measures writing quality by mapping it to number of years of formal education needed to easily understand the text in first reading.
4. *A count of spelling mistakes.* We base this on the Enchant spellchecker using the PyEnchant 1.6.6. package in Python. Harkness (2016) shows that spelling mistakes are associated with a lower likelihood of granting a loan in traditional channels because they serve as a proxy for characteristics correlated with lower income.
5. *The 1,052 bigrams from the open text box in each loan application.* We identify these following the text mining process described previously.

Because loan requests differ in length, and words differ in the frequency of appearance in our corpus, we normalize the frequency of a word appearance in a loan request to its appearance in the corpus and the number of words in the loan request using the term frequency-inverse document frequency (tf-idf) measure commonly used in information retrieval. The term frequency for word m in loan request j is defined by $tf_{mj} = X_{mj}/N_j$, where X_{mj} is the number of times word m appears in loan request j and N_j is the number of words in loan request j . This component controls for the length of the document. The inverse document frequency is defined by $idf_m = \log(D/M_m)$, where D is the number of loan requests and M_m is the number of loan requests in which word m appears. This term controls for the how often a word appears across documents. The tf-idf is given by $tf-idf_{mj} = tf_{mj} \times (idf_m + 1)$. Taken together, the tf-idf statistic provides a measure of how likely a word is to appear in a document beyond chance.

Financial and demographic variables. The second type of variables we consider are financial and demographic information, commonly used in traditional risk models. These include all information available to lenders on Prosper—loan amount, borrower’s credit grade (modeled as a categorical variable AA-HR), debt-to-income ratio, whether the borrower is a homeowner, the bank fee for payment transfers, whether the loan is a relisting of a previous unsuccessful loan request, and whether the borrower included a picture with the loan. We also control for the borrower’s geographical location to account for differences in economic environments. We grouped the borrowers’ states of residency into eight categories using the Bureau of Economic Analysis classification (plus a group for military personnel serving overseas). Finally, to fully account for all the information in loan requests, we extracted information included in the borrower’s profile pictures, such as gender (male, female, and “cannot tell”), age brackets (young, middle-aged, and old), and race (Caucasian, African American, Asian, Hispanic, or “cannot tell”) using human coders. For details about the coding

⁵ Because of stemming, words with fewer than three letters (e.g., “I”) may be kept due to longer stems (e.g., “I’ve”).

⁶ While n-grams with $n > 2$ (e.g., strings of three or more words) could have been part of our analyses, this would have increased dimensionality and computational difficulty substantially, which ultimately precluded their inclusion.

⁷ We note that our analyses are robust to increasing the number of words and bigrams that are included.

⁸ We classified a loan as “defaulted” if the loan status in Prosper is “Charge-off,” “Defaulted (Bankruptcy),” or “Defaulted (Delinquency).” We classified a loan as “paid” if it is labeled “Paid in full,” “Settled in full,” or “Paid.”

Table 1. Descriptive Statistics for Funded Loan Requests ($n = 19,446$).

Variables	Min	Max	Mean	SD	Freq.
Amount requested	1,000	25,000	6,506.9	5,732.4	
Debt-to-income ratio	0	10.01	.33	.86	
Lender interest rate	0	.350	.180	.077	
Number of words in description	0	766	207.9	137.4	
Number of words in title	0	13	4.595	2.012	
% of long words (6+ letters)	0%	71.4%	29.8%	6.4%	
SMOG	3.129	12	11.347	1.045	
Enchant spellchecker	0	54	2.355	3.015	
Number of prior listings	0	67	2.016	3.097	
Credit grade					
AA					.131
A					.134
B					.174
C					.215
D					.182
E					.084
HR					.081
Loan repayment (1 = paid, 0 = defaulted)					.669
Loan image dummy					.666
Debt-to-income missing dummy					.050
Group membership dummy					.253
Homeowner dummy					.470

procedure, see the Web Appendix.⁹ We note that while the Equal Credit Opportunity Act forbids discrimination based on race, age, and gender, we include these variables in the statistical model to examine the marginal value of the textual information that does not relate to demographics.

Lastly, we include the final interest rate for each loan as a predictor in our model.¹⁰ Arguably, in a financially efficient world, this final interest rate, which was determined using a bidding process, should reflect all the information available to lenders, which means that including the final interest rate in the models should render other variables insignificant predictors. However, given the possibility that Prosper's bidding mechanism allows for some strategic behavior by sophisticated lenders (thus not fully reflecting a market efficient behavior), our models test whether the text is predictive above and beyond the final interest rate. Table 1 presents summary statistics for the variables in our model.

Predicting Default

Our objective in this section is to evaluate whether the text that borrowers write in their loan request is predictive of their loan default up to three years post origination. To do so, we need to first build a strong benchmark—a powerful predictive model that includes the financial and

demographics information and maximizes the chances of predicting default using these variables. Second, we need to account for the fact that our model may include a very large number of predictors (over 1,000 bigrams). Given the large number of predictors, and the predictive nature of the task at hand, machine learning methods are most appropriate. In the subsequent section, as we aim to understand which words predict default, we combine the machine learning methods with data reductions methods (e.g., topic modeling) and standard econometric tools.

In evaluating a predictive model, it is common to compare alternative predictive models and choose the model that best predicts the desired outcome—loan repayment, in our case. From a purely predictive point of view, a better approach, commonly used in machine learning, is to train several predictive models and, rather than choose the best model, create an ensemble or stack the different models. Recent research has demonstrated the superior performance of ensemble models relative to individual models (Lee, Hosanagar, and Nair 2018). An ensemble of models benefits from the strength of each individual model and, at the same time, reduces the variance of the prediction.

The stacking ensemble algorithm includes two steps. In the first step, we train each model on the calibration data. Because of the large number of textual variables in our model, we employ a simultaneous variable selection and model estimation in the first step. In the second step, we build a weighting model to optimally combine the models calibrated in the first step.

We estimate five types of models in the first step. The models vary in terms of the classifier used and the approach to model variable selection. The five models are described next

⁹ Our judges also rated borrowers' attractiveness and trustworthiness on the basis of their picture (Pope and Sydnor 2011), but given the high degree of disagreements across raters, they are not included in our analyses.

¹⁰ An alternative measure would be the maximum interest rate proposed by the borrower. However, because this measure is highly correlated with the final lender rate, we include only the latter in the model.

and include two logistic regressions and three versions of decision tree classifiers.¹¹

Regularized logistic regressions. We estimate two logistic regressions—L1 and L2 regularization logistic regressions. The penalized logistic regression log-likelihood is:

$$L(Y|\beta, \lambda) = \sum_{j=1}^n (y_j \log[p(X_j|\beta)] + (1 - y_j) \log[1 - p(X_j|\beta)] - \lambda J(\beta)),$$

where $Y = \{y_1, \dots, y_n\}$ is the set of binary outcome variables for n loans (loan repayment), $p(X_j|\beta)$ is the probability of repayment based on the logit model, where X_j is a vector of textual, financial, and demographic predictors for loan j , β are a set of predictors' coefficients, λ is a tuning penalization parameter to be estimated using cross-validation on the calibration sample, and $J(\beta)$ is the penalization term. The L1 and L2 models differ with respect to the functional form of the penalization term, $J(\beta)$. In L1, $J(\beta) = \sum_{i=1}^k |\beta_i|$, while in L2, $J(\beta) = \sum_{i=1}^k \beta_i^2$, where k is the number of predictors. Therefore, L1 is the Lasso regression penalty (shrinks many of the regression parameters to exactly zero), and L2 is the ridge regression penalty (shrinks many parameters to small but nonzero values). Before entering the variables into these regression, we standardize all variables (Tibshirani 1997).

Tree-based methods (random forest and extra trees). There are three tree-based methods in the ensemble. We estimate two different random forest models, one with variance selection and the second with best feature selection as well as extremely randomized trees (extra trees). Both models combine many decision trees; thus, each of these tree-based methods is an ensemble in and of itself. The random forest randomly draws with replacements subsets of the calibration data to fit each tree, and a random subset of features (variables) is used in each tree. In the variance selection, random forest features are chosen on the basis of a variance threshold determined by cross validation (80/20 split). In the K-best feature selection, random forest features are selected on the basis of a χ^2 test. That is, we select the K-features with the highest χ^2 score. We use cross-validation (80/20 split) to determine the value of K. The random forest approach mitigates the problem of overfitting in traditional decision trees. The extra trees is an extension of the random forest, in which the cutoff points (the split) for each feature in the tree are also chosen at random (from a uniform distribution) and the best split among them is chosen. Due to the size of the feature space, we first apply the aforementioned K-best feature selection to select the features to be included in the extra trees (for more details, see the Web Appendix).

We used the scikit-learn package in Python (<http://scikit-learn.org/>) to implement the five classifiers on a random sample of 80% of the calibration data. For the logistic regressions,

we estimated the λ penalization parameter by grid search using a three-fold cross-validation on the calibration sample. For the tree-based methods, to limit overfitting of the trees, we randomized the parameter optimization (Bergstra and Bengio 2012) using a three-fold cross-validation on the calibration data to determine the structure of the tree (number of leaves, number of splits, depth of the tree, and criteria). We use a randomized parameter optimization rather than an exhaustive search (or a grid search) due to the large number of variables in our model. The parameters are sampled from a distribution (uniform) over all possible parameter values.

Model stacking and predictions. In the second step, we estimate the weights for each model to combine the ensemble of models using the remaining 20% of the calibration data. We use a simple binary logistic model to combine the different predictive models. Though other classifiers may be used, a logistic binary regression meta-classifier helps avoid overfitting and often results in superior performance (Whalen and Gaurav 2013). In our binary logistic regression model, repayment is the dependent variable and the probabilities of repayment for each loan by each of the five models in the ensembles from step one (the two logistic regularization regressions and the three decision tree methods) as predictors. The estimated parameters of the logistic regression provide the weights of each individual model in the ensemble. Specifically, the ensemble repayment probability for loan j can be written as:

$$p(\text{repayment}_j) = \frac{\exp(x_j'w)}{[1 + \exp(x_j'w)]},$$

where x_j is the vector of repayment probabilities $p(\text{repayment}_j | \text{model}_s)$ for each model s —from step one—and w are the estimated weights of each model in the logistic regression classifier.

We estimated an ensemble of the aforementioned five models and find the following weights for the different model: L1 = .040, L2 = .560, random forest K-best = .218, random forest variance select = .116, and extra trees = .066. To test whether the text that borrowers wrote in their loan requests is predictive of future default, we use a tenfold cross-validation. We randomly split the loans into ten equally sized groups, calibrate the ensemble algorithm on nine groups, and predict the remaining group. To evaluate statistical significance, we repeated the tenfold cross-validation ten times, using different random seeds at each iteration. By cycling through the ten groups and averaging the prediction results across the ten cycles and ten replications, we get a robust measure of 100 predictions. Because there is no obvious cutoff for a probability from which one should consider the loan as defaulted, we use the area under the curve (AUC) of the receiver operating characteristic (ROC) curve, a commonly used measure for prediction accuracy of binary outcomes. We further report the Jaccard index (e.g., Netzer et al. 2012) of loan default, which is defined as the number of correctly predicted defaulting loans divided by the total number of loans that were defaulted but were incorrectly predicted to be

¹¹ We also considered another decision tree (AdaBoost) as well as a support vector machine classifier but dropped them due to poor performance on our data.

Table 2. AUC for Models with Text Only, Financial and Demographics Information Only, and a Combination of Both.

	Model 1 Text Only	Model 2 Financial and Demo. Only	Model 3 Text, Financial, and Demo.	Improvement from Model 2 to Model 3
Overall AUC	66.69%	70.72%	72.60%	2.64%*
Low credit grades: E, HR	61.78%	62.62%	65.61%	4.77%*
Medium credit grades: B, C, D	62.84%	65.75%	68.06%	3.51%*
High credit grades: AA, A	72.34%	77.38%	78.88%	1.94%*
Jaccard index	37.97%	37.01%	39.50%	3.92%*
AUC of the Underlying Models of the Ensemble				
Logistic L1	67.75%	70.09%	71.66%	2.23%*
Logistic L2	68.09%	68.58%	72.09%	5.10%*
Random forest (variance selection)	64.62%	70.35%	70.85%	0.71%*
Random forest (best features selection)	66.14%	69.24%	71.13%	2.73%*
Extremely randomized trees (extra trees)	66.65%	69.81%	70.98%	1.66%*

*Significant improvements at $p < .01$.

Notes: All AUCs are averaged across ten replications of the tenfold mean. The plot of the ROC curve is for the average across randomly selected tenfold models, see Figure 1. The Jaccard index is calculated as $N00/(N01 + N10 + N00)$, where N00 is the number of correctly predicted defaults and N01 and N10 are the numbers of mispredicted repayments and defaults, respectively.

repaid, loans that were predicted to default but were repaid, and correctly predicted defaulted loans. This gives us an intuitive measure of hit rates of defaulting loans penalized for erroneous predictions of both type I and type II errors. Finally, building on research that shows credit scores have a lower predictive power of financial behavior for people with low scores (Avery et al. 2000), we report predictions for high (AA, A), medium (B, C), and low (D, E, HR) credit grades (while controlling for credit grades within each group).

We compare three versions of the ensemble: (1) a model calibrated only on the financial and demographic data, (2) a model that includes just the textual information (i.e., all the variables we created from the freely written text borrowers constructed) and ignores the financial and demographic information, and (3) a model that includes financial and demographic information together with the textual data. Comparing models 1 and 3 provides the incremental predictive power of the textual information over predictors commonly used in the financial industry. Comparing models 1 and 2 informs the degree of predictive information contained in the textual information relative to the financial and demographic information. In addition, we estimated separately the five predictive models that contribute to the ensemble (L1 and L2 regularization logistic regressions, the two random forest models, and the extra trees model) to assess the value of the textual information in different models.

Prediction Results

Table 2 details the average results of the tenfold cross-validation across ten random shufflings of the observations. The results we present constitute clean out-of-sample validation because in each fold we calibrate feature selection, model estimates, and the ensemble weights on 90% of the data and leave the remaining 10% of the data for validation. The results in Table 2 are the area under the ROC curve (or AUC) and the Jaccard index prediction measures. Figure 1 depicts the

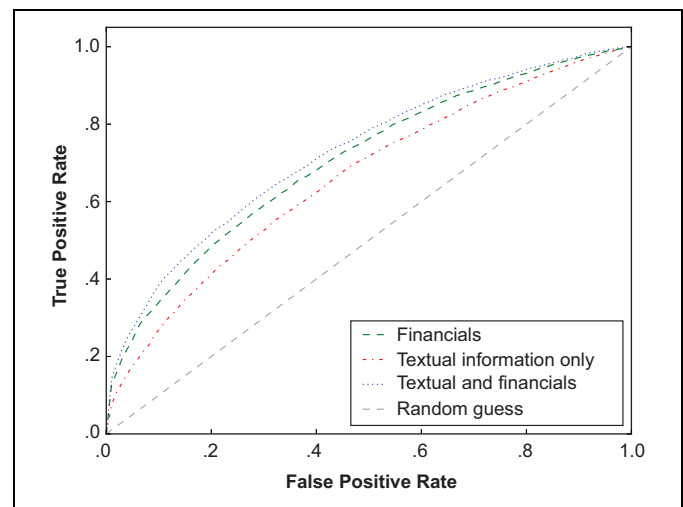


Figure 1. ROC curves for models with text only, financial and demographics information only, and a combination of both.

average ROC curve with and without the textual data for one randomly chosen tenfold cross-validation. As the ROC curve approaches the upper-left corner of the graph, the underlying model is better at predicting the outcome. The AUC of the model with textual, financial, and demographic information is 2.64% better than the AUC of the model with only financial and demographic information, and this difference is statically significant. In fact, the model with both textual and financial information has higher AUC in all 100 replications of the cross-validation exercise. Separating the sample by credit grade, we note that the textual information significantly improves predictions across all credit grade levels. However, the textual information is particularly useful in improving default predictions for borrowers with low credit levels. This result is consistent with Avery et al. (2000), who find credit scores to be least telling of people's true financial situation for those with low scores.

Interestingly, if we were to ignore the financial and demographic information and use only the borrower textual information, we obtain an AUC of 66.69% compared with an AUC of 70.72% for the model with only financial and demographic information. That is, brief, unverifiable “cheap talk” (Farrell and Rabin 1996) textual information provided by borrowers is nearly as predictive of default as the traditional financial and demographic information. This result is particularly impressive given the tremendous effort and expenditure involved in collecting the financial information relative to the simple method used to collect the textual information. This result may also suggest that textual information may be particularly useful in “thin file” situations, where the financial information about consumers is sparse (Coyle 2017).

The bottom part of Table 2 presents the predictive performance of each of the individual models in the ensemble. We observe that, for each of the models, the textual information significantly improves predictions in the validation sample above and beyond the model with the financial and demographic information only. However, the stacking ensemble model further improves predictions over each of the independent models. There are two key takeaways from this comparison. First, the textual information itself, independent of the machine learning model used, significantly contributes in predicting default over the traditional financial measures (Banko and Brill 2001). Second, combining models using an ensemble learning model further helps predict default. The reason for the improvement of the ensemble learning model relative to the individual model is that different models perform better in different aspects of the data. To better understand the performance of the ensemble model relative to the individual models, we compared the performance of the individual models by credit score, or the frequency of the words across loans (see Web Appendix Table A3). We find, for example, that the random forest variance and extra trees perform particularly well for low-credit-score loans, while the regularized logistic models perform well for the high-credit-score loans. Similarly, we find that the random forest model performs best for words that are less frequent across loans, but the regularized logistic models perform best for more frequent words.

To quantify the managerial relevance and financial implications of the improvement in predictive ability offered by the textual data, we conducted a back-of-the-envelope calculation. For each of the 19,446 granted loans we calculated the expected profit from investing \$1,000 in each loan based on the models with and without text. In calculating the expected profit, we assume that borrowers who default repay on average 25% of the loan before defaulting (based on estimates published by crowdfunding consulting agencies). The expected profits for loan j is

$$E(\text{profit}_j) = [1 - \text{Prob}(\text{repayment}_j)] \times [-.75 \times \text{amount_granted}_j + .25 \times \text{Interest_earned}_j] + [\text{Prob}(\text{repayment}_j)] \times \text{Interest_earned}_j,$$

(1)

where $\text{Prob}(\text{Repayment}_j)$ is the probability of repayment of loan j based on the corresponding model (with or without the text), amount_granted_j is the principal amount the lender grants for loan j (\$1,000 in our case), and Interest_earned_j is the interest rate paid to the lender based on loan j th final interest rate over three years for repaid loans and over three quarters of a year for defaulted loans (for simplicity we do not time-discount payments in years 2 and 3). For each of the two policies (based on the models with and without text) we sort the loans on the basis of their expected profit and select the top 1,000 loans with the highest expected return for each policy. Finally, we calculate the *actual* profitability of each lending policy based on the actual default of each loan in the data to calculate the return on the investment of the million dollars (1,000 loans time \$1,000 per loan). We find that the investment policy based on the model with the textual data returns \$57,571 more than the policy based on the financial and demographic information only. This is an increase of 5.75% in the ROI on the million dollars invested. Thus, while the improvement in default prediction for the model with the textual information might seem modest (nearly 3%) even though it is statistically significant, the improvement in ROI based on the textual information is substantial and economically meaningful.

We note that because we know only the repayment outcome for funded loans and because our model was trained to predict default only on the sample of funded loans, our back-of-the-envelope calculation can only assess the benefit of funding loans for loans that were funded. However, our model and data can provide some predictions with respect to the default likelihood of rejected (unfunded) loans. Indeed, when we compare the default likelihood distribution of rejected versus funded loans, we see that the default distribution, based on our model, is much higher for rejected loans relative to funded loans, but there is also a substantial overlap (see Web Appendix Figure A1). We also calculated a confusion matrix (see Web Appendix Table A4) based on the predicated default and expected profitability (following Equation 1) for loans that were actually granted versus not granted and whether these loans should have been granted (generate positive return). Of the 19,446 funded loan requests, our model recommends funding 60% (11,795 loans). In addition, of the 103,033 unfunded loan requests, our model recommends funding 21% (21,631 loans). Overall, based on our model we recommend granting 33,426 loans that are predicted to generate positive profits.

To summarize, the text borrowers write in their loan request can significantly improve predictions of loan default above and beyond all other available information, including the loan’s interest rate. We chose the ensemble-based predictive model to maximize predictive ability, but this model provides little to no interpretation of the parameter estimates, words, and topics that predict default. Next, we demonstrate how machine learning approaches combined with econometric models can be used beyond predictions and toward understanding which words and writing styles are most likely to appear in defaulting loans.

Words, Topics, and Writing Styles That Are Associated with Default

The result that text has a predictive ability similar in magnitude to the predictive ability of all other information is perhaps surprising, given that borrowers can write whatever they want. However, this result is consistent with the idea that people who differ in the way they think and feel also differ in what they say and write about those thoughts and feelings (Hirsh and Peterson 2009). We employed four approaches to uncover whether words, topics, and writing styles of defaulters differ from those who repaid their loan (based on the sample of 19,446 funded loans).

1. We use a naive Bayes classifier to identify the words or bigrams that most distinguish defaulted from fully paid loans. The advantage of the naive Bayes is in providing intuitive interpretation of the words that are most discriminative between defaulted and repaid loans. However, its disadvantage is that it assumes independence across predictors and thus cannot control for other variables.
2. To alleviate this concern, we use a logistic regression with L1 penalization, which reduces the dimensionality of the word space by setting some of the parameters to zero, to uncover the words and bigrams that are associated with default after controlling for the financial and demographic information. The L1 regression results corroborate the naive Bayes findings (correlation between the two analyses is .582, $p < .01$; see details in Web Appendix Table A5).
3. To look beyond specific bigrams and into the topics discussed in each loan, we use LDA analysis.
4. Finally, relying on a well-known dictionary, the LIWC (Tausczik and Pennebaker 2010), we identify the writing styles that are most correlated with defaulting or repaying the loan.

Words That Distinguish Between Loan Requests of Paying and Defaulting Borrowers

To investigate which words most discriminate between loan requests by borrowers who default versus repay the loan in full, we ran a multinomial naive Bayes classifier using the Python scikit-learn 3.0 package on bigrams (all possible words and bigrams) that appeared in at least 400 loans (1,052 bigrams). The classifier uses Bayes rule and the assumption of independence among words to estimate each word's likelihood of appearing in defaulted and paid loans. We then calculate the most "informative" bigrams in terms of discriminating between defaulted and repaid loans by calculating the bigrams with the highest ratio of $P(\text{bigram}|\text{defaulted})/P(\text{bigram}|\text{repaid})$ and the highest ratio of $P(\text{bigram}|\text{repaid})/P(\text{bigram}|\text{defaulted})$. Figures 2 and 3 present word clouds of the naive Bayes analysis of bigrams in their stemmed form (see underlying data in Web Appendix Table A6, Panels A and B). The size of each bigram corresponds to the likelihood that it will be included in a repaid loan request versus a defaulted loan request (Figures 2) or in a

defaulted loan request versus a repaid loan request (Figure 3). For example, the word "reinvest" in Figure 2 is 4.8 times more likely to appear in a fully paid than a defaulted loan request, whereas the word "God" in Figure 3 is 2.0 times more likely to appear in a defaulted than a fully paid loan request. The central cloud in each figure presents the most discriminant bigrams (cutoff ratio = 1.5), and the satellite clouds represent emerging themes based on our grouping of these words.

We find that, relative to defaulters, borrowers who paid in full were more likely to include the following in their loan application: (1) words associated with their financial situation, such as "reinvest," "interest," and "tax"; (2) words that may be a sign of projected improvement in financial ability, such as "graduate," "wedding," and "promote"; (3) relative words, such as "side," "rather," and "more than"; (4) long-term-time-related words, such as "future," "every month," and "few years"; and (5) "I" words, such as "I'd," "I'll," and "I'm." This indicates that borrowers who paid in full may have nothing to hide, have a brighter future ahead of them, and are generally truthful. The latter insight is based on research showing that the use of relative and time words as well as first-person "I" words is associated with greater candor because honest stories are usually more complex and personal (Newman et al. 2003). Dishonest stories, in contrast, are simpler, allowing the lying storyteller to conserve cognitive resources to focus on the lie more easily (Tausczik and Pennebaker 2010). Borrowers who repaid their loan also used words that indicate their financial literacy (e.g., "reinvest," "after tax," "minimum payment"). Indeed, higher financial literacy has been associated with lower debt (Brown et al. 2015).

As Figure 3 shows, borrowers who defaulted were more likely to mention the following:

1. Words related to financial hardships ("payday loan," "child support," and "refinance") and general hardship ("stress," "divorce," and "very hard"). This result is in line with Herzenstein, Sonenshein, and Dholakia (2011), who find that discussing personal hardship in the loan application is associated with borrowers who are late on their loan payments.
2. Words that explain their situation (e.g., "loan explain," "explain why") and discuss their work state (e.g., "hard work," "worker"). Providing explanations is often connected to past deviant behavior (Sonenshein, Herzenstein, and Dholakia 2011).
3. Appreciative and good-manner words toward lenders (e.g., "God bless," "hello") and pleading lenders for help (e.g., "need help," "please help"). Why is polite language more likely to appear in defaulted loan requests? One possibility is that it is not authentic. Indeed, Feldman et al. (2017) show that rude people are more trustworthy because their reactions seem more authentic.
4. Words referring to others (e.g., "god," "son," "someone"). The strong reference to others has been shown to exist in deceptive language style. Liars tend to avoid mentioning themselves, perhaps to distance themselves from the lie (Hancock et al. 2007; Newman




Figure 3. Words indicative of loan default.

Notes: The most common words appear in the middle cloud (cutoff = 1:1.5) and are then organized by themes. Starting on the top and moving clockwise: words related to explanations, external influence words and others, future-tense words, time-related words, work-related words, extremity words, words appealing to lenders, words relating to financial hardship, words relating to general hardship, and desperation/plea words.

predictive of default are the financial variables such as lender rate and credit score as well as words such as “payday loan,” “invest,” “hard,” “thank you,” “explain,” and “student.” This is an interesting result because it implies that talking about “payday loans” is not fully redundant with the information provided by one’s credit grade.

Relationship Between Words Associated with Loan Default and Words Associated with Loan Funding

It is possible that some borrowers strategically use words to convince lenders to fund their loans; therefore, we examine whether lenders are aware of such strategic behavior, and if not, which words slipped past lenders’ defenses and got them to fund overly risky loans that eventually defaulted. To investigate the relationship between words associated with loan default and words related to loan funding, we ran a naive Bayes analysis on the entire set of loans requests (122,479 funded and unfunded loan requests that included text) and assessed the bigrams with the highest ratio of $P(\text{bigram}|\text{funded})/P(\text{bigram}|\text{unfunded})$ and the highest ratio of $P(\text{bigram}|\text{unfunded})/P(\text{bigram}|\text{funded})$. For summary statistics of this data set, see Web Appendix Table A8; for the naive Bayes analysis results, see Web Appendix Table A9.

For each bigram, Figure 4 depicts its value on the ratio $P(\text{bigram}|\text{defaulted})/P(\text{bigram}|\text{repaid})$ versus its value on the ratio $P(\text{bigram}|\text{unfunded})/P(\text{bigram}|\text{funded})$. We named a few representative bigrams in the figure. A high correlation between the two ratios $P(\text{bigram}|\text{defaulted})/P(\text{bigram}|\text{repaid})$ and $P(\text{bigram}|\text{unfunded})/P(\text{bigram}|\text{funded})$ means that lenders are largely aware of the words that are associated with default. Results show a fairly strong correlation between the two ratios ($r = .354, p < .01$), suggesting that lenders are at least somewhat rational when interpreting and incorporating the text in their funding decisions. Examining the results more carefully, we find roughly three types of words: (1) words that have high likelihood of default and low likelihood of funding (e.g., “need help,” and “lost”) or low likelihood of default and high likelihood of funding (e.g., “excellent credit,” and “never miss”); (2) words that have a stronger impact on loan funding than on loan default (e.g., words related to explanation, such as “loan explain,” “situation explain,” “explain what,” and “explain why”); and (3) words that “tricked” lenders to fund loans that were eventually defaulted (words that were related to default more than to loan funding). The most typical words in this category are “god” and “god bless” as well as financial hardship words such as “payday” and “payday loan.”




Figure 4. Naive Bayes analysis for loan funding and loan default.

Table 3. The Seven LDA Topics and Representative Words with Highest Relevance.

LDA Topic	Words with highest relevance ($\lambda = .5$)
Employment and school	work, job, full, school, year, college, income, employ, student
Interest rate reduction	debt, interest, rate, high, consolidate, score, improve, lower
Expenses explanation	expense, explain, cloth, entertainment, cable, why, utility, insurance, monthly
Business and real estate	business, purchase, company, invest, fund, addition, property, market, build, cost, sell
Family	bill, try, family, life, husband, medical, reality, care, give, children, hard, daughter, chance, son, money, divorce
Loan details and explanations	loan, because, candidate, situation, financial, purpose, house, expense, monthly, income
Monthly payment	month, payment, paid, total, account, rent, mortgage, save, list, every, payday, budget

Notes: We chose the sample words on the basis of the relevance measure with $\lambda = .5$. For more exhaustive lists of words for each topic, see Web Appendix Table A10.

Analyzing the Topics Discussed in Each Loan Request and Their Relationship to Default

In Figures 2 and 3, we grouped the bigrams into topics on the basis of our own interpretation and judgment. However, several machine learning methods have been proposed to statistically combine words into topics based on their common cooccurrence in documents. Probably the most commonly used topic modeling approach is LDA analysis, which we apply to the complete data set of all loan requests (the 122,479 funded and rejected loans).

We use the online variational inference algorithm for the LDA training (Hoffman, Bach, and Blei 2010), following Griffiths and Steyvers (2004)'s settings and priors. For the LDA analysis, we used the 5,000 word stems that appeared most frequently across loan requests. Eliminating infrequent words mitigates the risk of rare-words occurrences and cooccurrence confounding the topics. Because the LDA analysis requires the researcher to determine the number of topics to be analyzed, we varied the number of topics between 2 and 30 and used model fit (the perplexity measure) to determine the final number of topics. We find that the model with seven topics had the best fit

(lowest perplexity). Table 3 presents the seven topics and the most representative words for each topic based on the relevance score of .5 (Sievert and Shirley 2014). Web Appendix Table A10 lists the top 30 words for each topic, and Web Appendix Figure A2 presents the perplexity analysis.

The topics we identify relate to the reason for the loan request, life circumstances, or writing style. We find three loan purpose topics: (1) employment and school, (2) interest rate consolidation, and (3) business and real estate. The other four topics are related to life circumstances and writing style: (4) expenses explanation, (5) family, (6) loan details explanation, and (7) monthly expenses. The monthly expenses topics are most likely related to a set of expenses Prosper recommended borrowers mention as part of their loan request during our data period.

Another advantage of the LDA is that it serves as a data-reduction technique, allowing us to use standard econometric methods (binary logistic regression) to relate text (topics) to default probabilities. Specifically, we ran a binary logistic regression with loan repayment = 1 and default = 0 as the dependent variable, the LDA probability of each topic

Table 4. Binary Regression with the Seven LDA Topics (Repayment = 1).

Financial and Loan-Related Variables	Estimate (SE)	Textual Variables	Estimate (SE)
Amount requested (in \$10 ⁵)	−7.08 (.35)*	Number of words in Description (in 10 ⁴)	−.00 (.001)
Credit grade: HR	−.79 (.08)*	Number of spelling mistakes	.00 (.00)
Credit grade: E	−.43 (.08)*	SMOG (in 10 ³)	−1.80 (1.60)
Credit grade: D	−.33 (.06)*	Words with six letters or more	−.72 (.37)*
Credit grade: C	−.17 (.05)*	Number of words in the title (in 10 ³)	−7.70 (8.00)
Credit grade: A	.76 (.08)*	Employment and school	2.26 (.43)*
Credit grade: AA	.24 (.07)*	Interest rate reduction	2.87 (.43)*
Debt to income	−.09 (.02)*	Expenses explanation	−.25 (.60)
Images	.04 (.04)	Business and real estate loan	.64 (.39)
Homeowner status	−.36 (.04)*	Family	−1.16 (.43)*
Lender interest rate	−5.57 (.30)*	Monthly payments	.99 (.41)
Bank draft fee annual rate	−39.12 (18.50)*		
Debt to income missing	−.25 (.07)*		
Group membership	−.20 (.04)*		
Prior listings	−.03 (.01)*	Intercept	2.13 (.35)*

Notes: Asterisk (*) indicates values significant at $p \leq .05$. For brevity, we do not report estimates for location, age, gender, or race. The topic loan detail explanation serves as baseline.

appearing in the loan (the topic “loan details explanation” serves as benchmark), and the same set of textual information metrics and financial and demographic variables used in the ensemble learning model described previously. Table 4 presents the results of a binary logistic regression with LDA topics and controls. We first observe that the financial and demographic variables are significant and in the expected direction: repayment likelihood increases as credit grades improve but decreases with debt-to-income ratio, home ownership, loan amount, and lender rate. The strong relationship between lender rate and repayment suggests some level of efficiency among Prosper lenders.

Relative to the “loan detail explanation” topic, we find that the topics of “employment and school,” “interest rate reduction,” and “monthly payment” are more likely to appear in repaid loan requests. These results corroborate the naive Bayes results that discussion of education is associated with lower default likelihood. Indeed, it is possible that traditional credit scores measures do not appropriately account for the positive effect of education on financial stability. The textual information provides lenders an indirect window into borrowers’ educations. The “family” topic, in contrast, was less likely to appear in repaid loans. Consistent with the naive Bayes analysis, we find that the topics of explaining one’s financials and loan motives and referring to family are associated with lower repayment likelihood. We also find that tendency to detail monthly expenses and financials is associated with higher likelihood of repayment, perhaps because providing such information is indeed truthful and forthcoming. Finally, although not the purpose of the LDA analysis, we find that the binary logistic model that includes the seven LDA topics probabilities in addition to the financial information fits the data and predicts default better than a model that does not include the LDA results, though it was worse than the ensemble model (for details, see the Web Appendix).

In summary, multiple methods converge to uncover themes and words that differentiate defaulted from paid loan requests. Next, we explore whether borrowers’ writing styles can shed light on the traits and states of defaulted borrowers.

Circumstances and Personalities of Those Who Defaulted

In this subsection, we rely on one of the more researched and established text analysis methods, the LIWC dictionary. This dictionary groups almost 4,500 words into 64 linguistic and psychologically meaningful categories such as tenses (past, present, and future); forms (I, we, you, she, or he); and social, positive, and negative emotions. Since its release in 2001, numerous researchers have examined and employed it in their research (for a comprehensive overview, see Tausczik and Pennebaker [2010]). The result of almost two decades of research is lists of word categories that represent the writing style of people with different personalities (Kosinski, Stillwell, and Graepel 2013; Pennebaker and King 1999; Schwartz et al. 2013; Yarkoni 2010), mental health states (Preotiuc-Pietro et al. 2015), emotional states (Pennebaker, Mayne, and Francis 1997), and many other traits and states.

The LIWC is composed of subdictionaries that can overlap (i.e., the same word can appear in several subdictionaries). We first calculate the proportion of stemmed words in each loan request that belong to each of the 64 dictionaries.¹³ We then estimated a binary logit model (loan repaid = 1 and loan defaulted = 0) to relate the proportions of words in each loan that appear in each dictionary to whether the loan was repaid, controlling for all financial and demographic variables used in our previous analyses. Results are presented in Table 5.

¹³ For this analysis, we did not remove words with fewer than three characters and infrequent words, as we are matching words to predefined dictionaries.

Table 5. Binary Regression with LIWC (Repayment = 1).

Financial and Basic Text Variables		LIWC Dictionaries			
	Amount requested ($\times 10^5$)	Swear words	Past words	Person pronoun words	
Credit grade: HR	-.6776 (.352)*	Filler words	Inhibition words	Work words	-.1517 (6.466)
Credit grade: E	-.8030 (.082)*	Perception words	Home words	Sexual words	.5035 (.898)
Credit grade: D	-.4177 (.080)*	Relative words	Hear words	They words	-10.6703 (10.580)
Credit grade: C	-.3088 (.061)*	Friend words	I words	Positive emotion words	-14.4516 (9.076)
Credit grade: A	-.1588 (.054)*	Anxiety words	Tentative words	Money words	.4510 (1.958)
Credit grade: AA	.7489 (.076)*	Negate words	Nonfluency words	Ingest words	.4857 (.720)
Debt to income	.2679 (.066)*	Insight words	Anger words	Verbs words	-3.8567 (5.029)
Debt to income missing	-.0843 (.010)*	"We" words	Achieve words	Adverbs words	.1246 (1.260)
Images	-.2391 (.074)*	Pronoun words	Inclusion words	Functional words	-3.043 (1.823)
Homeowner status	.0468 (.038)	Exclusion words	She/he words	Bios words	-1.1873 (1.791)
Lender interest rate	-.3493 (.037)*	Sad words	You words	Assent words	-1.6958 (2.576)
Bank draft fee annual rate	-5.618 (.301)*	Quantitative words	Cause words	Family words	-5.8254 (13.831)
Prior listings	-40.4880 (18.623)*	Articles	Social words	I pronoun words	-1.8106 (2.559)
Group membership	.0008 (.012)	Numbers words	Health words	Death words	-1.5238 (9.783)
Number of words in description	-.1627 (.042)*	Preposition words	Certain words	Body words	-18.7405 (10.460)
Number of spelling mistakes	-.0015 (.001)	Conjoint words	Present words	Religion words	-16.4249 (5.637)*
SMOG	-.0131 (.006)*	Auxiliary verbs words	Human words	Feel words	-19.8774 (6.590)*
Words with six letters or more	.0343 (.018)	Affect words	Space words	See words	-27.2482 (11.078)*
Number of words in the title	-.6324 (.508)	Discrepancy words	Future words	Leisure words	-12.2565 (10.940)
Intercept	-.0070 (.008)	Cognitive mechanism words	Motion words		-1.1835 (2.420)
	3.0565 (.435)*	Negative emotion words	Time words		

Notes: Asterisk (*) indicates values significant at $p \leq .05$. Standard errors are in parentheses. For brevity, we do not report estimates for location, age, gender, or race.

We begin by noting that all financial and demographic control variables are in the expected direction and are consistent with our LDA results. Fourteen LIWC dictionaries were significantly related to repayment behavior, and several of them mirror the naive Bayes and LDA results. To interpret our findings, we rely on previous research that leveraged the LIWC dictionary, which allows us to conclude that defaulted loan requests contain words that are associated with the writing styles of liars and of extroverts.

We begin with deception. From Table 5, we see that the following LIWC subdictionaries that have been shown to be associated with greater likelihood of deception are associated in our analysis with greater likelihood to default:

1. Present- and future-tense words. This result is similar to our naive Bayes findings. Indeed, prior research has shown that liars are more likely to use present- and future-tense words because they represent unconcluded situations (Pasupathi 2007).

2. Higher use of motion words (e.g., “drive,” “go,” “run”), which have been associated with lower cognitive complexity (Newman et al. 2003), and lower use of relative words (e.g., “closer,” “higher,” “older”), which are associated with greater complexity (Pennebaker and King 1999). Deceptive language has been shown to include more motion words and fewer relative words because it is less complex in nature.

3. Similar to our finding from the naive Bayes analysis that defaulters tend to refer to others and repayers tend to use more “I” words, we find that social words (e.g., “mother,” “father,” “he,” “she,” “we,” “they”) are associated with higher likelihood of default. Along these lines, Hancock et al. (2007) shows that the linguistic writing style of liars is reflected by lower use of first-person singular (“I”) and higher use of first-person plural (“we”) (see also Bond and Lee 2005; Newman et al. 2003). We note that in the context of hotel reviews, Ott, Cardie, and Hancock (2012) find higher use of “I” in fake reviews possibly because the posters did not have much to write about the hotel itself (because they have never been there), so they described their own activities.

4. Time words (e.g., “January,” “Sunday,” “morning,” “never”) and space words (e.g., “above,” “inch,” “north”) were associated with higher likelihood of default. These words have been found to be prevalent in deceptive statements written by prisoners because the use of such words draws attention away from the self (Bond and Lee 2005).

Taken together, we find that several of the LIWC dictionaries that have been previously found to be associated with deception are also negatively associated with loan repayment (positively associated with loan default). But do borrowers intentionally lie to lenders? One possibility is that people can predict, with some accuracy, future default months or even years ahead of time. If so, this would support the idea that defaulters are writing either intentionally or unconsciously to deceive lenders. However, there is another option: borrowers on Prosper might genuinely believe that they will be able to pay back the borrowed money in full. Indeed, previous research has suggested that people are often optimistic about future

outcomes (Weinstein 1980). What they may be hiding from lenders is the extent of their difficult situations and circumstances.

Our second observation is that the subdictionaries associated with the writing style of extroverts are also associated with greater likelihood of default. Extroverts have been shown to use more religious and body-related words (e.g., “mouth,” “rib,” “sweat,” “naked”; Yarkoni 2010), more social and human words (e.g., “adults,” “boy,” “female”; Hirsh and Peterson 2009; Pennebaker and King 1999; Schwartz et al. 2013; Yarkoni 2010), more motion words (e.g., “drive,” “go,” “run”; Schwartz et al. 2013), more achievement words (e.g., “able,” “accomplish,” “master”), and fewer filler words (e.g., “blah,” “like”; Mairesse et al. 2007)—all of which are significantly related to a greater likelihood of default in our analysis (see Table 5).

The finding that defaulters are more likely to exhibit the writing style of extroverts is consistent with research showing that extroverts are more likely to take risks (Nicholson et al. 2005), more likely to compulsively buy lottery tickets (Balanis 2002), and less likely to save (Nyhus and Webley 2001). Moreover, it is not a coincidence that the 14 LIWC dictionaries that were significantly correlated with default are correlated with both extroversion and deception. Previous literature has consistently documented that extroverts are more likely to lie not only because they talk to more people but also because these lies help smooth their interactions with others (Weiss and Feldman 2006).

We did not find a consistent and conclusive relationship between the LIWC dictionaries associated with the other big five personality traits and loan repayment. Similarly, results from other research on the relationship between LIWC and gender, age, and mental and emotional states did not consistently relate to default in our study. Finally, we acknowledge that there may be variables that are confounded with both the observable text and unobservable personality traits or states that are accountable for the repayment behavior. Nevertheless, from a predictive point of view, we find that the model that includes the LIWC dictionaries fits the data and predicts default better than a model that does not include the textual information (for details, see the Web Appendix).

General Discussion

The words we write matter. Aggregated text has been shown to predict market trends (Bollen, Mao, and Zeng 2011), market structure (Netzer et al. 2012), virility of news articles (Berger and Milkman 2012), prices of services (Jurafsky et al. 2014), and political elections (Tumasjan et al. 2010). At the individual text-writer level, text has been used to evaluate writers' state of mind (Ventrella 2011), identify liars (Newman et al. 2003) and fake reviews (Ott, Cardie, and Hancock 2012), assess personality traits (Schwartz et al. 2013; Yarkoni 2010), and evaluate mental health (Preotiuc-Pietro et al. 2015). In this article, we show that text has the ability to predict its writer's financial behavior in the distant future with significant accuracy.

Using data from an online crowdfunding platform, we show that incorporating loan application text into traditional models that predict loan default on the basis of financial and demographic information about the borrower significantly and substantially increases their predictive ability. Using machine learning methods, we uncover the words and topics borrowers often include in their loan request. We find that, at loan origination, defaulters used simple but wordier language, wrote about hardship, explained their situation and why they need the loan, and tended to refer to other sources such as their family, God, and chance. Building on prior research and the commonly used LIWC dictionary, we infer that defaulting borrowers write similarly to extroverts and liars. We obtained these results after controlling for the borrower's credit grade, which should capture the financial implications of the borrower's life circumstances, and the interest rate given to the borrower, which should capture the riskiness of different types of loans and borrowers. Simply put, we show that borrowers, consciously or not, leave traces of their intentions, circumstances, and personality in the text they write when applying for a loan—a kind of digital involuntary “sweat.”

Theoretical and Practical Contribution

Our work contributes to the recent but growing marketing literature on uncovering behavioral insights on consumers from the text they write and the traces they leave on social media (Humphreys and Jen-Hui Wang 2018; Matz and Netzer 2017). We demonstrate that the text consumers write at loan origination is indicative of their states and traits and predictive of their future repayment behavior. In an environment characterized by high uncertainty and high stakes, we find that verifiable and unverifiable data have similar predictive ability. While borrowers can truly write whatever they wish in the text box of the loan application—supposedly “cheap talk”—their word usage is predictive of future repayment behavior at a similar scale as their financial and demographic information. This finding implies that whether it is intentional and conscious or not, borrowers' writings seem to disclose their true nature, intentions, and circumstances. This finding contributes to the literature on implication and meaning of word usage by showing that people with different economic and financial situations use words differently.

Second, we make an additional contribution to the text analytics literature. The text-mining literature has primarily concentrated on predicting behaviors that occur at the time of writing the text, such as lying about past events (Newman et al. 2003) or writing fake reviews (Ott, Cardie, and Hancock 2012), but researchers have rarely attempted to predict writers' distant future behavior.

Third, our approach to predicting default relies on an automatic algorithm that mines individual words, including those without much meaning (e.g., articles, fillers), in the entire textual corpora. Previous work on narratives used to facilitate economic transaction (e.g., Herzenstein, Sonenshein, and Dholakia 2011; Martens, Jennings, and Jennings 2007) employed

human coders and therefore is prone to human mistakes and is not scalable, which limits its predictive ability and practical use. It may not be surprising that people who describe hardship in the text at loan origination are more likely to default, as prior research has shown. However, our work demonstrates that defaulters and repayers also differ in their usage of pronouns and tenses, and those seemingly harmless pronouns have the ability to predict future economic behaviors. Research has shown that unless agents are mindful of their pronoun usage, it is unlikely that they noticed or planned to use specific pronouns to manipulate the reader (Pennebaker 2011)—lenders, in our case.

Fourth, we provide evidence that our method of automatically analyzing free text is an effective way of supplementing traditional measures and even replacing some aspects of the human interaction of traditional bank loans. Textual information, such as in the loan applications we analyze, not only sheds light on previous behavior (which may be captured by measures such as credit scores) and provides it with context but also offers information about the future, which is unlikely to be captured by credit history reports. These future events may be positive (e.g., graduating, starting a new job) or negative (e.g., impending medical costs) and certainly affect lending decisions.

Furthermore, because lending institutions place a great deal of emphasis on the importance of models for credit risk measurement and management, they have developed their own proprietary models, which often have a high price tag due to data acquisition. Collecting text may be an effective and low-cost supplement to the traditional financial data and default models. Such endeavors may be particularly useful in “thin file” situations wherein historical data about customers' finances is sparse. A Government Accountability Office (GAO 2018) report, written at the request of Congress, delineates “alternative data usage” in loan decisions by financial technology (fintech) lenders (Prosper is one of the 11 lenders surveyed). “Alternative data” is defined as anything not used by the three credit bureaus, ranging from rent payments, education, and alma mater to internet browser history and social media activity. These data are either given freely by potential borrowers or bought from data aggregators. According to the report, currently there are no specific regulations regarding the use of these data, and fintech lenders use consulting and law firms specializing in fair lending issues to test their models for compliance with fair lending laws. These lenders explain that alternative data allow them to lend money to those who otherwise would not be eligible, including small businesses requiring a small loan and individuals with low credit scores. Thus, while fintech lenders voiced their need for more guidance from the Consumer Financial Protection Bureau, the use of alternative data is unlikely to go away and, if anything, their reliance on it will only increase.

The objective of the current research is to explore the predictive and informative value of the text in loan applications. However, as mentioned in the GAO report, the way lenders use such information to decide which borrowers should be granted

loans must be carefully assessed and comply with fair lending laws and their interpretation by the Consumer Financial Protection Bureau.

Avenues for Future Research

Our research takes the first step in automatically analyzing text to predict default and therefore initiates multiple research opportunities. First, we focus on predicting default because this is an important behavior that is less idiosyncratic to the crowdfunding platform whose data we analyze (compared with lending decisions). Theoretically, many aspects of loan repayment behavior, which are grounded in human behavior (e.g., extroversion; Nyhus and Webley 2001), should be invariant to the type of loan and lending platform, whereas other aspects may vary by context. Yet because we analyze data from one platform, we present a case study for the potential value of textual analysis in predicting default. Future research should extend our work to different populations, other types of unsecured loans (e.g., credit card debt), and secured loans (e.g., mortgages).

Second, our results should be extended to other types of media and communication, such as phone calls or online chats. It would be interesting to test how an active conversation (two-sided correspondence vs. one-sided input) may affect the results. In our data, the content of the text is entirely up to borrowers—they disclose what they wish in whichever writing style they choose. In a conversation, the borrower may be prompted to provide certain information.

Third, we document specific words and themes that might help lenders avoid defaulting borrowers and help borrowers better express themselves when requesting the loan. In line with the market efficiency hypothesis, if both lenders and borrowers internalize the results we documented, these results may change. In other words, the Lucas critique (Lucas 1976) that historical data cannot predict a change in economic policy because future behavior changes as policy changes may apply to our situation. In what follows we discuss this critique on both demand for loans (the borrowers' side) and supply of loans (lenders' side). We begin with borrowers and wish to emphasize two aspects. First, an important objective of our study is to explore and uncover the behavioral signals or traces in the text that are associated with loan default. This objective is descriptive rather than prescriptive and hence not susceptible to the Lucas critique. Second, for the Lucas critique to manifest three conditions need to occur: (1) the research findings need to fully disseminate, (2) the agent needs to have sufficient incentive to change her behavior, and (3) the agent needs to be able to change her behavior based on the proposed findings (Van Heerde, Dekimpe, and Putsis 2005). With respect to the first point, evidence from deception detection mechanisms (Li et al. 2014) as well as macroeconomics (Rudebusch 2005) suggest that dissemination of research results rarely fully affect behavior. Moreover, borrowers in our context vary substantially in terms their socio-economic and education level, which makes it nearly impossible to change the behavior of all or even the majority of them. As for the second and third points, borrowers

indeed have a strong incentive to change their behavior due to the decision importance; however, they might not necessarily be able to do that. Even if our results are fully disseminated, people's use of pronouns and tenses (the finding at the heart of our research) has been found to be largely subconscious and changing it demands heavy cognitive load (Pennebaker 2011). Considering the supply of loans, we acknowledge that lenders have an incentive to learn about possible tactics borrowers employ to garner trust and empathy when lending is economically not warranted. If lenders have the ability to learn about word usage that enables bad loan requests to be sufficiently similar to good ones and hence get funded (when they should not have been), they are likely fund a different mix of loans. Put differently, if our findings are disseminated then a different set of loans would have been funded—some that were rejected would have been funded and some that were funded would have been rejected. Future research should explore how word usage changes over time and the resulting mix of funded and unfunded loans.

Finally, while we are studying the predictive ability of written text regarding a particular future behavior, our approach can be easily extended to other behaviors and industries. For example, universities might be able to predict students' success based on the text in the application (beyond people manually reading the essays). Similarly, human resource departments and recruiters can use the words in the text applicants write to identify promising candidates.

To conclude, borrowers leave meaningful signals in the text of loan applications that help predict default, sometimes years after loan origination. Our research adds to the literature using text mining to better understand consumer behavior (Humphreys and Jen-Hui Wang 2018), and especially in the realm of consumer finance (Herzenstein, Sonenshein, and Dholakia 2011).

Acknowledgment

We thank you for reading our paper, we politely ask that you cite our work, and we promise to work hard and cite you back. God bless!

Author Contributions

Oded Netzer, Alain Lemaire, and Michal Herzenstein contributed equally to this work.

Associate Editor

Michel Wedel

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Financial support from Columbia Business School and Lerner College at the University of Delaware is greatly appreciated.

References

- Agarwal, Sumit and Robert Hauswald (2010), "Distance and Private Information in Lending," *Review of Financial Studies*, 23 (7), 2757–88.
- Agarwal, Sumit, Paige M. Skiba, and Jeremy Tobacman (2009), "Payday Loans and Credit Cards: New Liquidity and Credit Scoring Puzzles," *American Economic Review*, 99 (2), 412–17.
- Anderson, Jon, Stephen Burks, Colin DeYoung, and Aldo Rustichini (2011), "Toward the Integration of Personality Theory and Decision Theory in the Explanation of Economic Behavior," presented at the IZA Workshop: Cognitive and Non-Cognitive Skills, Bonn, Germany (January 25–27), <https://pdfs.semanticscholar.org/db58/6683a98173dba870d5e8fc93920fa867a1d2.pdf>.
- Arya, Shweta, Catherine Eckel, and Colin Wichman (2013), "Anatomy of the Credit Score," *Journal of Economic Behavior and Organization*, 95 (November), 175–85.
- Avery, Robert B., Raphael W. Bostic, Paul S. Calem, and Glenn B. Canner (2000), "Credit Scoring: Statistical Issues and Evidence from Credit-Bureau Files," *Real Estate Economics*, 28 (3), 523–47.
- Balabanis, George (2002), "The Relationship Between Lottery Ticket and Scratch-Card Buying Behaviour, Personality and Other Compulsive Behaviours," *Journal of Consumer Behaviour*, 2 (1), 7–22.
- Banko, Michele and Eric Brill (2001), "Scaling to Very Very Large Corpora for Natural Language Disambiguation," in *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics*. Toulouse, France: Association for Computational Linguistics, 26–33.
- Barber, Brad M. and Terrance Odean (2001), "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment," *Quarterly Journal of Economics*, 116 (1), 261–92.
- Berger, Jonah and Katherine L. Milkman (2012), "What Makes Online Content Viral?" *Journal of Marketing Research*, 49 (2), 192–205.
- Bergstra, James and Yoshua Bengio (2012), "Random Search for Hyper-Parameter Optimization," *Journal of Machine Learning Research*, 13 (1), 281–305.
- Berneth, Jeremy, Shannon G. Taylor, and Harvell Jackson Walker (2011), "From Measure to Construct: An Investigation of the Nomological Network of Credit Scores," *Academy of Management Proceedings*, 2011 (1), 1–6.
- Bollen, Johan, Huina Mao, and Xiaojun Zeng (2011), "Twitter Mood Predicts the Stock Market," *Journal of Computational Science*, 2 (1), 1–8.
- Bond, Gary D. and Adrienne Y. Lee (2005), "Language of Lies in Prison: Linguistic Classification of Prisoners' Truthful and Deceptive Natural Language," *Applied Cognitive Psychology*, 19 (3), 313–29.
- Brown, Meta, John Grigsby, Wilbert van der Klaauw, Jaya Wen, and Basit Zafar (2015), "Financial Education and the Debt Behavior of the Young," *Federal Reserve Bank* (New York), Report #634.
- Coyle, Tim (2017), "The Thin Filed, Underbanked: An Untapped Source of Buyers," *Credit Union Times* (September 22), <https://www.cutimes.com/2017/09/22/the-thin-filed-underbanked-untapped-source-of-buyers/?sreturn=20190103095108>.
- DePaulo, Bella M., James J. Lindsay, Brian E. Malone, Laura Muhlenbruck, Kelly Charlton, and Harris Cooper (2003), "Cues to Deception," *Psychological Bulletin*, 129 (1), 74–118.
- Dholakia, Utpal, Leona Tam, Sunyee Yoon, and Nancy Wong (2016), "The Ant and the Grasshopper: Understanding Personal Saving Orientation of Consumers," *Journal of Consumer Research*, 43 (1), 134–55.
- Farrell, Joseph and Matthew Rabin (1996), "Cheap Talk," *Journal of Economic Perspectives*, 10 (3), 103–18.
- Feldman, Gilad, Huiwen Lian, Michal Kosinski, and David Stillwell (2017), "Frankly, We Do Give a Damn: The Relationship Between Profanity and Honesty," *Social Psychological and Personality Science*, 8 (7), 816–26.
- Fernandes, Daniel, John G. Lynch Jr., and Richard G. Netemeyer (2014), "Financial Literacy, Financial Education, and Downstream Financial Behaviors," *Management Science*, 60 (8), 1861–83.
- Freitas, Antonio L., Nira Liberman, Peter Salovey, and E. Tory Higgins (2002), "When to Begin? Regulatory Focus and Initiating Goal Pursuit," *Personality and Social Psychology Bulletin*, 28 (1), 121–30.
- GAO (2018), "Financial Technology: Agencies Should Provide Clarification on Lenders' Use of Alternative Data," research report (December), <https://www.gao.gov/assets/700/696149.pdf>.
- Griffiths, Thomas L. and Mark Steyvers (2004), "Finding Scientific Topics," *Proceedings of the National Academy of Sciences*, 101 (1), 5228–35.
- Gross, Jacob P.K., Cekic Osman, Don Hossler, and Nick Hillman (2009), "What Matters in Student Loan Default: A Review of the Research Literature," *Journal of Student Financial Aid*, 39 (1), 19–29.
- Hancock, Jeffrey T., Lauren E. Curry, Saurabh Goorha, and Michael Woodworth (2007), "On Lying and Being Lied to: A Linguistic Analysis of Deception in Computer-Mediated Communication," *Discourse Processes*, 45 (1), 1–23.
- Harkness, Sarah K. (2016), "Discrimination in Lending Markets: Status and the Intersections of Gender and Race," *Social Psychology Quarterly*, 79 (1), 81–93.
- Herzenstein, Michal, Scott Sonenshein, and Utpal M. Dholakia (2011), "Tell Me a Good Story and I May Lend You My Money: The Role of Narratives in Peer-to-Peer Lending Decisions," *Journal of Marketing Research*, 48 (Special Issue), S138–49.
- Hildebrand, Thomas, Manju Puri, and Jörg Rocholl (2017), "Adverse Incentives in Crowdfunding," *Management Science*, 63 (3), 587–608.
- Hirsh, Jacob B. and Jordan B. Peterson (2009), "Personality and Language Use in Self-Narratives," *Journal of Research in Personality*, 43 (3), 524–27.
- Hoffman, Matthew, Francis R. Bach, and David M. Blei (2010), "Online Learning for Latent Dirichlet Allocation," in *Advances in Neural Information Processing Systems*, <https://papers.nips.cc/paper/3902-online-learning-for-latent-dirichlet-allocation.pdf>.
- Humphreys, Ashlee and Rebecca Jen-Hui Wang (2018), "Automated Text Analysis for Consumer Research," *Journal of Consumer Research*, 44 (6), 1274–1306.
- Jurafsky, Dan, Victor Chahuneau, Bryan R. Routledge, and Noah A. Smith (2014), "Narrative Framing of Consumer Sentiment in Online Restaurant Reviews," *First Monday*, 19 (4), <https://firstmonday.org/ojs/index.php/fm/article/view/4944/3863>.

- Kosinski, Michal, David Stillwell, and Thore Graepel (2013), "Private Traits and Attributes Are Predictable from Digital Records of Human Behavior," *Proceedings of the National Academy of Sciences*, 110 (15), 5802–05.
- Kupor, Daniella M., Kristin Laurin, and Jonathan Levav (2015), "Anticipating Divine Protection? Reminders of God Can Increase Nonmoral Risk Taking," *Psychological Science*, 26 (4), 374–84.
- Lee, Dokyun, Kartik Hosanagar, and Harikesh S. Nair (2018), "Advertising Content and Consumer Engagement on Social Media: Evidence from Facebook," *Management Science*, 64 (11), 5105–31.
- Li, Jiwei, Myle Ott, Claire Cardie, and Eduard Hovy (2014), "Towards a General Rule for Identifying Deceptive Opinion Spam," in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*. Baltimore: Association for Computational Linguistics, 1566–76.
- Lucas, Robert (1976), "Econometric Policy Evaluation: A Critique," in *The Phillips Curve and Labor Markets*, *Carnegie-Rochester Conference Series on Public Policy*. New York: Elsevier, 19–46.
- Lynch, John G., Jr., Richard G. Netemeyer, Stephen A. Spiller, and Alessandra Zammit (2010), "A Generalizable Scale of Propensity to Plan: The Long and the Short of Planning for Time and for Money," *Journal of Consumer Research*, 37 (1), 108–28.
- Mairesse, François, Marilyn A. Walker, Matthias R. Mehl, and Roger K. Moore (2007), "Using Linguistic Cues for the Automatic Recognition of Personality in Conversation and Text," *Journal of Artificial Intelligence Research*, 30 (1), 457–500.
- Martens, Martin L., Jennifer E. Jennings, and P. Devereaux Jennings (2007), "Do the Stories They Tell Get Them the Money They Need? The Role of Entrepreneurial Narratives in Resource Acquisition," *Academy of Management Journal*, 50 (5), 1107–32.
- Matz, Sandra and Oded Netzer (2017), "Using Big Data as a Window into Consumers' Psychology," *Current Opinion in Behavioral Sciences*, 18 (December), 7–12.
- Mayer, Christopher, Karen Pence, and Shane M. Sherlund (2009), "The Rise in Mortgage Defaults," *Journal of Economic Perspectives*, 23 (1), 27–50.
- McLaughlin, G. Harry (1969), "SMOG Grading—A New Readability Formula," *Journal of Reading*, 12 (8), 639–46.
- McAdams, Dan P. (2001), "The Psychology of Life Stories," *Review of General Psychology*, 5 (2), 100–123.
- Mehl, Matthias R., Samuel D. Gosling, and James W. Pennebaker (2006), "Personality in Its Natural Habitat: Manifestations and Implicit Folk Theories of Personality in Daily Life," *Journal of Personality and Social Psychology*, 90 (5), 862–77.
- Netemeyer, Richard G., Dee Warmath, Daniel Fernandes, and John G. Lynch Jr. (2018), "How Am I Doing? Perceived Financial Well-Being, Its Potential Antecedents, and Its Relation to Overall Well-Being," *Journal of Consumer Research*, 45 (1), 68–89.
- Netzer, Oded, Ronen Feldman, Jacob Goldenberg, and Moshe Fresko (2012), "Mine Your Own Business: Market-Structure Surveillance Through Text Mining," *Marketing Science*, 31 (3), 521–43.
- Newman, Matthew L., James W. Pennebaker, Diane S. Berry, and Jane M. Richards (2003), "Lying Words: Predicting Deception from Linguistic Styles," *Personality and Social Psychology Bulletin*, 29 (5), 665–75.
- Nicholson, Nigel, Emma Soane, Mark Fenton-O'Creevy, and Paul Willman (2005), "Personality and Domain-Specific Risk Taking," *Journal of Risk Research*, 8 (2), 157–76.
- Norvilitis, Jill M., Michelle M. Merwin, Timothy M. Osberg, Patricia V. Roehling, Paul Young, and Michele M. Kamas (2006), "Personality Factors, Money Attitudes, Financial Knowledge, and Credit-Card Debt in College Students," *Journal of Applied Social Psychology*, 36 (6), 1395–1413.
- Nyhus, Ellen K. and Paul Webley (2001), "The Role of Personality in Household Saving and Borrowing Behaviour," *European Journal of Personality*, 15 (S1), S85–S103.
- Oppenheimer, Daniel M. (2006), "Consequences of Erudite Vernacular Utilized Irrespective of Necessity: Problems with Using Long Words Needlessly," *Applied Cognitive Psychology*, 20 (2), 139–56.
- Ott, Myle, Claire Cardie, and Jeff Hancock (2012), "Estimating the Prevalence of Deception in Online Review Communities," in *Proceedings of the 21st Conference on World Wide Web*. New York: Association for Computing Machinery, 201–10.
- Palmer, Christopher (2015), "Why Did So Many Subprime Borrowers Default During the Crisis: Loose Credit or Plummeting Prices?" working paper, University of California, Berkeley.
- Pasupathi, Monisha (2007), "Telling and the Remembered Self: Linguistic Differences in Memories for Previously Disclosed and Previously Undisclosed Events," *Memory*, 15 (3), 258–70.
- Pennebaker, James W. (2011), "The Secret Life of Pronouns," *New Scientist*, 211 (2828), 42–45.
- Pennebaker, James W. and Anna Graybeal (2001), "Patterns of Natural Language Use: Disclosure, Personality, and Social Integration," *Current Directions in Psychological Science*, 10 (3), 90–93.
- Pennebaker, James W. and Laura A. King (1999), "Linguistic Styles: Language Use as an Individual Difference," *Journal of Personality and Social Psychology*, 77 (6), 1296–1312.
- Pennebaker, James W., Tracy J. Mayne, and Martha E. Francis (1997), "Linguistic Predictors of Adaptive Bereavement," *Journal of Personality and Social Psychology*, 72 (4), 863–71.
- Pennebaker, James W. and Lori D. Stone (2003), "Words of Wisdom: Language Use over the Life Span," *Journal of Personality and Social Psychology*, 85 (2), 291–301.
- Pompian, Michael M. and John M. Longo (2004), "A New Paradigm for Practical Application of Behavioral Finance: Creating Investment Programs Based on Personality Type and Gender to Produce Better Investment Outcomes," *Journal of Wealth Management*, 7 (2), 9–15.
- Pope, Devin G. and Justin R. Sydnor (2011), "What's in a Picture? Evidence of Discrimination from Prosper.com," *Journal of Human Resources*, 46 (1), 53–92.
- Preotiuc-Pietro, Daniel, Johannes Eichstaedt, Gregory Park, Maarten Sap, Laura Smith, and Victoria Tobolsky, et al. (2015), "The Role of Personality, Age and Gender in Tweeting about Mental Illnesses," in *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*. Denver: Association for Computational Linguistics, 21–30.

- Rudebusch, Glenn D. (2005), "Assessing the Lucas Critique in Monetary Policy Models," *Journal of Money, Credit, and Banking*, 37 (2), 245–72.
- Rugh, Jacob S. and Douglas S. Massey (2010), "Racial Segregation and the American Foreclosure Crisis," *American Sociological Review*, 75 (5), 629–51.
- Rustichini, Aldo, Colin G. DeYoung, Jon E. Anderson, and Stephen Burks (2016), "Toward the Integration of Personality Theory and Decision Theory in Explaining Economic Behavior: An Experimental Investigation," *Journal of Behavioral and Experimental Economics*, 64 (October), 122–37.
- Schwartz, H. Andrew, Johannes C. Eichstaedt, Margaret L. Kern, Lukasz Dziurzynski, Stephanie M. Ramones, and Megha Agrawal, et al. (2013), "Personality, Gender, and Age in the Language of Social Media: The Open-Vocabulary Approach," *PLOS One*, 8 (9), e73791.
- Sengupta, Rajdeep and Geetesh Bhardwaj (2015), "Credit Scoring and Loan Default," *International Review of Finance*, 15 (2), 139–67.
- Shah, Anuj K., Sendhil Mullainathan, and Eldar Shafir (2012), "Some Consequences of Having Too Little," *Science*, 338 (6107), 682–85.
- Sievert, Carson and Kenneth E. Shirley (2014), "LDAvis: A Method for Visualizing and Interpreting Topics," in *Proceedings of the Workshop on Interactive Language Learning, Visualization, and Interfaces*. Baltimore: Association for Computational Linguistics, 63–70.
- Sonenshein, Scott, Michal Herzenstein, and Utpal M. Dholakia (2011), "How Accounts Shape Lending Decisions Through Fostering Perceived Trustworthiness," *Organizational Behavior and Human Decision Processes*, 115 (1), 69–84.
- Sussman, Abigail B. and Rourke L. O'Brien (2016), "Knowing When to Spend: Unintended Financial Consequences of Earmarking to Encourage Savings," *Journal of Marketing Research*, 53 (5), 790–803.
- Tausczik, Yla R. and James W. Pennebaker (2010), "The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods," *Journal of Language and Social Psychology*, 29 (1), 24–54.
- Tibshirani, Robert (1997), "The Lasso Method for Variable Selection in the Cox Model," *Statistics in Medicine*, 16 (4), 385–95.
- Toma, Catalina L. and Jeffrey T. Hancock (2012), "What Lies Beneath: The Linguistic Traces of Deception in Online Dating Profiles," *Journal of Communication*, 62 (1), 78–97.
- Tumasjan, Andranik, Timm O. Sprenger, Philipp G. Sandner, and Isabell M. Welp (2010), "Predicting Elections with Twitter: What 140 Characters Reveal About Political Sentiment," in *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*. Palo Alto, CA: Association for the Advancement of Artificial Intelligence, 178–85.
- Van Heerde, Harald J., Marnik G. Dekimpe, and William P. Putsis Jr. (2005), "Marketing Models and the Lucas Critique," *Journal of Marketing Research*, 42 (1), 15–21.
- Ventrella, Jeffrey (2011), *Virtual Body Language*. Pittsburgh: ETC Press.
- Wei, Yanhao, Pinar Yildirim, Christophe Van den Bulte, and Chrysanthos Dellarocas (2015), "Credit Scoring with Social Network Data," *Marketing Science*, 35 (2), 234–58.
- Weiss, Brent and Robert S. Feldman (2006), "Looking Good and Lying to Do It: Deception as an Impression Management Strategy in Job Interviews," *Journal of Applied Social Psychology*, 36 (4), 1070–86.
- Weinstein, Neil D. (1980) "Unrealistic Optimism About Future Life Events," *Journal of Personality and Social Psychology*, 39 (5), 806–20.
- Whalen, Sean and Gaurav Pandey (2013), "A Comparative Analysis of Ensemble Classifiers: Case Studies in Genomics," in *Data Mining (ICDM), 2013 IEEE Conference*. New York: Institute of Electrical and Electronics Engineers, 807–16.
- Yarkoni, Tal (2010), "Personality in 100,000 Words: A Large-Scale Analysis of Personality and Word Use Among Bloggers," *Journal of Research in Personality*, 44 (3), 363–73.