


Screening Peers Softly: Inferring the Quality of Small Borrowers

Asim Ijaz Khwaja, Rajkamal Iyer, Erzo F. P. Luttmer,
Kelly Shue

CID Working Paper No. 259
March 2013

© Copyright 2013 Khwaja, Asim; Iyer, Rajkamal; Luttmer, F.P.
Erzo; Shue, Kelly and the President and Fellows of Harvard
College



Working Papers

Center for International Development
at Harvard University

Screening Peers Softly: Inferring the Quality of Small Borrowers

Rajkamal Iyer Asim Ijaz Khwaja Erzo F. P. Luttmer Kelly Shue*

March 29, 2013

Abstract

The recent banking crisis highlights the challenges faced in credit intermediation. New online peer-to-peer lending markets offer opportunities to examine lending models that primarily cater to small borrowers and that generate more types of information on which to screen. This paper evaluates screening in a peer-to-peer market where lenders observe both standard financial information and soft, or nonstandard, information about borrower quality. Our methodology takes advantage of the fact that while lenders do not observe a borrower's exact credit score, we do. We find that lenders are able to predict default with 45% greater accuracy than what is achievable based on just the borrower's credit score, the traditional measure of creditworthiness used by banks. We further find that lenders effectively use nonstandard or soft information and that such information is relatively more important when screening borrowers of lower credit quality. In addition to estimating the overall inference of creditworthiness, we also find that lenders infer a third of the variation in the dimension of creditworthiness that is captured by the credit score. This credit-score inference relies primarily upon standard hard information, but still draws relatively more from softer or less standard information when screening lower-quality borrowers. Our results highlight the importance of screening mechanisms that rely on soft information, especially in settings targeted at smaller borrowers.

JEL codes: *D53, D8, G21, L81*

Keywords: *Peer-to-peer Credit Markets, Market-based Lending, Screening, Market Inference, Information and Hierarchies, Soft Information.*

* Iyer: MIT Sloan School of Management riyer@mit.edu; Khwaja: Harvard Kennedy School and NBER akhwaja@hks.harvard.edu; Luttmer: Dartmouth and NBER erzo.fp.luttmer@dartmouth.edu; Shue: University of Chicago kelly.shue@chicagobooth.edu. We are extremely grateful to Prosper.com for being so generous with their time in answering our queries. We thank Liran Einav, Raymond Fisman, Matthew Gentzkow, Lawrence Katz, Atif Mian, Enrichetta Ravina, David Scharfstein, Jesse Shapiro, Jeremy Stein, three anonymous referees and seminar participants at Harvard, the NBER finance meetings and UCLA Anderson for helpful comments. We thank David Robinson for excellent research assistance. The views expressed in this paper are solely our own.

I. Introduction

An important function of credit markets is to screen borrowers and allocate credit efficiently. Traditionally, the credit score provided by rating agencies has been the main tool banks use to screen smaller borrowers. These scores are compiled using sophisticated models based on the payment history of the borrower along with other verified information, such as the number of credit lines and the outstanding balances. However, the recent banking crisis highlights some of the challenges of traditional credit screening in predicting default. In addition, one of the difficulties faced in allocating credit to smaller borrowers is that the credit score is primarily based on historical repayment history, and is therefore very susceptible to small shocks to borrowers' financial conditions. This often creates difficulties for smaller borrowers in accessing credit.

A variety of new lending models offer potentially valuable insights on how to best address shortcomings of traditional screening mechanisms. Peer-to-peer online lending platforms provide a non-hierarchical, market-based mechanism that facilitates screening by aggregating information on borrower creditworthiness over multiple individual lenders. One distinguishing feature of these markets is the availability of nonstandard or "soft" information, which may provide valuable information about borrower creditworthiness (Petersen, 2004; Berger et al., 2002). In this paper, we evaluate how these market-based screening mechanisms compare to screening based on credit score and traditional methods. Is the screening by these markets better than the screening achievable based on just the credit score? Does the market screen high-quality borrowers more effectively than low-quality ones? How important is the role of soft/nonstandard information in screening borrowers, and does this role depend on borrower quality?

This paper uses a setting that provides an ideal laboratory to quantify collective inference by individual, non-expert, market participants. It also allows us to estimate the extent and nature of inference arising from different sources of information, such as hard and soft information. The setting used is that of an online peer-to-peer lending market, Prosper.com, where borrowers post loan listings and where multiple individual lenders bid to fund a portion of the loan at a desired interest rate. Lenders have access to standard hard financial information commonly used by banks, such as the borrower's income and number of past delinquencies. In addition, lenders can view nonstandard information, such as the maximum interest rate the borrower is willing to pay as well as softer and less quantifiable information, such as the borrower's picture and a textual description of his/her reasons for the loan application. The interest rate for a funded loan is determined through

sequential bidding and reflects the lenders' collective perception of the quality and, hence, the creditworthiness of the borrower.¹

We exploit a unique feature of the Prosper marketplace in our proprietary data: while lenders only see the borrower's aggregate credit category, we as econometricians observe a borrower's exact credit score – a much finer measure of the borrower's underlying creditworthiness. We first examine the efficiency of screening in these markets by comparing the power of the interest rate (proxying lenders' inference) set by market participants in predicting default against the default predictability obtained by using the exact credit score of the borrower. In theory, the credit score should be the best available aggregator of the standard financial variables in terms of predicting default because it is based upon a sophisticated prediction model, estimated using the same (and possibly more extensive) type of standard hard financial data. However, individual Prosper lenders may be able to improve upon the predictive power of the credit score because they make use of nonstandard and softer borrower information in addition to standard financial variables. At the same time, the interest rate set by lenders could be less predictive of default than the credit score because lenders tend to be non-expert individuals, may be driven by their personal biases, and lack access to a larger pool of data on which a credit score is based.

We find that the lenders in these markets are able to substantially outperform the credit score in terms of predicting default. We first show that the market interest rate on loans explains more variation in ex-post default than the credit score can explain. We then present a more formal comparison using tools from signal detection that are common in credit scoring. Specifically, we construct "Receiver Operator Curves" (ROCs) and show that the "area under the curve" (AUC) - a simple metric used to judge the screening power of a screening score – is both large in an absolute sense and also significantly higher for the market interest rate than for the borrower's exact credit score. In particular, the interest rate set by lenders predicts default 45% more accurately than the borrower's credit score. In addition, we compare lender inference to the best possible predictor that an econometrician could construct using all available codeable data (hard and soft). We find that the market interest rate (our sufficient statistic for lenders' inference) also screens favourably relative to this more demanding benchmark as it exceeds 80% of the AUC of the econometrician.

To address a possible concern that Prosper lenders directly use a noisy proxy for the exact credit score in forming their inference (lenders observe seven aggregate credit score categories), we

¹ The loan is funded only if the total amount bid equals or exceeds the amount requested by the borrower, and the final interest rate is determined by the highest reservation interest rate among the set of lenders that bid successfully.

examine AUC curves corresponding to interest rates set within each observed credit category. We find that even within credit categories, the AUC for the interest rate remains high and significantly outperforms the predictive power of the credit score.

Next, we examine whether the extent of lender inference differs based on borrower quality. We find that inference is greater in the higher credit categories (better borrowers) than in the lower ones. However, the interest rate set by lenders is a better predictor of default than the credit score across all credit categories. We also explore how lenders weight standard financial versus nonstandard/soft information in forming their predictions of default. We find that both sources of information are important in screening, but that inference from soft/nonstandard information is relatively more important when assessing worse borrowers.

To the extent that a higher interest rate leads directly to increased default (as in Stiglitz and Weiss, 1981), our results may not only reflect lender inference but also reverse causality. We test whether the interest rate has a causal effect on default using credit-category borders as instruments for exogenous changes in the interest rate. The intuition for the instrument is that, at the exogenously defined borders, there is a sharp jump in interest rates even though borrower quality is continuous. We do not find any evidence of reverse causality. Similarly, the interpretation of our estimates as lender inference would be threatened if lenders directly learn borrowers' exact credit scores from self-reported borrower information in the listing text or through public and private communication via Prosper's "questions-and-answers" feature. While this channel is unlikely because Prosper strongly discourages borrowers from revealing detailed personal information and a text search through all listing text does not reveal any self-reported credit scores, we further examined this possible channel by restricting our sample to the period before the introduction of the question-and-answer feature, and find similar results. Our results also hold under other sample restrictions/splits that account for periods where Prosper introduced policy/information changes.

The results above highlight the ability of lenders to infer borrower creditworthiness along dimensions *not* captured by the credit score. Inference beyond the credit score is important because the credit score is primarily based upon hard information (e.g., past repayment history) and will miss other predictors of borrower quality. However, we are also interested in how well lenders can infer the information content of the credit score itself. In the remainder of the paper, we present a complementary analysis of how well lenders infer creditworthiness along the dimension that is directly captured by the credit score. An advantage of doing so is that it allows us to develop a

methodology to obtain precise magnitudes of inference arising from different sources of information.

We find that, within a given credit category (spanning 40 points in the credit score), lenders are able to infer a third of the difference in creditworthiness that is captured by a borrower's exact credit score. This effect is economically significant because such a degree of inference allows lenders to offer a rate that is 140 basis points lower for borrowers at the top of a typical credit category than for borrowers at the bottom of that category. Given that the credit score is computed based on proprietary formulas developed by credit bureaus and not all variables that go into the computation are available to lenders, it is by no means obvious that lenders can piece together the information provided in the listing and infer a third of the true credit score.² However, we estimate that lenders infer as much as 69% of what they could have potentially extracted from the information provided on the Prosper website.

We find that, along the credit score dimension, lenders base most of their inference on standard financial variables. Yet, soft/nonstandard variables also contribute to inference. Of the soft/nonstandard variables, we find that lenders draw the most inference from the maximum interest rate that a borrower posts she is willing to pay for the loan. This rate is likely to serve as a credible signal that satisfies the single-crossing property because (i) borrowers posting too low a rate risk not having the loan funded and (ii) it is costlier for lower-quality borrowers to risk not having the loan funded as they have fewer alternate funding options. Our results suggest that, consistent with the models of cheap talk, individuals pay greater attention to the more credible signals sent.³

As before, we also find a high degree of inference from the non-coded component of the listing, especially among the lower credit categories. In general, coding soft information is challenging because it is difficult to quantify the information content of pictures or lengthy personal text descriptions. An advantage of our methodology is that we can measure the inference drawn from such information without explicitly coding it, since this inference is computed as a "residual," that is, the variation of interest rates with the exact credit score that remains after controlling for a very flexible functional form of coded information.

² The R-squareds of regressions of credit scores within each credit category on a flexible specification of all hard information variables are low (average R^2 of around 0.3), suggesting that it is not trivial to reconstruct the credit score using the hard information variables available to lenders.

³ The borrower maximum rate also censors our observations when the interest rate that the market requires to fund a listing exceeds the borrower maximum rate. As we explain in more detail in the methodology section, our estimation strategy corrects for this mechanical censoring effect.

We should note that, while overall inference is high, we also document mistakes in inference. Lenders also offer lower interest rates to borrowers who post personal pictures and friend endorsements, even though pictures and endorsements are empirically uncorrelated with underlying borrower creditworthiness. An alternative interpretation is that these are not mistakes but instead reflect charitable motives of some lenders.

Our paper contributes to the literature that examines the importance of soft information in screening. Liberti and Mian (2009), Agarwal and Hauswald (2010), and Rajan et al. (2013) find that greater hierarchical distance discourages the use of subjective and more abstract information in banks. Our findings that suggest that the use of soft information for screening is better in online credit markets, where the hierarchical distance between the borrower and lender is small, is consistent with these papers. Our paper further adds to this literature by examining the differences in the importance of soft information for screening across borrowers in different credit categories. In addition, our paper quantifies the relative magnitude of soft versus hard information used by market participants for screening borrowers. It also quantifies the increase in the accuracy in assessing borrower creditworthiness that arises from the use of soft information. By decomposing the extent of inference from different sources of information, the paper sheds light on the relevance of different types of information in markets (Crawford and Sobel, 1982; Farrell and Rabin, 1996; Berger et al., 2002; and Petersen, 2004). Our identification of mistakes in inference contributes to the literature on the possibility of manipulation in asset markets, e.g., through the provision of false, unverifiable soft information by borrowers (Camerer, 1998, and Strumpf and Rhode, 2003).

Our work complements the recent literature that specifically examines lending in peer-to-peer markets. Pope and Sydnor (2011), Ravina (2012), and Theseira (2009) examine whether these markets display discrimination based on personal attributes, such as race and physical appearance. Peer-to-peer markets may also make better use of social network information. While Freedman and Jin (2010) find evidence of adverse selection due to informational problems faced by lenders in Prosper, they also find that social networks (endorsements by friends) can help alleviate these problems. In a similar spirit, Lin et al. (2013) find that stronger and more verifiable relational networks help reduce the adverse selection problems in Prosper. In contrast to these papers that carefully document lending behavior in peer-to-peer markets, our focus is instead on evaluating the screening ability of these markets and decomposing the extent of inference along different information sources.

More broadly, our paper contributes to the literature that examines information aggregation, inference, and learning in markets. There are several theoretical papers that focus on information aggregation through prices (Grossman, 1976; Townsend, 1978; Grossman and Stiglitz, 1980; and Vives, 1993, 1995). Another strand of literature focuses on learning in decentralized markets (Wolinsky, 1990; Duffie and Manso, 2007; and Duffie, Malamud, and Manso, 2009). On the empirical front, Biais, Hillion, and Spatt (1999) and Davies (2003) examine learning in the pre-opening period in equity markets. There are also several experimental papers that examine price formation in asset markets (Plott and Sunder, 1988; Forsythe and Lundholm, 1990; Bronfman et al., 1996; Cao, Ghysels, and Hatheway, 2000; and Hanson et al., 2006). A related strand of literature examines prediction markets. Small election markets, like the Iowa electronic markets, and event markets, which rely on aggregating information from a relatively small number of non-expert individuals, seem to provide reasonably accurate predictions (Wolfers and Zitzewitz, 2004). Our results show that, despite not being financial experts, individual lenders in peer-to-peer markets can infer a significant fraction of underlying borrower creditworthiness using both soft and hard information.

Although this paper focuses on one particular market, it examines an informative setting for studying alternative screening mechanisms more broadly. Our results suggest that even in markets with non-expert participants and with individual participants making mistakes, market participants collectively perform quite well. The magnitude of inference from hard and soft information regarding borrower creditworthiness is high, and has significantly greater predictive power than the credit score. Our results highlight the importance of credible soft information over and above the use of hard information. In banks, loan officers typically acquire soft information for larger clients during the screening process. However, this process is time consuming and is often bypassed when screening smaller borrowers and in the automated underwriting process. Our results underscore the need to design better mechanisms to incorporate soft information in banking systems that rely on more rule-based lending. Given peer-to-peer markets' ability to effectively screen borrowers, and given their non-collateral-based lending structure, such markets can offer a potential capital source for small borrowers who may otherwise be limited to more costly sources of finance, such as payday lenders and credit-card debt.

II. Context and Data

A. Context

Peer-to-peer lending on the Internet enables individual lenders to locate individual borrowers and vice-versa. In the U.S., there are around twelve active, online peer-to-peer lending websites. Furthermore, in Europe and Asia, online peer-to-peer lending markets are on the rise.⁴ In this paper, we exploit unique data from Prosper.com, an online peer-to-peer lending marketplace that was founded in February 2006. It focuses on U.S. clients and intermediates capital mostly between individual lenders and small borrowers. Prosper has funded over \$440 million in loans and currently has 1,590,000 members.

All Prosper loans are personal, three-year fixed-rate, unsecured loans. Borrowers request loans by creating public listings on the Prosper.com website. They choose the amount of money to request (up to \$25,000) and the duration of the loan listing (3, 5, 7, or 10 days). The online listing consists of three components: pictures, listing text, and credit information. The pictures and text contain unverified soft information provided voluntarily by the borrower. Often, borrowers describe why they need a loan, why they are good credit risks, and their income and expenditure flows. Some borrowers also post optional pictures of themselves or of themes related to their loan purpose. The third listing component, credit information, contains verified hard information obtained by Prosper through a credit check. The credit information section contains information on each borrower's delinquencies, credit lines, home ownership status, debt, inquiries, and public records. A sample listing is provided in Appendix D.

The credit information also contains the borrower's credit category. According to the Prosper.com website, "A credit category is what potential lenders use to measure your likelihood of repaying money you have borrowed based on your past history." Prosper assigns each borrower to one of seven credit categories based on the borrower's Experian ScoreX PLUS credit score. Of particular importance for the empirical strategy used in this paper is that the exact credit score is not observed by Prosper lenders or borrowers: participants in the Prosper marketplace observe only credit categories. The relationship between credit scores and credit categories is shown below.⁵

Category:	HR	E	D	C	B	A	AA
Score:	520-559	560-599	600-639	640-679	680-719	720-759	760-900

⁴ See http://en.wikipedia.org/wiki/Peer-to-peer_lending.

⁵ The above credit-category chart reflects the Prosper classification at the end of our sample period. A major change in credit-category criteria occurred on February 12, 2007. Prior to the credit criteria change, the credit categories were set such that: HR(0-539), E(540-600). After February 12, 2007, credit scores below 520 were disqualified and the credit-category stratification was finalized to the numbers described in the chart. For consistency of results, we restrict our sample to the period after February 12, 2007. However, results are robust to using the sample from before February 12, 2007 (see Table 4).

In addition, borrowers can join borrower groups led by “group leaders.” The ratings and financial rewards of group leaders depend on the payment profiles of the group’s members. Therefore, group leaders often pledge to exert social pressure on group members to repay loans. Group leaders can write public messages endorsing the borrower and can bid on group members’ loans. In addition, borrowers can become friends with other registered Prosper users. These friends can add public friend endorsement texts to listings and can cast friend bids on listings.

After listings are posted, lenders can browse through Prosper’s website for listings to bid on. Multiple lenders can bid on and fund each listing. Lenders can bid on portions of listings (\$50 minimum) and set their reservation rates, the lowest interest rate at which they are willing to fund the listing. The bidding begins at the maximum interest rate the borrower is willing to pay. The listing is funded only if the total amount of money bid by lenders matches or exceeds the loan amount requested by the borrower. If the total amount bid by lenders is greater than the amount requested by the borrower, the interest rate is bid down. Lenders with lower reservation interest rates are given priority in the bidding hierarchy. The final interest rate is determined by the highest reservation interest rate among the set of lenders that successfully bid for the loan.

After the listing is funded and approved by the borrower, the borrower begins to make monthly payments that are divided across lenders in proportion to each lender’s winning bid size. The borrower never directly interacts with the lenders, and all payments are routed via Prosper. If a borrower is late in making payments or defaults on the loan, his behavior is reported to the major credit agencies and the borrower’s credit rating suffers. If the borrower is late for four or more months, Prosper sells the loan to a collection agency and splits the proceeds among the lenders.

B. Data

Our dataset contains all credit information variables displayed on a borrower’s loan listing, as well as the text of the listing and the complete history of each borrower’s loan repayment stream. In addition, our data includes the credit score (unobserved by lenders and borrowers) for each

borrower.⁶ Our sample contains all listings posted between February 12, 2007 and October 16, 2008.⁷ The sample covers 194,033 listings, 17,212 of which were funded.

Table 1 provides summary statistics of the variables used in our analysis. We provide statistics for both the universe of listings (funded and unfunded) and the set of funded listings (listings that resulted in loans). We further divide the set of variables into standard financial variables and soft/nonstandard variables. The standard financial variables include hard information from the borrower’s credit report that is typically used by traditional banks. As expected, funded listings tend to have borrowers with better credit scores – in particular, funded listings tend to have far fewer “high risk” borrowers (those in the lowest credit categories). Still, 30.6% of the funded listings default at some point in the 3-year duration of the loan, with the default rate ranging from 14.7% in credit category AA to 51.6% in credit category HR. Because defaults often occur after some of the principal has already been repaid, and because some of the principal gets recovered when a defaulted loan is sold off to a collection agency, the fraction of the principal repaid at the end of the loan term is higher than one minus the default rate. The fraction repaid is 79.7% on average, and ranges from 91.0% in credit category AA to 62.5% in credit category HR. Among the universe of listings, the average loan amount requested is \$8015. The maximum interest rate borrowers are willing to pay is 21% on average. Lower credit categories and higher debt-to-income ratios are disproportionately represented among Prosper listings. For example, the average listing corresponds to a debt-to-income ratio of 54%. Funded listings tend to have better credit variables because listings representing individuals with better credit variables are much more likely to be funded. The debt-to-income ratio among funded listings is substantially lower at 33%.

The soft/nonstandard variables capture soft information that may be difficult to fully quantify, as well as information that is quantifiable but not typically used by banks, i.e., nonstandard variables that represent borrower choices. Borrower choice variables include the maximum interest rate the borrower is willing to pay, the listing duration (number of days the listing remains public), and listing category (e.g., debt consolidation or student loan). We also code basic proxies for soft

⁶ Note that even borrowers do not have access to the exact Experian ScoreX PLUS credit score obtained from the credit rating agency because it is not available for purchase by borrowers. We are able to work with this data under a non-disclosure agreement that safeguards the confidential and proprietary nature of some of the variables in the dataset.

⁷ Prosper entered a “quiet period” in October 2008, during which it ceased making new loans in anticipation of an SEC cease-and-desist procedure. Prosper emerged from the quiet period in July 2009 using a new system of classifying prospective borrowers into credit categories. We therefore do not use data on loans originating after October 2008. More importantly, because we want to observe the full repayment profile of each 3-year loan, we can only use data on loans that originated 3 or more years ago. We also use data from May 2006 to February 12, 2007 as part of a robustness check. However, we exclude data from this period in our baseline sample because the credit-category boundaries changed on February 12, 2007. See Section 2, Part A for more details.

information, such as whether the borrower posts a picture or the number of words used in the listing text descriptions. We code the soft information in order to roughly estimate the relative importance of pictures, listing text, friend endorsements, etc., for lender inference. However, we do not attempt to fully quantify the large selection of soft information available in Prosper listings. Rather, as we explain in the next section, we develop a methodology to measure how much inference is drawn from *uncoded* sources of listing content.

III. Methodology

A. Estimating Screening Performance

We begin by examining the ability of the Prosper marketplace to infer borrower quality, as proxied by ex-post loan performance. Under the assumption that the objective of marginal lenders on the Prosper marketplace is to maximize the returns on their portfolios, the interest rate is the market’s best predictor of loan performance. In the results section, we discuss the consequences of relaxing this assumption. We measure the quality of a screening method as (a) the simple goodness-of-fit (R^2) from a linear regression of ex-post loan performance on the predictor of loan performance used by that screening method (Appendix A details why this, and not the regression coefficient, is the appropriate statistic) and, more formally, as (b) the area under receiver operator curves (ROCs), a technique that is standard in the signal detection literature and commonplace in commercial financial banking markets. We describe the latter in detail below. We prefer using ROC curves over R^2 , as they provide a more interpretable estimate of inference, although one downside of the former is that it requires a binary outcome measure. This poses no problem when we use default as the measure of ex-post loan performance, but it precludes us from using ROC curves with the fraction of loan repaid as the outcome measure. In that case, we will present only the R^2 measures.⁸

While these methods provide an absolute and interpretable sense of how well a classification/screening method performs, it is also helpful to compare the quality of inference by Prosper lenders to other standard benchmarks. The primary benchmark we use is the ability of a borrower’s credit score to predict default. The credit scores is the most common instrument used by banks to screen borrowers and, hence, serves as a useful benchmark. In addition, we will go beyond

⁸ Alternatively, we could define a fraction repaid “threshold value” to provide a binary outcome classification. Doing so would provide results that are analogous to those obtained by using default.

the credit score and develop a more challenging benchmark based on what an econometrician would do if he had access to all the coded information available in the Prosper marketplace.

Receiver Operator Curves

If the outcome of interest, in our case default, is a binary indicator, then a simple and very standard way of measuring the performance of a screening procedure is to compare how accurately the procedure can identify both 0 and 1 values of the outcome indicator. Consider the following example below, where the columns give the actual value of the outcome (default = 0 or 1) and the rows are the predicted value given by the screening procedure.

		Actual Value	
		1	0
Predicted Value	1	True Positive	False Positive
	0	False Negative	True Negative

In a perfect classification, only the diagonal entries that are populated, i.e., both the True Positive Rate = $(TP/(FN+TP))$; the “sensitivity” of a screening test) and the True Negative Rate = $(TN/(FP+TN))$; the “specificity” of a screening test) are 100%.⁹

Now consider an actual screening procedure that creates a continuous predictor for a binary outcome. Consider a hypothetical threshold value above which we assign one value of the outcome and below which, the other value. Consider credit scores as an example. We may say individuals above a score of 600 are predicted not to default (outcome=0), while those below are predicted to default (outcome=1).¹⁰ Each threshold value will generate a different table above. By construction, as the threshold value changes, we face a (potential) tradeoff between the True Positive and True Negative Rates. In the example given, it is easy to see that as we lower the threshold, we predict more people as non-defaulters. Setting the threshold as low as possible thus results in everyone

⁹ Note that this terminology is analogous to what is more typically used in statistics – Type I and Type II errors. A Type I error occurs when we incorrectly reject a true hypothesis (this corresponds to a false positive), and a Type II error occurs when we incorrectly accept a false hypothesis (this corresponds to a false negative).

¹⁰ While credit scores are a continuous measure of the probability of default, it is important to highlight that what we actually observe is always a binary outcome (i.e., the borrower either defaults or does not default). Hence, the credit scoring models that are used to construct a credit score actually follow this procedure, i.e., they estimate the optimal weights for the underlying variables that go into forming a credit score (such as repayment history, level of indebtedness, demographic information, etc.). The model is then judged by how well it classifies the ultimate outcome – a binary indicator of default.

being classified as a 0 (non-defaulter), and gives a 100% TNR but a 0% TPR. Likewise, when we increase the threshold value to its strictest possible value, we end up with a 100% TPR (in that every true defaulter is correctly identified simply because we are classifying everyone as defaulters) but a 0% TNR.

A Receiver Operator Curve (ROC) is a convenient way of illustrating this tradeoff at each possible threshold. As we vary the threshold, the ROC curve plots the True Positive Rate on the y-axis against the False Positive Rate ($= 1 - \text{TNR}$). Figure 1 illustrates this for a hypothetical screening procedure, with each point on the curve showing the TPR and FPR for a particular threshold value. As explained above, this curve will always start at the origin and end at the right hand corner. A “perfect” screening tool is represented by a single point on the top-left hand corner (TPR=100% and FPR=0%), while the worst possible screening tool would be one that is essentially random, causing the curve to lie exactly on the 45° line.¹¹

More specifically, the cardinal measure used to quantify the accuracy of a screening tool is the area under the curve (AUC), which ranges from 1 (perfect system) to 0.5 (worst possible screening tool - purely random predictions). This corresponds closely to a Gini coefficient ($\text{Gini} = 2 * \text{AUC} - 1$), and the literature in credit scoring uses both the AUC and the Gini as a way of quantifying how good a screening tool is. We will present results using the AUCs.

While there are no obvious benchmarks for levels of AUC since it depends on the usage context (and cannot be readily compared across different samples with varying underlying population default rates), it is the most common metric used in the credit-scoring industry. As a rough rule of thumb, an AUC of 0.6 or greater is generally considered desirable in information-scarce environments, while AUCs of 0.7 or greater are the goal in more information-rich environments.

We estimate the AUCs (and show the ROC curves) using the market interest rate for a given loan as determined in the Prosper marketplace, since this acts as a simple metric by which the market judges the quality of a loan (i.e., higher interest rate loans are deemed to be of lower quality).

¹¹ A perfect tool would correctly rank all borrowers. Hence, tightening the threshold would at first only reject bads (move along the y-axis while the x-axis stays at 0). Only after tightening the threshold to the point that all of the bads are rejected (the top-left hand corner of the graph) would further tightening start rejecting the goods, moving the curve along the y-axis until all applicants are rejected. Conversely, the ROC curve for a screening tool with no predictive power would not distinguish bads from goods, causing them to be evenly distributed across all scores/values of the screening tool. Thus, starting with the worst value of the score and tightening the threshold would lead to a rejection of both goods and bads in equal proportion (to their population), and the resulting ROC curve would be along the diagonal line. Note that this diagonal is the worst a screening tool could do. If the ROC curve is strictly below the diagonal then one can simply invert the score and come up with a better (than random) screening mechanism.

We will compare the AUC of the market interest rate to the AUC using the credit score of the borrower. Note that a unique feature of the Prosper marketplace is that, while we as econometricians observe the credit score, lenders only observe aggregate credit categories. In addition, we will also compare the market interest rate AUC to the AUC of the best possible score that an econometrician could create based on a regression of ex-post default on all the coded characteristics of the borrower listing. This latter AUC approximates an upper bound for the predictive power of the best possible screening tool given all coded information available.¹² In creating the best possible screening score, one must be cautious about over-fitting. Given small samples, it is possible to achieve high in-sample AUCs but have poor predictive power out of sample. Therefore, we estimate the best possible score using a split sample approach. We regress ex-post default on all observed borrower characteristics for a randomly chosen half of the sample to obtain the regression coefficients. We then combine these coefficients with the observed borrower characteristics of the other half of the sample to create the best possible score.

Finally, we test whether the AUCs for the market interest rate and the other benchmark screening tools are statistically different from one another using the non-parametric approach developed in DeLong, DeLong, and Clarke-Pearson (1988) and implemented in the STATA routine “roccomp.”

B. Estimating Sources of Inference

In addition to estimating how well the market is able to screen, we investigate how lenders weight different sources of information in forming their screening measure (the market interest rate). Of particular interest is the relative importance placed on standard hard financial variables (such as debt-to-income ratio and past delinquencies) that banks typically use in addition to the credit score, and soft/nonstandard variables that banks typically would not use but are more common in P2P marketplaces (such as a borrower’s number of friend endorsements, maximum rate at which they are willing to borrow, and self-descriptions/pictures).

We estimate the marginal contribution to inference from each of three sets of information as described in Section 2: standard financial variables (including the credit-category bins that lenders observe), nonstandard variables, and uncoded listing content (which represents soft information that

¹² Our estimate of the best possible screening tool uses only coded information. There may be some aspects of soft information, such as the information content of a picture, that we are unable to code. In later sections, we develop a methodology to measure inference from uncoded listing content. For now, we note that it is possible to outperform our best possible score by effectively using uncoded information contained in softer sources, such as pictures or listing text.

is difficult to quantify). We define the marginal contribution as the extent to which inference would improve if lenders were able to observe one additional set of information, conditional on having already observed the other two sets of information. Note that if the information content of one set of variables is strongly correlated with the information content of another set of variables, the marginal contribution from both will be low (because it is already contained in the other set of information). As our results will show below, this is indeed the case for inference from higher credit categories.

We measure the information from uncoded listing content contained in the interest rate as the residual from a regression of the interest rate on a flexible functional form of all the coded listing content (standard financial variables and nonstandard variables). This residual thus allows us to quantify the contribution of uncoded listing content, and we refer to it as the “deduced measure for uncoded information.”

To estimate the marginal inference drawn from standard financial information, we estimate a first-stage regression of the interest rate for the loan on a flexible functional form of nonstandard variables and the deduced measure for uncoded information. The residual from this first-stage regression only contains variation in the interest rate that comes from standard financial variables that cannot be explained by the other two sets of information. We then estimate an ROC curve using the residual from this first-stage estimation. The ROC curve’s corresponding AUC represents the marginal inference from standard financial variables. Alternatively, we can measure marginal inference using R-squareds instead of AUC. In a second-stage specification, we regress default on the residual from the first-stage regression. The R^2 from the second-stage estimation measures the marginal inference drawn from standard financial variables.

We follow the same procedure to measure the marginal contribution of inference from nonstandard variables, except that now the first-stage regression has controls for standard financial variables and the deduced measure for uncoded content. Similarly, to measure inference from uncoded listing content, we use a first-stage regression that has standard financial variables and nonstandard variables as controls.

C. Inferring the Unobserved Credit Score

The previous two sections describe methods that compare the inference of borrower quality in the Prosper marketplace against the predictive powers of the credit score. We thus measure whether lenders are able to infer borrower creditworthiness as measured by ex-post default along

dimensions *not* captured by the credit score. Inference beyond the credit score is important because the credit score is primarily based on hard information (e.g., past repayment history) and may miss other valuable predictors of borrower quality. However, we are also interested in how well lenders are able to infer the information content of the credit score itself, given that they can only observe aggregate credit categories. In this section, we describe a complementary methodology that measures how well lenders are able to infer creditworthiness along the dimension that is directly captured by credit score and provides an exact decomposition of inference by source of information.

The main idea is straightforward. Lenders observe credit categories but not the exact credit score. Consequently, if they offer loans at lower interest rates to borrowers who have better credit scores *within* a given credit category, then lenders must have correctly inferred from other information that these borrowers are more creditworthy than others in the same credit category. We can precisely quantify lenders' inference of creditworthiness along the dimension captured by the credit score by comparing the degree to which the interest rate declines with the exact credit score within credit categories to the overall decline in the interest rate across credit categories.¹³ A benefit of this methodology is that it provides an exact decomposition of the magnitude of inference of borrower creditworthiness by source of information.¹⁴

We illustrate our empirical methodology with a stylized graph of the relationship between the exact credit score and the market interest rate. The x-axis of Figure 2 plots the borrower's exact credit score, which is a proxy for one dimension of creditworthiness. Since the repayment probability is higher for more creditworthy people, the market interest rate should fall monotonically in the credit score if lenders could observe the true score (as shown by the dashed blue line). If the credit-score categories were the only information that lenders observed, the interest rate would be constant within categories and would only jump at the category borders. Thus, if we observe that the interest rate falls *within* credit-score categories, it must be the case that lenders are able to infer information about borrowers' creditworthiness from information other than the categorical credit-score variable (as illustrated by the discontinuous, downward-sloping, solid red line).

¹³ While the context is different, our method of using information not available to Prosper lenders to measure inference is similar to Farber and Gibbons (1996) and Altonji and Pierret (2001), who estimate employer inference of worker quality using AFQT scores, which are observed by the econometrician but not by the economic agents.

¹⁴ Unlike the decomposition described in the previous section (which provides marginal contributions that do not sum to total inference), this decomposition of inference by information sources is exact. We are able to do an exact decomposition here because we measure inference along the credit score dimension as a slope coefficient, and slope coefficients can be exactly decomposed into variables that contribute to it through the standard omitted bias formula. In contrast, we measured inference in the previous section by R-squareds and AUCs, which cannot be exactly decomposed.

The degree to which lenders infer creditworthiness from this other information is given by the amount by which the interest rate falls within credit-score categories relative to the total drop in interest rates, both within and between credit-score categories. In the figure, the interest rate drops by an amount β within each credit-score category and drops discontinuously by an amount α at each credit-score boundary. Hence, the total drop over one credit category (including one boundary) equals $\alpha + \beta$. Of this total drop, the interest rate falls by β due to the change in creditworthiness that lenders inferred from information other than credit category. We denote the fraction of information learned from all sources other than credit category by the symbol $\gamma \equiv \beta / (\alpha + \beta)$, and refer to γ as the amount of “inference” made by lenders along the credit-score dimension.

In this stylized setup, the following regression yields parameter estimates α and β from which the fraction of information inferred, γ , can be calculated:

$$InterestRate_i = \mu + \alpha Cat(CreditScore_i) + \beta CreditScore_i / CatSize + \varepsilon_i \quad (1)$$

where $InterestRate_i$ is the interest rate charged on loan i , $CreditScore_i$ is the exact credit score of the borrower of loan i , and $Cat(.)$ is a scalar that denotes the category of the credit score; because there are 7 credit-score categories, $Cat(.)$ takes on the integers 1 through 7. $CatSize$ is a constant that is equal to the range of credit scores that each credit category spans. Finally, ε denotes the error term, and the remaining Greek symbols are parameters to be estimated.

If we move from the starting point of one credit category to the starting point of the next category, the interest rate changes by α at the credit-category border (because $Cat(CreditScore_i)$ increases by one at the border) and changes by β within the credit category (since $CreditScore_i / CatSize$ increases by exactly one within each credit category). The fraction of this total change that lenders infer from information other than the credit-score categories is given by $\gamma = \beta / (\alpha + \beta)$. Thus, a γ of zero means that lenders are not at all able to infer creditworthiness from information other than the credit-score categories, whereas a γ of one implies that lenders are perfectly able to infer creditworthiness from the information provided. Our methodology does not rule out perverse values of γ : negative values of γ indicate that lenders interpret information that is related to higher exact credit scores as signs of lower creditworthiness, and values of γ greater than one mean that lenders place too much value on information indicating higher creditworthiness.

In Appendix B, we present a more formal analysis of the case (which we actually consider) where we allow for the underlying relationship between interest rate and exact credit score to be

non-linear and the bin sizes of credit categories to be of unequal size. We also detail how this method allows us to decompose inference along the credit-score dimension by source of information. For this decomposition, we include in equation (1) a flexible functional form in variables for each source of information. If an information source contributes to inference, it will lower the coefficient on the credit score. The standard omitted variable bias formula then determines the exact amount by which each included control variable lowers the coefficient on the credit score, which corresponds to that control variable’s contribution to inference. While it may seem challenging to quantify or code qualitative data (such as pictures and other personal details), an advantage of our methodology is that we can still derive the contribution of such information: the contribution of uncoded information is inferred from the remaining relation between the exact credit score and interest rate within credit categories while controlling for a flexible functional form of all coded information.

IV. Results

We first examine the extent to which the interest charged by lenders is able to predict loan performance. We compare the predictive power of the interest rate set by market participants against the predictive power of two benchmarks: the exact credit score of the borrower and the best predictor achievable by an econometrician relying on all coded listing content. Next, we explore how the accuracy of lender inference differs across borrower quality, and examine the weight placed by lenders on different sources of information (standard financial versus soft/nonstandard information) for inference. Finally, we focus on inference along the credit-score dimension of creditworthiness because, along that dimension, we can provide an exact decomposition of inference by information source.

A. How Well Does the Market Screen?

Table 2 first examines the predictive relationship between interest rates and loan performance. Since we observe the full three-year history of the loans (and only consider matured loans), our outcomes of interest are the default rate (an indicator for whether the loan is over 3 months late) and the fraction of the principal repaid after the end of its term. The functional form of our specification is derived in Appendix A and uses $1/(1+r)$ as our independent variable, where r denotes the interest rate. In panel A, we find that the interest rate is indeed a significant predictor of

default and fraction repaid. The adjusted R^2 of the regressions are 8% and 9% respectively. To provide a benchmark for the performance of the interest rate, we examine default predictability using exact credit score in Panel B. We find that the credit score also significantly predicts default and fraction repaid, but note that the adjusted R^2 obtained using the interest rate is nearly double of that obtained using just the credit score (8% using the interest rate versus 4% using the credit score for default as an outcome variable).¹⁵

These initial regressions show that the interest rate set by market participants is a better predictor of individual creditworthiness than the exact credit score. To rule out the possibility that the interest rate merely performed better because it happened to have a better fitting functional form than the credit score, we re-ran the credit score regressions with a very flexible functional form (a 7-part spline with breakpoints at the credit-category boundaries). This did not meaningfully increase the R^2 , which suggests that market participants are able to use the available hard and soft information on Prosper listings to infer creditworthiness along dimensions not captured by the credit score. Given that the credit score we use is specially designed for the subprime market, the ability of market participants to collectively outperform the credit score in predicting loan outcomes is remarkable. This finding is even more remarkable given the possibility that lenders do not exclusively try to maximize the returns on their loans but also have charitable or other motives. In such cases, the interest rate consists of a component reflecting a lender's inference plus a "noise" term reflecting other possible motives. As a result, the interest rate is not as good a predictor of loan outcome as the inference component by itself.

One possible concern with the finding above is that higher interest rates by themselves could lead to borrower default by increasing the burden on borrowers (Stiglitz and Weiss, 1981). To address this concern, we estimate the causal effect of the interest rate on default using credit-category borders as instruments. The intuition for the instrument is that there is a sharp jump in the interest rate at the exogenously defined borders (AA, A, B, etc.), but that underlying creditworthiness should be smooth at the borders because the credit score does not change discontinuously there. Panel C shows no evidence of a causal effect of the interest rate on loan performance, and we conclude that our results are not driven by reverse causation.

As detailed in the methodology section, a more formal way to compare the predictive power of two screening mechanisms is through ROC curves. Recall that an AUC of 0.6 or above is

¹⁵ See also Adams et al. (2009), Einav et al. (2013) and Keys et al. (2010), who find that credit score is a good predictor of default.

generally considered to be a reasonable benchmark in information-poor environments, such as screening of small borrowers. The first two panels of Figure 3 show the ROC curves for the interest rate and credit score. We find that the area under the ROC curve (AUC) for the interest rate is 0.6824, while the AUC using the exact credit score is lower at 0.6254. Note that a completely uninformative screening system has an AUC of 0.5 and that even a 0.01 improvement in AUC is considered a noteworthy gain in the credit scoring industry. A straightforward calculation shows that lenders are able to predict default with 45% greater accuracy than what is possible by using just the borrower's credit score, an improvement that is significant at the 1% level.¹⁶ This is particularly noteworthy since the credit score in question – Experian ScoreX PLUS credit score – is designed by Experian as a special score (not even available to borrowers themselves) that is supposed to be better tailored to the types of borrowers on Prosper and reduces the need for additional information.¹⁷

As alternative benchmarks, we also construct the ROC curve based on the best possible score that an econometrician would construct if he used (i) all observable standard financial variables, including the exact credit score, or (ii) all available coded information, i.e., both standard financial variables and nonstandard/soft variables, but not uncoded soft information. These scoring systems offer a maximum AUC of 0.719 and 0.724, respectively. In comparison to these more demanding benchmarks, Prosper lenders still perform fairly well – they infer more than 80% of the variation in creditworthiness estimated by the econometrician.¹⁸

Screening by Borrower Quality

The results thus far show that, on average, the interest rate is a substantially better predictor of loan performance than credit score. We now explore whether the superior performance of the Prosper marketplace participants (over credit-rating agencies) holds across different ranges of borrower quality. Table 3 presents the analogous regressions to Table 2 separately for low-quality

¹⁶ We calculate the percentage improvement as $(0.6824-0.5)/(0.6254-0.5) = 1.45$, where 0.5 is subtracted from both the interest rate and credit score AUC because 0.5 is the AUC under a non-informative (random) scoring system.

¹⁷ Experian claims that the ScoreX PLUS provides “a stronger separation of good and bad accounts and classifies more bad accounts into the worst-scoring ranges.” They further note that “traditional risk models typically are unable to score a significant percentage of consumers due to insufficient credit information,” but that “with ScoreX PLUS, almost all consumers can be effectively scored to rank order risk, thereby reducing the need for manual review” and that “in multiple market validations with traditional risk models, ScoreX PLUS performed better in over 90 percent of the head-to-head comparisons.” See http://annualcreditreport.experian.com/products/scorex_plus.html and http://annualcreditreport.experian.com/products/pdf/scorex_plus.pdf.

¹⁸ We again adjust both AUC's by 0.5, which corresponds to zero inference, and estimate that $(0.6824-0.5)/(0.7243-0.5)=0.81$ and $(0.6824-0.5)/(0.7190-0.5)=0.83$.

borrowers (credit categories HR, E, D, and C) and high-quality borrowers (credit categories B, A, and AA).¹⁹ Panels A and B show that the interest rate explains a higher fraction of the variation in default among high-quality borrowers than among low-quality borrowers. The adjusted R^2 from the regression of default on interest rate for low-quality borrowers is 3.5%, while that for high-quality borrowers is 11%. Similarly, the adjusted R^2 for fraction repaid on interest rate for low-quality borrowers is 4.6%, while that for high-quality borrowers is 12%. Panels C and D examine the differences in inference using credit score and find similar patterns. The adjusted R^2 for default predictability for low-quality borrowers using credit score is 1.9%, while that for high-quality borrowers is 2.8%. The adjusted R^2 for fraction repaid using credit score for low-quality borrowers is 2.5%, while that for high-quality borrowers is 2.8%. Altogether, the results show that overall predictability is higher for high-quality borrowers and that the level of predictability obtained using interest rates is in each case higher than that obtained using credit score.

Figures 4 and 5 show the ROC curves for low- and high-quality borrowers. In line with the earlier results, we find that the AUC using the interest rate is larger than the AUC using the credit score for both low- and high-quality borrowers. The differences are large – the interest rate outperforms the credit score by 52 percent for low-quality borrowers and by 100 percent for high-quality borrowers. As before, we also present the more demanding benchmark of the best prediction an econometrician would produce using all available coded information, and we find that the market performs well against this benchmark for both the low- and high-quality borrowers.

Robustness

A potential concern with the analysis above is that lenders on Prosper do observe credit categories, even if they do not observe the exact credit score. One could argue that rough credit categories already incorporate much of the information embedded in credit scores and, hence, the comparison between the predictive powers of the interest rate (which may include information on credit categories) and credit score will always give the interest rate somewhat of an edge. To address this concern, we examine the ROC curves and AUC *within* each credit category bin. By examining how well the interest rate and credit score predict borrower quality within a given credit category, we can be assured that the Prosper lenders are not exploiting any credit-score information. Rows (1) through (7) of Table 4 show that for each of the seven credit categories, the interest rate

¹⁹ We chose this categorization because it provides us with a roughly equal number of loans in each category.

outperforms the exact credit score in predicting default, with the difference being significant at the 1-percent level in six out of seven credit categories.

The remaining columns in Table 4 confirm that our main finding – that the AUC for the market interest rate is higher than that for credit score – holds across a range of sub-samples that address various robustness concerns. Row (9) addresses the concern that the interest rate may be more predictive of default than the credit score because borrowers may directly inform lenders of their exact credit score. Prosper allows borrowers to post information in the listing and also has a facility for questions and answers (intermediated via Prosper). However, because all of this information is unverified, borrowers would have an incentive to report the highest possible credit score within a credit category. Moreover, in an automated text search of listing text, we did not find any instance of borrowers' reporting their credit scores. Additionally, in personal communications with Prosper staff, we were told that great care was taken by Prosper to purge any personal references or personal information, such as credit score or Social Security number. Finally, we estimate lender inference in the sample period (prior to February 12, 2007) when there was no facility for questions and answers. As the results show, the interest rate also significantly outperforms the credit score in this earlier period.

Another potential concern is that Prosper introduced several changes in its policy over the sample period and that these may, in turn, affect our inference estimates and interpretation. For example, one could imagine that suggested ranges provided by Prosper to the borrowers in setting the borrower maximum rate might impact the extent of inference. Also, Prosper introduced portfolio plans that could have a similar impact if the portfolio lenders were guided by Prosper. However, our results suggest that these changes are not a concern in practice. In rows (10) and (11), we estimate the AUCs for the sample before and after these changes, finding similar results. Another concern might be that borrowers in some states are subject to usury laws (Rigbi, 2013). These laws may create an artificial ceiling on interest rates and impact the extent of inference. In row (12), we also estimate inference in the period without usury law restrictions, and we again find similar results.

We also carry out several other robustness checks. To address the concern that some borrowers are affiliated with groups in which group members might know each other and share personal information, we estimate inference in row (13) for the sample restricted to borrowers who are not affiliated with any group. In addition, to make sure that the inference is not driven by learning about individual borrowers from previous listings or other loans availed by the same borrower (e.g., default observed in previous loans), we estimated the AUCs for a sample restricted to

first-time listings (row 14) and to first-time loans (row 15). We again find similar results in terms of the interest rate outperforming credit score in predicting default.²⁰

B. Sources of Inference

In this section, we examine how lenders use different types of information to construct the interest rate. In particular, how much weight do they place on standard financial information versus soft/nonstandard information in forming an opinion of borrower quality?

We group information into two broad categories of interest: standard financial variables (information generally used by banks to screen small borrowers) and soft/nonstandard information (information that is either difficult to quantify or unique to P2P markets). Generally speaking, standard financial variables tend to be hard, verifiable, and readily coded. We provide the details and summary statistics of variables included in this category in Table 1.²¹ Meanwhile, soft/nonstandard information tends to be subjective, non-financial, potentially harder to verify, and more likely to behave like signals. Of particular interest are the various “softer” pieces of information such as pictures, individual background, description, and online exchanges, which are readily identifiable, but much harder to code in a way that is suitable for empirical analysis. For example, one may be able to code whether a listing has a picture or even attributes about the picture, but it is impossible to code all aspects of the picture from which lenders might draw information. We are nevertheless able to quantify the information of uncoded information that is incorporated into the interest rate. We do so by creating the deduced measure of uncoded listing content as the residual of a regression of the interest rate on a flexible functional form of all sources of coded information. We are thus able to distinguish between the coded content of soft/nonstandard information and the uncoded content, which leaves us with three sources of information: standard financial variables (which are coded), soft/nonstandard *variables* (which are coded), and the deduced measure of uncoded listing content (which also reflects soft/nonstandard information).

To estimate the marginal contributions drawn from each type of these three types of information, we follow the methodology described in Section III. Our estimates measure marginal contributions for each type of information, assuming that lenders already see the other two types of

²⁰ We would caution against making too much of differences in AUC across rows/samples in a given column since the sample size (and baseline default) varies noticeably across these samples.

²¹ For the sake of brevity, Table 1 does not provide summary statistics for 66 borrower occupation dummies and 52 borrower state-of-residence dummies (50 states, District of Columbia, and Puerto Rico). However, these variables are included as controls in the relevant specifications in Table 5, Table 7, and the Appendix tables.

information. Panel A of Table 5 presents the fraction of inference on the default rate. Columns (1) to (3) report the results for low-quality borrowers (bottom four credit grades) while columns (4) to (6) report the results for high-quality borrowers (top three credit grades).

In row (1), we focus on the inference using the AUC measure for the marginal contribution of each source. We find that all three types of information contribute toward inference. For low-quality borrowers, the AUC from standard financial variables is 0.573, whereas the corresponding figure for high-quality borrowers is 0.591. Thus, inference from standard financial variables is 0.019 lower for low-quality borrowers than for high-quality borrowers. Note that, in AUC terms, even a 0.019 drop is considered fairly substantial. For low-quality borrowers, the AUCs for soft/nonstandard variables and uncoded listing content are 0.542 and 0.538. For high-quality borrowers the corresponding figures are 0.545 and 0.543, which are only slightly lower than those for the low-quality borrowers. In other words, the marginal contribution to inference from soft/nonstandard sources remains basically stable across different credit categories, while that of hard information clearly drops in lower credit categories. As a result, the marginal contribution of soft/nonstandard information for inference is *relatively* more important for lower-quality borrowers than for higher-quality borrowers.

In row (2), we measure the marginal inference using R-squareds. We find similar results to those obtained using AUC measures. Inference is important for all three sources of information, but soft/nonstandard information is relatively more important for lower-quality borrowers.

In panel B, we report the marginal inference using R-squareds for fraction repaid as the outcome variable. Recall that we cannot use AUC (our preferred approach) to compare inference when using fraction repaid, as it is not a binary variable. As in panel A, we find that all three sources of information are important for inference. Moreover, as before, the marginal contribution to inference from standard financial variables is smaller for the low credit categories (an adjusted R^2 of 0.014 for the low versus 0.018 for the high credit categories) while, if anything, inference from the uncoded listing content is *higher* in the low credit categories (an adjusted R^2 of 0.006 for the low versus 0.004 for the high credit categories). However, the results are somewhat different for the soft/nonstandard variables since, here, unlike what we found using default as an outcome, the inference from soft/nonstandard variables drops in the low credit category (an adjusted R^2 of 0.008 in the low versus 0.011 in the high credit categories).

The broad message from these results, though, is essentially the same. Both standard financial and soft/nonstandard information contribute on the margin towards inference (and

therefore add value above each other). In addition, we find that standard financial variables are less predictive of loan performance for lower-quality borrowers than for higher-quality borrowers, both absolutely and relative to soft/nonstandard information. In sum, these results highlight the value of complementing hard information with soft/nonstandard information, especially when lenders screen low-quality borrowers.

C. Quantifying Inference along the Credit-Score Dimension

The previous results highlight the ability of lenders to infer borrower creditworthiness along dimensions *not* captured by the credit score. In particular, we show that inference from soft/nonstandard information is an important reason why lenders are able to set interest rates that are more predictive of borrower quality than the credit score, which is based entirely on hard information (e.g., past repayment history). However, we also want to understand lenders' ability to infer the information content of the credit score itself, given that the credit score is considered the standard by banks in terms of aggregating hard information for the purposes of predicting default. Here we exploit the fact that lenders observe the credit score only in categories but do not observe its exact value. Within each credit-category bin, we estimate the degree to which lenders are able to infer differences in creditworthiness across borrowers along the dimension of creditworthiness that is directly captured by credit score. A further benefit of this methodology is that it provides a precise decomposition of inference that arises from different sources of information, rather than lower bounds as presented in the previous sections.

Estimating Overall Inference Magnitude

Before turning to our regressions, we present the empirical analogue to Figure 2 (and Appendix C. Figure C.1). In Figure 6, we plot raw market interest rates against credit score. As is clear from the figure, the average interest rate declines by about 18 percentage points as we move from an average interest rate of about 26% at the lowest credit scores to an average interest rate of about 8% at the highest credit scores. Importantly, the figure shows that the interest rate also declines with credit score *within* credit categories, suggesting that lenders are able to infer creditworthiness along the credit-score dimension from other listing information. In addition, there are discrete jumps in the interest rate at the credit-category boundaries, which show that lenders do not perfectly infer the full information content of the credit score.

To test the significance of the decline in interest rates within credit categories, we first run a simple OLS regression of the market interest rate on credit score/40 and credit category (measured as a variable that is 1 for category HR, 2 for category E, ... , and 7 for category AA). Column (1) of Table 6 presents this regression. The coefficient on the credit score/40 shows that the interest rate falls by 0.55 percentage points within the typical credit category, which has a width of 40 points in the credit score. This decline is highly statistically significant and confirms the intuition from the figure that lenders are able to infer variation in creditworthiness within credit categories from other information in the listing. The coefficient on the credit category shows that the interest rate falls by a statistically significant 2.17 percentage points at the typical credit-category border. Of the 18.3 percentage point fall in the interest rate from the lowest to the highest credit score, 13.1 percentage points ($= 6 \times 2.17$) occurs at the category borders, and the remaining 5.2 percentage points occur within credit categories. Hence, a first take on the magnitude of inference would be that lenders are able to infer $5.2/18.3 = 28\%$ of the variation in creditworthiness (along the dimension of credit score) from other listing information.

There are two reasons why the analysis so far is only suggestive. First, the regression in column (1) has a rigid functional form that imposes a constant slope of interest rate with respect to credit score and a constant size of jumps in interest rate at the credit-category boundaries. To relax these functional form restrictions, we will estimate the more flexible model as specified in equation (B.1) in Appendix B. Second, and more fundamentally, the market interest rate is a censored variable: it is only observed when the interest rate at which lenders are willing to lend is lower than the maximum interest rate that the borrower has specified. Hence, the market interest rate could mechanically fall within a credit category if borrowers with higher credit scores within a credit category specify lower borrower maximum rates and if the rate at which lenders are willing to lend has a random component. Such a decline would reflect borrower behavior rather than lender inference. To capture only lender behavior, we need to estimate how the offer rate – i.e., the uncensored interest rate at which lenders are willing to lend – varies with credit score within credit categories. If the loan occurs, the market rate is equal to the offer rate. If the listing remains unfunded, we infer that the offer rate exceeds the borrower maximum rate. To properly take this censoring issue into account, we will estimate the regression as a censored regression, in which the censoring takes place at the listing-specific borrower maximum rate.

Column (2) of Table 6 implements our preferred approach (equation (B.1) in the methodology section of Appendix B) and directly estimates the extent of inference that takes place.

While we allow for a flexible form that estimates inference separately for each credit category, we focus on the combined gamma, as discussed in the methodology section. The results show that, on average, lenders are able to infer a third (0.33) of the difference in creditworthiness (along the dimension measured by credit score) between the most creditworthy and the least creditworthy borrowers within a given credit category.

We consider this to be a considerable achievement on the part of Prosper lenders for the following reason. A benchmark for our inference estimate of 0.33 is the amount of inference regarding the credit score that could have been attained if lenders had optimally used all coded information from the listing content. This benchmark, estimated using the method described in Appendix B, is 0.42. Because the benchmark is only for coded listing content and our inference estimate may also be partly based on uncoded listing content, a fairer comparison is to relate the inference based solely on coded content with our benchmark. As we will see later (Table 7), we estimate an inference of 0.04 from uncoded content, so $0.33 - 0.04 = 0.29$ is inferred from coded content. Thus, lenders were able to infer $0.29 / 0.42 = 69\%$ of what was attainable given the information provided on the Prosper website. This strikes us as a significant achievement: lenders not only infer creditworthiness beyond what is captured by credit score, but are also able to infer more than two-thirds of the information available that is directly captured by the credit score.

To understand the economic significance of this inference result, note that the α s and β s sum to 39 percentage points. In other words, the mean offered interest rate falls by 39 percentage points as we go from the lowest credit score (520) to the highest (900), which corresponds to a 411 basis-point decline ($= 3900 \times 40 / (900 - 520)$) for a typical 40-point credit category.²² The inference estimate of $\gamma = 0.330$ means that lenders infer about a third of the 411 basis-point decline in the offer rate from information other than credit category. This finding implies that lenders are (correctly) willing to offer an interest rate that is 137 ($= 0.330 \times 411$) basis points lower to the borrowers with the highest credit score within a credit category relative to the borrowers with the lowest credit score in that same category, despite being unable to observe the exact credit score. The (partial) inference is thus economically meaningful. Even if we relax our assumption that lenders exclusively try to maximize their returns by allowing for charitable or other motives, gamma remains an unbiased estimate of inference as long as charitable motives do not vary systematically with credit score within

²² This decline in the offer rate is greater than the decline in the market interest rate because censoring is much more severe in the lowest credit categories than in the highest credit categories. In particular, only 1.8% of listings are funded in the lowest credit category, while 30.9% of listings are funded in the highest credit category.

a credit-category bin. If charitable motives systematically decrease with credit score within a bin, the interest rate would fall less within a bin than would be the case absent charitable motives, and gamma would be an underestimate of true inference.

While we focus on the combined gamma, we should note that there is considerable variation in the gammas measuring inference within each credit category and that one can reject that they are all equal. The results from column (2) show that all but one of the category-specific gammas are positive and that six of the seven gammas are statistically significant at the ten-percent level or better. The inference is the largest (0.45) for the highest credit category. However, we caution against making too much of the comparisons between the separate gammas for each category, since each individual estimate is not precisely estimated given the smaller sample sizes that one necessarily faces within each credit category. Our preferred approach is therefore to compare high and low credit categories by grouping individual ones, and we will do so later.

The fact that inference is incomplete ($\gamma < 1$) implies that borrowers just below a category boundary pay a significantly higher interest rate than borrowers just above the boundary. One might therefore expect that Prosper disproportionately attracts listings by individuals with credit scores in the lower ranges of each category. Freedman and Jin (2010) present evidence consistent with such adverse selection. Adverse selection, however, does not bias our estimates since we observe exact credit score and our estimator does not depend on the density of observations by credit score within a category.

An Exact Decomposition of Inference

Table 7 presents the results of our decomposition of inference along the dimension of creditworthiness captured by the credit score. The first column presents the results from a single regression (equation B.4 in the methodology section in Appendix B) that decomposes the total combined gamma into components that are explained by specific variables in the listing. The last column presents the p-value from a test of whether the combined gamma is equal across the low and high categories.

We start by presenting analogous results from our baseline specification in Table 6 (column 2). As before, the total combined gamma is 0.33.²³ We find that the gamma for the lower credit

²³ In the first line of Table 7, we report the sum of all the components of γ . As noted in the methodology section in Appendix B, the decomposition of gamma into its components only holds in expectation in the case of a censored

categories is 0.24, while the gamma for the higher credit categories is 0.42. An F-test rejects equality of estimates between the high and the low credit categories, suggesting that there is differential inference across credit categories. These results confirm the results from Table 3, where we also found that lenders were better able to infer borrower creditworthiness for high-quality borrowers than for low-quality borrowers. The next rows present the contributions that the standard and nonstandard/soft variables make to the total combined gamma. We report both the aggregate gammas for these sub-categories and the gammas for the variables within each sub-category that show the largest (in magnitude) inference. Appendix C Table C.1 presents the individual gammas for all the variables separately. Under our assumption that lenders seek to maximize returns, the gamma for a single variable is a valid estimate of inference drawn from that variable along the credit-score dimension. If, however, a variable affects charitable motives, then the gamma on that variable will overestimate inference if the correlation between that variable and charitable motives is of the same sign as the correlation between that variable and the credit score, and underestimate inference otherwise.

The first column in Table 7 shows that, for all credit categories combined, standard financial (i.e., hard) variables contribute to most of the inference. It is not surprising that less inference is drawn from the nonstandard/soft variables, as we measure the contribution of each type of information towards inference of creditworthiness only along the dimension of credit score, which, after all, is a summary statistic of hard information. As shown in previous sections, soft/nonstandard information is an important contributor to inference of aspects of creditworthiness that are not captured by the credit score. Moreover, as we discuss below, soft/nonstandard information is relatively more important for inference about the credit score for the sample of lower-quality borrowers. In fact, close to half of the inference in lower credit categories comes from *uncoded* listing content.

Consistent with our earlier results, we take away four main points from the decomposition of the total gamma and the comparison of this decomposition between high and low credit categories.

First, along the credit-score dimension, lenders learn more from standard banking variables, which are financial and “hard,” than from variables that are voluntarily posted by borrowers. This is not unexpected since one would, *ex ante*, think that the former are not only more directly related to a

regression. As a result, the estimate of the sum of the components, 0.328 from equations (B.4) and (B.5), is close but not identical to the direct estimate of gamma, 0.330 from equation (B.1), that we presented in Table 6.

borrower's creditworthiness, but are also verified and therefore less subject to the possible "cheap talk" concerns of voluntarily posted and unverified information.

Second, we find that, among the set of standard financial variables, most of the inference is indeed drawn from variables that traditionally proxy for the likelihood of borrower distress. The number of current delinquencies, the number of credit inquiries in the last six months, the amount delinquent, and the debt-to-income ratio are variables that have high inference content. Examining whether the inference from these variables is similar across the low and high credit categories, we find that the inference for the number of current delinquencies, amount delinquent, and number of credit inquiries in the last six months is greater in the lower credit categories. However, for the debt-to-income ratio, there is greater relative inference in the higher credit categories.²⁴

Third, as before, inference from soft/nonstandard variables is relatively more important for lower credit categories (14% of overall inference) than for higher credit categories. Among the coded soft/nonstandard variables, inference content is highest for the borrower maximum rate (the maximum interest rate the borrower is willing to pay to get the loan funded). In fact, it is the second most important inference variable among all forty (including standard financial variables) that we examined. The average inference from the borrower maximum rate is 0.064 (or 19% of total inference) across all credit categories and is greater for lower credit categories (33.9%) than for higher credit categories (10.2%). It is not surprising that the borrower maximum rate accounts for more inference than other types of soft/nonstandard information, as the borrower maximum rate can serve as a credible signal of creditworthiness. As one would expect, we find that borrowers who choose a lower borrower maximum rate have a lower probability of their listing being funded, even conditional on credit score (results not reported). Since more creditworthy borrowers likely have better "outside" borrowing options, it is less costly for them, relative to less creditworthy borrowers, to post a lower borrower maximum rate. While establishing this as a separating equilibrium requires

²⁴ Consider why the magnitude of the inference changes for a given variable across the high and low credit categories. Each variable's contribution to inference can be thought of as the product of two coefficients – the (partial) coefficient from a regression of interest rate on the variable (that reflects how lenders value this variable) and the coefficient from a regression of the variable on credit score (that reflects how borrower attributes/choices are related to their credit score). Thus, inference may increase for a variable across credit categories if either (or both) of the coefficients increases. For example, in the case of current delinquencies, an examination of these coefficients shows that the large magnitude in lower categories is primarily driven by the fact that credit score is more strongly (negatively) associated with current delinquencies in the lower credit categories. Conversely, debt-to-income accounts for a greater fraction of inference in the higher credit categories because the partial coefficient from a regression of interest rate on debt-to-income is greater in magnitude in higher credit categories. This reflects the fact that lenders place more weight on debt-to-income as credit score increases.

further assumptions that we do not have the data to test for, it does strongly suggest that such a single crossing property may in fact be generated in equilibrium.

The fourth main finding from Table 7 concerns the importance of inference from uncoded information, which is soft/nonstandard by definition. While the gamma on uncoded listing content is insignificant for the whole sample, we estimate a statistically significant gamma of 0.096 (39% of total inference) from uncoded sources in the lower credit categories. Thus, in the lower credit categories, lenders draw inferences from subjective listing content that we cannot readily code. This result is consistent with our earlier finding from Table 5 that the marginal contribution to inference from uncoded listing content is most important for the lower credit categories. We also find similar results (not reported) when we estimate inference from uncoded listing content for sub-samples where we expect softer information to be of more importance: listings with images; listings where the borrower has at least one delinquency recorded in the last seven years; and listings where the number of characters in the listing text exceeds 900 (median among funded loans). In addition, we find similar results when we estimate the inference from uncoded listing content using specifications in which we use linear controls or cubic controls for all of our x variables (results not reported), suggesting that this estimate is robust to the functional form of the control variables. Thus, we conclude that uncoded listing content plays an important role for inference for weaker borrowers.

We finally note that, as evidenced by several variables that show negative inference, lenders do not always draw the correct inference from the information in the listing. For some variables, like amount requested, this negative inference likely reflects inference along other dimensions of creditworthiness not captured by credit score. This would be the case if, for a given credit score, larger loan amounts increase default likelihood.²⁵ For other variables such as posting a picture, writing more, or being endorsed by more friends, this negative inference may be indicative of mistakes lenders make (to the extent that we believe lenders are driven by profit motives). An

²⁵ Amount requested displays large negative inference in lower categories but large positive inference in higher categories. While we would normally interpret negative inference as reflecting systematic lender mistakes (for example, incorrectly believing that a variable representing a negative borrower attribute is positively correlated with credit score and mistakenly offering lower interest rates for higher values of that variable), in the case of amount requested, we believe that amount requested is also likely to have an impact through a non-credit-score dimension of creditworthiness. Unlike other variables, which mostly proxy for a borrower's attributes, amount requested is a feature of the loan. On the one hand, a higher amount requested likely predicts higher credit score because creditworthy individuals may believe that they can ask for larger amounts (which is generally the case in our data). On the other hand, all else equal, one expects that those who borrow more are more likely to default because they face larger repayment obligations. Thus, amount requested affects interest rates both through the credit-score dimension of creditworthiness and through the loan-size dimension.

alternative interpretation is that lenders do know that a borrower is more likely to default but still offer her a better interest rate due to charitable motives.

V. Conclusion

Our results show that lenders in peer-to-peer markets are able to effectively infer borrowers' creditworthiness using the rich information set that these markets provide. We find that the interest rate set by lenders in these markets predicts borrower default 45% more accurately than the exact credit score. We also find that lenders rely on nonstandard or soft sources of information in their screening process and that such information is relatively more important when screening borrowers of lower quality. In addition, the use of credible signals (like borrower maximum rate) that are nonstandard in banking contexts suggests that enhancing the opportunity for borrowers to post credible signals can further facilitate the screening process. While our findings are reassuring in that they show that these markets partially infer true creditworthiness and are not entirely influenced by "cheap" talk, there is the caveat that lenders in these markets sometimes appear to make incorrect inferences.

The broader issue, though, is to what extent peer-to-peer markets can complement traditional lenders, such as banks. We offer direct evidence that the market is able to infer dimensions of creditworthiness beyond credit score. While lenders do rely on standard banking variables for their inference, they also draw substantial additional inference from credible soft, subjective, or nonstandard information. The uncollateralized nature of lending and the ability of lenders to partly screen suggest that peer-to-peer markets can indeed complement existing lending models and improve access to credit, particularly for small individual borrowers who may otherwise be limited to more costly sources of finance, like payday lenders and credit-card debt. Furthermore, the effective use of soft/nonstandard information in screening borrower creditworthiness in these markets highlights the value to banks of incorporating credible, but softer or more subjective information in the screening of smaller, lower-quality borrowers.

References:

- Adams, William, Liran Einav, and Jonathan Levin, 2009, "Liquidity Constraints and Imperfect Information in Subprime Lending," *American Economic Review*, 99(1): 49-84.
- Agarwal, Sumit, and Robert Hauswald, 2010, "Distance and Private Information in Lending." *Review of Financial Studies*, 23(7): 2757-2788.
- Altonji, Joseph G., and Charles R. Pierret, 2001, "Employer Learning and Statistical Discrimination," *Quarterly Journal of Economics*, 116(1): 313-350.
- Berger, Allen N., and Gregory F. Udell, 2002, "Small Business Credit Availability and Relationship Lending: The Importance of Organizational Structure," *Economic Journal*, 112(477): F32-F53.
- Berger, Allen N., Nathan H. Miller, Mitchell A. Petersen, Raghuram G. Rajan and Jeremy C. Stein, 2002, "Does Function follow Organizational Form? Evidence from the Lending Practices of Large and Small Banks," *Journal of Financial Economics*, 76(2): 237-269.
- Biais, Bruno, Pierre Hillion, and Chester Spatt, 1999, "Price Discovery and Learning During the Preopening Period in the Paris Bourse," *Journal of Political Economy*, 107(6): 1218-1248.
- Bronfman, Corinne, Kevin McCabe, David Porter, Stephen Rassenti, and Vernon Smith, 1996, "An Experimental Examination of the Walrasian 'Tatonnement' Mechanism," *Rand Journal of Economics*, 27(4): 681-699.
- Camerer, Colin, 1998, "Can Asset Markets Be Manipulated? A Field Experiment with Racetrack Betting," *Journal of Political Economy*, 106(3): 457-82.
- Cao, Charles, Eric Ghysels, and Frank Hatheway, 2000, "Price Discovery without Trading: Evidence from Nasdaq Preopening," *Journal of Finance*, 56(3): 1339-1365.
- Crawford, Vincent, and Joel Sobel, 1982, "Strategic Information Transmission," *Econometrica*, 50(6): 1431-1451.
- Davies, Ryan J., 2003, "The Toronto Stock Exchange Preopening Session," *Journal of Financial Markets*, 6(4): 491-516.
- DeLong, Elizabeth R., David M. DeLong, and Daniel L. Clarke-Pearson, 1988, "Comparing the Areas Under Two or More Correlated Receiver Operating Characteristic Curves: A Nonparametric Approach," *Biometrics*, 44(3): 837-845.
- Duffie, Darrell, and Gustavo Manso, 2007, "Information Percolation in Large Markets," *American Economic Review Papers and Proceedings*, 97(2): 203-209.
- Duffie, Darrell, Semyon Malamud, and Gustavo Manso, 2009, "Information Percolation with Equilibrium Search Dynamics," *Econometrica*, 77(5): 1513-1574.

- Einav, Liran, Mark Jenkins, and Jonathan Levin, 2013, “The Impact of Credit Scoring on Consumer Lending,” *RAND Journal of Economics*, forthcoming.
- Farber, Henry S., and Robert Gibbons, 1996, “Learning and Wage Dynamics,” *Quarterly Journal of Economics*, 111(4): 1007-1047.
- Farrell, Joseph, and Rabin Matthew, 1996, “Cheap Talk,” *Journal of Economic Perspectives*, 10(3): 103-118.
- Forsythe, Robert, and Russell Lundholm, 1990, “Information Aggregation in an Experimental Market,” *Econometrica*, 58(2): 309–347.
- Freedman, Seth, and Ginger Z. Jin, 2010, “Learning by Doing with Asymmetric Information: Evidence from Prosper.com,” Working paper, University of Maryland.
- Grossman, Sanford, 1976, “On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information,” *Journal of Finance*, 31(2): 573-585.
- Grossman, Sanford, and Joseph Stiglitz, 1980, “On the Impossibility of Informationally Efficient Markets,” *American Economic Review*, 70(3): 393-408.
- Hanson, Robin, Ryan Oprea, and David Porter, 2006, “Information Aggregation and Manipulation in an Experimental Market,” *Journal of Economic Behavior & Organization*, 60(4): 449–459.
- Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, 2010, “Did Securitization Lead to Lax Screening? Evidence From Subprime Loans,” *Quarterly Journal of Economics*, 125(1): 307-362.
- Liberti, Jose M., and Atif R. Mian, 2009, “Estimating the Effect of Hierarchies on Information Use,” *Review of Financial Studies*, 22(10): 4057-4090.
- Lin, Mingfeng, Siva Viswanathan, and N. R. Prabhala, 2013, “Judging Borrowers by the Company they Keep: Friendship Networks and Information Asymmetry in Online Peer-to-Peer Lending,” *Management Science*, 59(1): 17-35.
- Petersen, Mitchell, 2004, “Information: Hard and Soft,” Working paper, Northwestern University.
- Plott, Charles R., and Shyam Sunder, 1988, “Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets,” *Econometrica* 56(5): 1085–1118.
- Pope, Devin G., and Justin Sydnor, 2011, “What’s in a Picture? Evidence of Discrimination from Prosper.com,” *Journal of Human Resources*, 46(1): 53-92.
- Rajan, Uday, Amit Seru, and Vikrant Vig, 2013, “The Failure of Models that Predict Failure,” *Journal of Financial Economics*, forthcoming.

- Ravina, Enrichetta, 2012, "Love & Loans: The Effect of Beauty and Personal Characteristics in Credit Markets," Working Paper, Columbia GSB.
- Rigbi, Oren, 2013, "The Effect of Usury Laws: Evidence from the Online Loan Market," *Review of Economics and Statistics*, forthcoming.
- Stiglitz, Joseph E., and Andrew M. Weiss, 1981, "Credit Rationing in Markets with Imperfect Information," *American Economic Review*, 71(3): 393-410.
- Strumpf, Koleman, and Paul Rhode, 2003, "Historical Presidential Betting Markets," *Journal of Economic Perspectives*, 18(2): 127–141.
- Theseira, Walter, 2009, "Competition to Default: Racial Discrimination in the Market for Online Peer-to-Peer Lending," Dissertation, Wharton.
- Townsend, Robert M., 1978, "Market Anticipations, Rational Expectations, and Bayesian Analysis," *International Economic Review*, 19(2): 481-94.
- Vives, Xavier, 1993, "How Fast Do Rational Agents Learn?" *Review of Economic Studies*, 60(2): 329-347.
- Vives, Xavier, 1995, "The Speed of Information Revelation in a Financial Market Mechanism," *Journal of Economic Theory*, 67(1): 178–204.
- Wolfers, Justin, and Eric Zitzewitz, 2004, "Prediction Markets," *Journal of Economic Perspectives*, 18(2): 107-126.
- Wolinsky, Asher, 1990, "Information Revelation in a Market with Pairwise Meetings," *Econometrica*, 58(1): 1-23.

Appendix A: Inference from the relationship between interest rate and ex-post defaults

This appendix formally derives what determines the relationship between interest rates and *ex-post* default. In particular, it shows that the slope of this relationship does not depend on the precision with which lenders can predict future defaults. The appendix further shows that the slope of *ex-post* defaults with respect to the interest rate becomes steeper if lenders systematically underreact to available information and that it becomes flatter if lenders base interest rates on factors unrelated to default probabilities.

To establish the determinants of the slope of *ex-post* defaults with respect to interest rates, we develop a stylized model that takes lenders to be approximately risk neutral with respect to idiosyncratic default risks over small bets (typically around \$100) in the Prosper online marketplace.²⁶ Lenders thus require an interest rate r such that their expected returns are equal to the “risk-free” rate of return, r_f :

$$(1 + r) = E[(1+r)(1-(1-\kappa)W)], \quad (\text{A.1})$$

where W is an indicator variable that equals 1 if the loan defaults and $\kappa \in [0, 1)$ denotes the fraction of principal and interest that the lender still receives in case of default. Solving equation (A.1) for the expected default rate yields the expected default rate as a function of the interest rate:

$$E[W | r] = \frac{1}{1 - \kappa} - \frac{1 + r_f}{1 - \kappa} (1 + r)^{-1}. \quad (\text{A.2})$$

This expression indicates that we should expect a linear relationship between observed default probabilities and $(1+r)^{-1}$. In particular, if lenders are rational, the regression coefficient on $(1+r)^{-1}$ in a regression with default realizations as the dependent variable can be interpreted as $-(1 + r_f)/(1-\kappa)$. Note that the relationship in (A.2) only requires that lenders are rational and maximize expected returns, but does not depend on the precision with which lenders form the expectation of the default probability. Intuitively, if lenders have very little information on which to base their expectations, the variance in $(1+r)^{-1}$ will be low. In that case, however, the variance in $E[W | r]$ will be low as well, so that the slope of the relationship between these two variables is not affected. This implies that the slope of the relationship between $E[W | r]$ and $(1+r)^{-1}$ cannot inform us about the strength of lenders’ signals of the default probability.

To show more formally that the relationship between $E[W | r]$ and $(1+r)^{-1}$ does not depend on the precision of rational lenders’ perceptions of loan default probabilities, we now model the process by which lenders assess loan default probabilities. Suppose lenders get a signal s of the true default probability of a loan. Assume that the signal equals the true (unobservable) default probability with probability π and that the signal is completely uninformative (i.e., a random draw from the distribution of true default probabilities) with probability $1-\pi$, so:

$$s = \omega \quad \text{with probability } \pi, \text{ and} \\ s \sim f(\omega) \quad \text{with probability } 1-\pi,$$

²⁶ The risk premium for *systematic* default risk is included in r_f .

where ω denotes the true, but unobservable, default probability of a loan and $f(\omega)$ denotes the population distribution of true default probabilities. Given this information structure, the true default probability conditional on observing signal s is:

$$E[W|s] = \pi s + (1-\pi) \omega_m, \quad (\text{A.3})$$

where ω_m is the unconditional mean default probability. To derive the interest rate as a function of the observed signal, we take the expectation in the no-arbitrage condition with respect to signal s :

$$(1 + r_f) = E[(1+r)(1-(1-\kappa)W) | s]. \quad (\text{A.4})$$

Solving (A.4) with respect to the interest rate yields the interest rate as a function of the signal:

$$r(s) = [1/(1 + r_f) - (\pi s + (1-\pi) \omega_m)(1-\kappa)/(1 + r_f)]^{-1} - 1. \quad (\text{A.5})$$

We invert equation (A.5) to find the signal that gave rise to the observed interest rate r :

$$s(r) = 1/(\pi(1-\kappa)) - (1 + r_f) / (\pi(1-\kappa)(1+r)) - \omega_m(1-\pi)/\pi. \quad (\text{A.6})$$

The true default probability of loans with interest rate r is the true default probability given signal s :

$$E[W|r] = E[W|s(r)] = \pi s(r) + (1-\pi) \omega_m. \quad (\text{A.7})$$

Substituting equation (A.6) into equation (A.7) yields:

$$E[W|r] = \frac{1}{1-\kappa} - \frac{1+r_f}{1-\kappa}(1+r)^{-1},$$

which is identical to equation (A.2). This formalizes the intuition that for rational lenders the relationship between $E[W|r]$ and $(1+r)^{-1}$ does not depend on π , the precision of the signal that lenders receive about a loan's default probability.

Next, we allow for two departures from rationality, and examine how these departures affect the relationship between $E[W|r]$ and $(1+r)^{-1}$. First, we allow lenders to misperceive the strength of the signal they receive. Second, we allow lenders to set interest rates for reasons unrelated to the expected return (i.e., include pure noise in the interest rates chosen). One can think of lenders misperceiving the strength of the signal as lenders only perceiving a fraction of the signal (and not realizing that the unobserved fraction is perfectly correlated with the observed fraction). We think of the signal as the information content in the listing and assume that lenders only observe a fraction τ of the signal. Now, the lender's *perceived* default expectation conditional on receiving signal s is:

$$E_L[W|s] = \pi\tau s + (1-\pi\tau) \omega_m, \quad (\text{A.8})$$

where ω_m is the unconditional mean default probability and the subscript L on the expectation operator denotes that the expectation is formed by the lenders. As before (see equation A.3), the true default expectation for a loan with signal s is:

$$E[W | s] = \pi s + (1-\pi) \omega_m.$$

To derive the interest rate as a function of the observed signal, we take the expectation in the no-arbitrage condition with respect to signal s but take into account that lenders only base their inference on the fraction τ of the signal:

$$(1 + r_f) = E_L[(1+r)(1-(1-\kappa)W) | s]. \quad (\text{A.9})$$

Substituting (A.8) into (A.9) and solving for r yields the interest rate as a function of the signal:

$$r(s) = [1/(1 + r_f) - (\pi\tau s + (1-\pi\tau) \omega_m)(1-\kappa)/(1 + r_f)]^{-1} - 1. \quad (\text{A.10})$$

We invert equation (A.10) to find the signal that gave rise to the observed interest rate r :

$$s(r) = 1/(\pi\tau(1-\kappa)) - (1 + r_f) / (\pi\tau(1-\kappa)(1+r)) - \omega_m(1-\pi\tau)/(\pi\tau). \quad (\text{A.11})$$

The true default probability of loans with interest rate r is the true default probability given signal s :

$$E[W | r] = E[W | s(r)] = \pi s(r) + (1-\pi) \omega_m. \quad (\text{A.12})$$

Substituting (A.11) into (A.12) yields:

$$E[W | r] = \frac{1}{(1-\kappa)\tau} - \frac{1-\tau}{\tau} \omega_m - \frac{1+r_f}{(1-\kappa)\tau} (1+r)^{-1}. \quad (\text{A.13})$$

Equation (A.13) implies that, as τ becomes smaller, the slope of default realizations with respect to $(1+r)^{-1}$ becomes steeper, and the relationship becomes infinitely steep (i.e., undefined) for $\tau = 0$. The intuition is that when lenders underreact to the signal ($\tau < 1$), they reduce the variance of the interest rates they set, so the data points on the X -axis get compressed. However, the true default probabilities do not get compressed since they do not depend on the weight that the lender places on the signal. With data points on the X -axis getting compressed, but those on the Y -axis unaffected, the relationship becomes steeper.

Next, we allow lenders to use information that is not related to the signal to set interest rates. Statistically, one can think of this as noise, which might be economically interpreted as motivations for lending unrelated to returns (altruism, taste-based discrimination, etc.). We model this unrelated information by the error term η , and write:

$$(1+r)^{-1} = (1+r^*)^{-1} + \eta, \quad (\text{A.14})$$

where $(1+r)^{-1}$ is the actual interest rate charged and $(1+r^*)^{-1}$ is the interest rate lenders would have charged if they behaved purely to maximize expected returns (though they may still misperceive the strength of the signal). Thus, r^* is given by equation (A.10) above, so that:

$$(1+r^*)^{-1} = 1/(1 + r_f) - (\pi\tau s + (1-\pi\tau) \omega_m)(1-\kappa)/(1 + r_f). \quad (\text{A.15})$$

Let the error term be uncorrelated with the “underlying” interest rate $(1+r^*)^{-1}$, and let the variance of η be denoted by σ_η^2 . The variance in $(1+r^*)^{-1}$ can be found by noting that the only stochastic term in (A.15) is s , and that the variance of s is equal to the variance of the true default probability, denoted by σ_ω^2 . The variance of $(1+r^*)^{-1}$ is thus equal to:

$$\sigma_{(1+r^*)^{-1}}^2 = \left(\pi\tau(1-\kappa) / (1+r_f) \right)^2 \sigma_\omega^2. \quad (\text{A.16})$$

The relationship between true default probability and $(1+r^*)^{-1}$ is still given by (A.13). If we regress the observed ex-post default rates on the actual interest rate $(1+r)^{-1}$, which is a noisy measure of the underlying interest rate $(1+r^*)^{-1}$, the coefficient on the actual interest rate will be attenuated, with degree of attenuation given by the standard formula for attenuation bias from classical measurement error. The coefficient of a regression of default probabilities on $(1+r)^{-1}$ will therefore be:

$$-\frac{1+r_f}{(1-\kappa)\tau} \cdot \left[1 - \frac{\sigma_\eta^2}{\sigma_{(1+r^*)^{-1}}^2 + \sigma_\eta^2} \right] = -\frac{1+r_f}{(1-\kappa)\tau} \cdot \left[1 - \frac{\sigma_\eta^2}{\left(\pi\tau(1-\kappa) / (1+r_f) \right)^2 \sigma_\omega^2 + \sigma_\eta^2} \right]. \quad (\text{A.17})$$

Expression (A.17) shows that if lenders base their interest rates partly on information unrelated to default probabilities (“noise”), the relationship between default realizations and $(1+r)^{-1}$ will become flatter. The intuition is that the noise spreads out the data points on the X-axis, leading to a flatter regression line. Thus, the two departures from rationality – (i) incomplete inference from the signal and (ii) noise in setting the interest rate – have opposite predictions on the slope of the relationship between default realizations and $(1+r)^{-1}$. Since we have no way of estimating σ_η^2 and σ_ω^2 , we cannot use our estimate of the slope to estimate τ , the degree of inference that lenders make from the signals they observe. This is why we focus on the R^2 and AUC measures instead.

Appendix B: Detailed methodology to estimate inference along the credit-score dimension

The benefit of the stylized setup shown in Figure 2 and the corresponding regression is their simplicity. However, if the true credit score were observable, the underlying relationship between interest rate and exact credit score could very well be non-linear. Moreover, credit categories are not all of equal size. Appendix Figure C.1 depicts this more realistic situation. The dashed blue line shows a possible underlying relationship between interest rate and exact credit score for the hypothetical scenario in which exact credit score is observable by lenders. This relationship is now allowed to be non-linear. As a result of this non-linearity, the slope of the observed relationship between market interest rate and credit score need not be the same within each credit category, and the jump in market interest rate at the category borders may vary. The solid red line depicts the estimated relationship between market interest rate and exact credit score. This line falls by β_k within category k and falls by α_k at the border between category $k-1$ and category k .

To determine the amount of inference, we first calculate the total fall in interest rate over each credit category. To do so, we need to decide what part of the jump of size α_k at the border between category $k-1$ and category k can be attributed to category $k-1$ and what part to category k . It appears most natural to attribute this jump proportionally to the size of each category, but results are similar when we attribute it evenly across the two bordering categories. Let λ_k denote the size of category $k-1$ as a fraction of the combined size of categories $k-1$ and k . Then, the part of the drop in interest rate at the border of categories $k-1$ and k that is attributed to category k is equal to $(1-\lambda_k)\alpha_k$. Similarly, the part of the drop at the next category border that is attributed to category k is $\lambda_{k+1}\alpha_{k+1}$. Since the interest rate falls by β_k within category k , the total drop in interest associated with category k is $\delta_k = (1-\lambda_k)\alpha_k + \lambda_{k+1}\alpha_{k+1} + \beta_k$.²⁷ The fraction of information inferred within this category, γ_k , can then be calculated as β_k / δ_k .

To estimate these parameters, we regress the interest rate on a spline in the exact credit score and cumulative dummies for the credit-score categories:

$$InterestRate_i = \mu + \sum_{k=2}^N \alpha_k I_k^{Cum}(CreditScore_i) + \sum_{k=1}^N \beta_k FracGap_k(CreditScore_i) + \varepsilon_i, \quad (B.1)$$

where $InterestRate_i$ is the interest rate charged on loan i , $CreditScore_i$ is the exact credit score of the borrower of loan i , $I_k^{Cum}(CreditScore_i)$ are cumulative credit-score dummies, and $FracGap_k$ is a variable that increases linearly with exact credit score within credit category k and that is constant everywhere else. The coefficient α_k measures the jump in interest rate at the credit-score boundary between credit categories $k-1$ and k , the coefficient β_k measures the change in interest rate within category k , and ε_i is the error term. Formally, we define $I_k^{Cum}(CreditScore_i)$ as an indicator variable that equals one if borrower i is in credit category k or higher:

²⁷ By definition, we can neither estimate a jump at the lower border of the bottom credit category nor at the upper border of the top credit category. When calculating the gammas for the first (bottom) and seventh (top) category, we assume that jumps at the lower and upper borders are of equal size: We assume that $(1-\lambda_1)\alpha_1$ equals $\lambda_2\alpha_2$ and that $\lambda_8\alpha_8$ equals $(1-\lambda_7)\alpha_7$.

$$I_k^{Cum}(CreditScore_i) = \begin{cases} 0 & \text{if } CreditScore_i < c_k \\ 1 & \text{if } CreditScore_i \geq c_k \end{cases}, \quad (\text{B.2})$$

where c_k is the credit score that forms the boundary between categories $k-1$ and k . Formally, $FracGap_k(CreditScore_i)$ is defined as:

$$FracGap_k(CreditScore_i) = \begin{cases} 0 & \text{if } CreditScore_i \leq c_k \\ \frac{CreditScore_i - c_k}{c_{k+1} - c_k} & \text{if } c_k < CreditScore_i \leq c_{k+1} \\ 1 & \text{if } c_{k+1} < CreditScore_i \end{cases}, \quad (\text{B.3})$$

Thus, $FracGap_k$ increases linearly from 0 to 1 as we move from the lowest to the highest credit score within category k . Further, $FracGap_k$ is 0 for values below c_k and equals 1 for all credit scores above c_{k+1} .

When we estimate equation (B.1), the test $\beta_k = 0$ tests the hypothesis that lenders are not able to infer variation in creditworthiness within category k (along the dimension measured by exact credit score) from all the information provided in the listing. Since the estimates of the β_k may be relatively imprecise, we also test the joint hypothesis that all β_k are equal to zero. Because the coefficients α_k measure the jumps in interest rate at the credit-score boundaries, we can reject the hypothesis that lenders are perfectly able to infer creditworthiness (along the dimension measured by exact credit score) from the information on the listing if these α s are statistically significant.

Because we estimate the γ parameters separately for each credit category, they are each based on relatively few observations. As a result, the parameters may not be estimated very precisely for particular categories, even if they are jointly significant. We therefore also present a combined γ estimate, which is the weighted average across credit-score categories of γ_k , where the weights are the precision with which the parameter is estimated in each category.

When we estimate equation (B.1), we hope to recover the effect of the listing characteristics on the interest rate that lenders require to compensate them for the perceived credit risk of that listing. If this interest rate exceeds the maximum interest rate that the borrower is willing to pay (as specified by the variable *borrower maximum rate*), the listing will not be funded, and we consequently do not observe the interest rate that lenders require. Thus, our observations of the interest rate are censored at the borrower maximum rate.²⁸ This censoring problem would bias our estimates of inference if we were to estimate equation (B.1) using ordinary least squares. Instead, we estimate equation (B.1) as a censored regression with the censoring occurring at the borrower maximum rate specified by each listing. A censored regression, which is a generalization of the Tobit model, rests on the implicit assumption that listings that were not funded would have been funded at some interest rate larger than the observed borrower maximum rate. If the error term has a homoskedastic and normal distribution, estimates from a censored regression will yield consistent estimates of the parameters determining the interest rates that lenders require to fund a listing.

To estimate the highest gamma that could have been achieved if lenders had perfectly used all hard and coded soft information available in the listing sheet, we first regress default realizations on a very flexible functional form of all standard financial variables and soft/nonstandard variables.²⁹

²⁸ State usury laws limit the maximum interest rate that borrowers may set for loans (most states allow a maximum interest rate of 36%). Thus, when state usury caps censor the market interest rate, the usury cap censors at the borrower maximum rate.

²⁹ Specifically, we use a flexible set of 215 controls for standard financial and nonstandard variables, as described in

We use this regression to predict a default probability for each loan. Because these predicted default probabilities are based on information that was observable to the lenders, lenders could also have made these predictions themselves (if they had used all coded listing content optimally). Next, we rerun our baseline specification (equation B.1) using the predicted default probability as the dependent variable instead of the interest rate. The gamma from this baseline specification with predicted default as the dependent variable tells us what fraction of predictable default along the credit score dimension occurs *within* credit categories. Hence, we consider it a benchmark for our baseline estimates because it tells us what gamma the lenders could have achieved had they adjusted their interest rate for the default risk that they could have predicted based on coded listing content.

A caveat to our methodology is that we measure inference of credit score *within* credit categories – our measurements necessarily take the size of credit categories as given. Prosper exogenously sets credit categories to represent 40-point increments in credit score. While it is not obvious whether this poses a systematic concern, it is possible that our estimates could be different if, for example, Prosper created a different credit category for every ten points in credit score instead of for every 40 points. For example, for very narrow credit-score bins, lenders may exert less effort to infer differences in credit score within bins and hence (optimally) show lower learning inference. Thus, in a sense, the credit category needs to be large enough for inferring quality *within* it to be economically meaningful. In our case, there is substantial variation in credit quality within credit categories. A forty-point change in credit score represents a large range in creditworthiness, as evidenced by the average 400-basis-point decline per credit category in the mean interest rate offered by lenders.

Decomposing Inference by Source of Information

The inference parameter γ measures the contribution of all sources of information on the Prosper website, whether or not this information can be coded as a quantitative variable. To measure the contributions of various information sources, we add to regression (B.1) controls for all the coded content of the listing:

$$InterestRate_i = \mu + \sum_{k=2}^N \alpha_k I_k^{Cum}(CreditScore_i) + \sum_{k=1}^N \beta_k^{Resid} FracGap_k(CreditScore_i) + \sum_{m=1}^M x_i^m \varphi^m + \varepsilon_i, \quad (B.4)$$

where x_i^m denotes the m^{th} variable for the coded listing content of borrower i , and φ^m denotes the corresponding regression coefficient.³⁰ In regression (B.4), the fitted interest rate can change with credit score within a credit category for two reasons. First, even after controlling for all the observable characteristics, there still may be a residual correlation between exact credit score and interest rate within a credit category due to inference from uncoded listing content. Since we measure exact credit score within credit categories by *FracGap*, this residual correlation is measured by β_k^{Resid} . Second, the fitted interest rate may vary within a credit category because (i) listings with higher values of *FracGap* may have different observable characteristics and (ii) the interest rate responds to these characteristics. We measure component (i) – the degree to which observable

Table 1, which are further interacted with the seven credit category dummies.

³⁰ In all specifications, we define the x variables to be specific within credit categories, which means that we estimate the φ coefficients for the control variables separately by credit category. We correct the α coefficients for any jumps in the interest rate at credit-category boundaries that are absorbed by the interactions of x and the credit categories or for jumps in the x variables themselves. This correction ensures that the α coefficients fully capture the jumps in the interest rate at the category boundaries.

characteristic x^m varies with $FracGap$ – by running a regression of the observations of x^m within category k on $FracGap_k$ and a constant term. We denote the coefficient on $FracGap_k$ in this bivariate regression by θ_k^m . We measure component (ii) – the degree to which the interest rate responds to characteristic x^m – by the regression coefficient φ^m . The total contribution of variable x^m to the relationship between $FracGap$ and interest rate within category k is given by the product of these two components: $\beta_k^m \equiv \theta_k^m \varphi^m$.

We decompose our original estimate β_k from the regression without the controls for coded listing characteristics (regression B.1) as follows:³¹

$$\beta_k = \beta_k^{Resid} + \sum_{m=1}^M \theta_k^m \varphi^m \equiv \beta_k^{Resid} + \sum_{m=1}^M \beta_k^m. \quad (\text{B.5})$$

In equation (B.5), $\sum_{m=1}^M \beta_k^m$ is the part of the within-category drop in interest rates that can be attributed to coded information, while the remainder is explained by uncoded information. Thus, rather than attempting to code the qualitative information (quantification of which, by definition, will be highly imperfect), we infer its information content from β_k^{Resid} , which measures the extent to which the interest rate varies with exact credit score within credit-score categories after controlling for all coded information. To ensure that β_k^{Resid} reflects uncoded information, rather than omitted higher-order terms of the x variables, we include all x variables as quadratics and interact them with credit-category indicators.³² Instead of reporting each single β_k^m , we report a sum of the β s that correspond to standard financial variables and a sum of the β s that correspond to soft/nonstandard variables. We also include β_k^{Resid} , which measures the contribution of uncoded information, with the soft/nonstandard variables. Finally, the corresponding inference parameters, γ_k^m , are calculated by dividing each type of β_k by δ_k .

We should note that this decomposition is accurate, provided that listing characteristic x^m affects interest rates only through the aspect of creditworthiness captured by credit score. Alternately, φ^m may capture an effect of x^m on the interest rate that is mediated both through the credit-score dimension and another dimension of creditworthiness. In that case, we would ascribe less (more) inference to x^m if it has a similar (opposite) impact on this other dimension of creditworthiness (compared to the credit-score dimension).

³¹ This is an application of the standard omitted variable bias formula. The omitted variable bias formula holds by construction if the equation is estimated by OLS. However, because we estimate our model as a censored regression, the omitted variable bias decomposition holds only in expectation. As a result, our decomposition will not add up exactly.

³² In addition, we include dummy variables for each of the following variables taking on a value of zero: Number of Current Delinquencies, Number of Delinquencies in Last 7 Years, Number of Public Record Requests in Last 10 Years, Number of Public Records in Last 12 Months, and Revolving Credit Balance. Amount Delinquent and Revolving Credit Balance are not introduced as quadratics, but as logs with dummies for values equal to zero and values less than or equal to 100. Each of these variables (including dummies and higher-order terms) is interacted with a full set of credit-category indicators.

Table 1: Summary Statistics

	All Listings		Funded Listings	
	Mean	S.D.	Mean	S.D.
Loan Outcomes				
Annual Lender Interest Rate			0.166	0.068
Default Dummy			0.306	
<i>Credit Category HR</i>			0.516	
<i>Credit Category E</i>			0.424	
<i>Credit Category D</i>			0.356	
<i>Credit Category C</i>			0.318	
<i>Credit Category B</i>			0.305	
<i>Credit Category A</i>			0.234	
<i>Credit Category AA</i>			0.147	
Fraction of Loan Repaid			0.797	0.334
<i>Credit Category HR</i>			0.625	0.406
<i>Credit Category E</i>			0.708	0.377
<i>Credit Category D</i>			0.762	0.352
<i>Credit Category C</i>			0.793	0.334
<i>Credit Category B</i>			0.798	0.329
<i>Credit Category A</i>			0.852	0.292
<i>Credit Category AA</i>			0.910	0.235
Credit Scores and Categories				
Credit Score	609.5	73.8	676.0	74.5
Credit Category Dummies				
<i>Credit Category HR</i>	0.343		0.068	
<i>Credit Category E</i>	0.164		0.074	
<i>Credit Category D</i>	0.178		0.173	
<i>Credit Category C</i>	0.136		0.211	
<i>Credit Category B</i>	0.082		0.183	
<i>Credit Category A</i>	0.055		0.140	
<i>Credit Category AA</i>	0.044		0.152	
Standard Financial Variables				
Amount Requested (\$)	8015	6577	6761	5788
Number of Current Delinquencies	2.89	4.54	0.77	2.28
Number of Delinquencies, Last 7 Years	9.68	15.78	4.30	10.52
Number of Public Record Requests, Last 10 Years	0.57	1.20	0.33	0.83
Total Number of Credit Lines	25.61	14.57	24.30	14.29
Number of Credit Score Inquiries, Last 6 Months	3.71	4.45	2.38	3.35
Amount Delinquent (\$)	3191	12662	855	4504
Bank Card Utilization (total balances/total limits)	0.63	0.42	0.54	0.37
Number of Public Records, Last 12 Months	0.07	0.34	0.03	0.22
Number of Current Credit Lines	8.52	6.08	9.70	5.89
Number of Open Credit Lines	7.51	5.41	8.34	5.22
Revolving Credit Balance (\$)	13446	33874	16773	38030
Debt-to-Income Ratio	0.54	1.37	0.33	0.90
Fraction Homeowners	0.37		0.48	
Credit History Age (years)	13.3	7.1	13.4	7.2
Length of Current Employment Status (months)	20.91	51.90	22.73	53.52
Personal Annual Income Dummies				
<i>N/A or Unable to Verify</i>	0.053		0.025	
<i>Not Employed</i>	0.021		0.007	
<i>\$1- \$24,999</i>	0.163		0.120	
<i>\$25,000 - \$49,999</i>	0.402		0.372	
<i>\$50,000 - \$74,999</i>	0.211		0.253	
<i>\$75,000 - \$99,999</i>	0.078		0.117	
<i>\$100,000+</i>	0.064		0.101	

Table 1 - Continued: Summary Statistics

	All Listings		Funded Listings	
	Mean	S.D.	Mean	S.D.
Employment Status Dummies				
<i>Full-Time</i>	0.812		0.859	
<i>Part-Time</i>	0.041		0.040	
<i>Self-Employed</i>	0.096		0.074	
<i>Retired</i>	0.028		0.020	
<i>Not Employed</i>	0.023		0.008	
Soft/Nonstandard Variables				
Borrower Maximum Interest Rate	0.21	0.09	0.21	0.08
Duration of Loan Listing Dummies				
<i>3 Days</i>	0.044		0.037	
<i>5 Days</i>	0.046		0.055	
<i>7 Days</i>	0.693		0.661	
<i>10 Days</i>	0.218		0.247	
Listing Category Dummies				
<i>Not Available</i>	0.386		0.380	
<i>Debt Consolidation</i>	0.281		0.262	
<i>Home Improvement Loan</i>	0.024		0.033	
<i>Business Loan</i>	0.098		0.100	
<i>Personal Loan</i>	0.114		0.121	
<i>Student Loan</i>	0.025		0.024	
<i>Auto Loan</i>	0.017		0.017	
<i>Other</i>	0.056		0.063	
Bank Draft Annual Fee Dummy	0.010		0.007	
Borrower Lists City of Residence Dummy	0.11		0.16	
Borrower Provides Image Dummy	0.54		0.69	
Characteristics of Listing Text				
<i>HTML Character Number</i>	283	271	309	350
<i>Text Character Number</i>	963	716	1106	806
<i>Average Word Length</i>	4.63	0.58	4.59	0.55
<i>Average Sentence Length</i>	122.75	97.14	106.96	68.62
<i>Number of Numerics</i>	13.03	11.31	14.49	14.32
<i>Percent of Words Misspelled</i>	0.03%	0.03%	0.03%	0.04%
<i>Number of Dollar Signs</i>	8.98	5.78	8.49	7.25
<i>Percent of Listing as Signs</i>	0.23%	0.88%	0.46%	1.26%
Number of Characters in Listing Title	30.76	13.74	32.36	13.54
Member of Group Dummy	0.18		0.30	
Group Leader Reward Rate Dummies				
<i>0%</i>	0.916		0.867	
<i>0.25%</i>	0.002		0.010	
<i>0.50%</i>	0.015		0.046	
<i>0.75%</i>	0.001		0.002	
<i>1.00%</i>	0.034		0.047	
<i>1.50%</i>	0.004		0.007	
<i>2.00%</i>	0.019		0.017	
<i>3.00%</i>	0.006		0.003	
<i>4.00%</i>	0.003		0.001	
Number of Friend Endorsements	0.324	0.769	0.519	0.973
Observations	194,033		17,212	

For brevity, we do not summarize 66 occupation and 52 state dummies (including District of Columbia and Puerto Rico). These are included as controls in the relevant specifications in Tables 5 and 7, and the Appendix tables. *Default Dummy* equals one if the loan is three or more months late at the end of the three-year loan term. *Percent of Listings as Signs* refers to the percentage of the listing composed of non-alpha-numeric characters. *HTML Character Number* refers to the number of html formatting characters in the listing and reflects the extent to which borrowers formatted their listings. *Public Records* includes bankruptcies, judgments, tax liens, court records, and overdue child support. *Bank Draft Annual Fee Dummy* equals one if the borrower elected to pay a 1% annual fee for not using the electronic funds transfer option.

Table 2: Interest Rates and Loan Performance

Panel A: OLS - Do Interest Rates Predict Loan Performance?	Default - 3 or more months late	Fraction Repaid
1/(1 + Interest Rate)	-1.525 *** (0.038)	1.173 *** (0.029)
N	17212	17212
Adjusted R ²	0.0814	0.0914
Panel B: OLS - Do Credit Scores Predict Loan Performance?	Default - 3 or more months late	Fraction Repaid
Exact Credit Score/100	-0.129 *** (0.004)	0.096 *** (0.003)
N	17212	17212
Adjusted R ²	0.0432	0.0456
Adjusted R ² using 7-part spline in credit score	0.0442	0.0476
Panel C: IV - Do Interest Rates Causally Affect Loan Performance?	Default - 3 or more months late	Fraction Repaid
1/(1 + Interest Rate)	0.166 (0.415)	0.061 (0.298)
N	17212	17212
First Stage F-stat	58.45	58.45

Default is a dummy for whether the loan is 3 or more months late as of three years after the loan is initiated (all loans have three-year maturities). Fraction Repaid measures the fraction of the principal that has been repaid after three years, not including missed interest payments. Formally, fraction repaid is defined as (principal - netdefaults)/principal where netdefaults is the principal balance minus loan sale proceeds and forfeited group rewards. Specifications in Panel A regress these measures of loan performance on $1/(1+r)$, where r is the 3-year interest rate that lenders receive on the loan. Panel B regresses loan performance measures on exact credit score. We also report the r-squared from a second specification allowing for a 7-part spline in exact credit score. Panel C examines whether interest rate has a causal effect on loan performance using credit-category dummies as the excluded instruments. We report the second-stage results from a two-stage least squares regression of default or fraction repaid on $1/(1+r)$ with controls for a spline in credit score (kinks in the spine are set at credit category boundaries). $1/(1+r)$ is instrumented with credit category dummies (the excluded instruments). Standard errors are allowed to be clustered by borrower (some borrowers hold more than one loan) and are in parentheses with * significant at 10%; ** significant at 5%; and *** significant at 1%.

Table 3: Variation in Inference by Borrower Quality

	Default (1)	Fraction Repaid (2)
Panel A: Low Credit Categories (HR - C)		
1/(1 + Interest Rate)	-1.220 *** (0.067)	1.045 *** (0.051)
N	9041	9041
Adjusted R ²	0.0349	0.0460
Panel B: High Credit Categories (B - AA)		
1/(1 + Interest Rate)	-2.201 *** (0.069)	1.602 *** (0.052)
N	8171	8171
Adjusted R ²	0.1124	0.1229
Panel C: Low Credit Categories (HR - C)		
Exact Credit Score/100	-0.157 *** (0.012)	0.132 *** (0.009)
N	9041	9041
Adjusted R ²	0.0191	0.0243
Adjusted R ² using 7-part spline in credit score	0.0192	0.0249
Panel D: High Credit Categories (B - AA)		
Exact Credit Score/100	-0.164 *** (0.010)	0.113 *** (0.007)
N	8171	8171
Adjusted R ²	0.0282	0.0277
Adjusted R ² using 7-part spline in credit score	0.0282	0.0278

Panels A and B estimate the relationship between $1/(1+r)$ and loan performance separately for the sample of low-quality borrowers (credit categories HR - C) and high-quality borrowers (credit categories B - AA). Panels C and D estimate the relationship between exact credit score and loan performance separately for low- and high-quality borrowers. All variables are as defined in Table 2. Standard errors are allowed to be clustered by borrower (some borrowers hold more than one loan) and are in parentheses with * significant at 10%; ** significant at 5%; and *** significant at 1%.

Table 4: Robustness of Measure of Inference

Estimation Sample		AUC using Interest Rate	AUC using Credit Score	P-value Test of Equality	N
(1)	Credit category HR	0.578 (0.017)	0.514 (0.017)	0.009	1169
(2)	Credit category E	0.548 (0.016)	0.517 (0.016)	0.167	1274
(3)	Credit category D	0.592 (0.011)	0.534 (0.011)	0.000	2971
(4)	Credit category C	0.602 (0.010)	0.515 (0.010)	0.000	3627
(5)	Credit category B	0.659 (0.010)	0.535 (0.011)	0.000	3149
(6)	Credit category A	0.734 (0.012)	0.511 (0.014)	0.000	2414
(7)	Credit category AA	0.780 (0.012)	0.591 (0.015)	0.000	2608
(8)	Baseline sample (All listings 2/12/2007 - 10/16/2008)	0.682 (0.004)	0.625 (0.005)	0.000	17212
(9)	Period without question and answers (Pre 2/12/2007)	0.743 (0.018)	0.687 (0.020)	0.000	767
(10)	Period before suggested borrower maximum rate and portfolio plans (Pre 10/30/2007)	0.728 (0.007)	0.663 (0.007)	0.000	5852
(11)	Period after suggested borrower maximum rate and portfolio plans (Post 10/30/2007)	0.664 (0.005)	0.602 (0.006)	0.000	11360
(12)	Period without state usury law restrictions on interest rates (Post 4/15/2008, excl. TX and SD)	0.658 (0.007)	0.600 (0.008)	0.000	6420
(13)	Sample restricted to listings with no group affiliation	0.676 (0.005)	0.608 (0.006)	0.000	12061
(14)	Sample restricted to listings posted by borrowers with no previous Prosper listings	0.695 (0.007)	0.615 (0.008)	0.000	6350
(15)	Sample restricted to listings posted by borrowers with no previous Prosper loans	0.682 (0.004)	0.623 (0.005)	0.000	15303

This table examines the robustness of our inference estimates from Table 2, using default as our outcome variable. All variables are as described in Table 2. Rows (1) through (7) restrict our sample to borrowers within each of the seven credit categories, while Row (8) shows estimates using the full baseline sample. Row (9) restricts our sample to the period before public and private questions were allowed between borrowers and lenders (pre February 12, 2007). This ensures that inference is measured from lender inference rather than from possible direct exchanges of credit score information between borrowers and lenders. Note that our baseline sample excludes the pre February 12, 2007 period because credit-category cutoffs changed on February 12, 2007. Rows (10) and (11) restrict our sample to the periods before and after Prosper added (a) a web application to suggest borrower maximum rates to borrowers and (b) an application allowing automatic bids on loans through lender portfolio plans (pre and post October 30, 2007). Representatives from Prosper have confirmed that Prosper does not use exact credit score in its calculations of suggested borrower maximum rate or its implementation of lender portfolio plans. Row (12) restricts our sample to the period after Prosper became exempt from most state usury laws which capped the maximum interest rate (post April 15, 2008) and excludes the two states, Texas and South Dakota, for which usury laws are still enforced. Row (13) restricts the sample to listings posted by borrowers with no group affiliations. Rows (14) and (15) restrict the sample to listings posted by borrowers with no previous Prosper listing or loan (funded listing), respectively. These tests confirm that our measurements of inference do not depend on information about the past repayment and listings history of borrowers who apply for more than one loan. Standard errors are allowed to be clustered by borrower (some borrowers apply for more than one loan) and are in parentheses with * significant at 10%; ** significant at 5%; and *** significant at 1%.

Table 5: Sources of Inference

	Low Credit Categories			High Credit Categories		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Predicting Default	Marginal Contribution to Inference by:			Marginal Contribution to Inference by:		
	standard financial variables	soft / nonstandard variables	uncoded listing content	standard financial variables	soft / nonstandard variables	uncoded listing content
AUC	0.5726	0.5421	0.5381	0.5916	0.5452	0.5345
Adjusted R ²	0.0140	0.0049	0.0038	0.0196	0.0057	0.0040
Panel B: Predicting Fraction Repaid	Marginal Contribution to Inference by:			Marginal Contribution to Inference by:		
	standard financial variables	soft / nonstandard variables	uncoded listing content	standard financial variables	soft / nonstandard variables	uncoded listing content
Adjusted R ²	0.0140	0.0077	0.0055	0.0184	0.0109	0.0042

Panel A estimates the marginal contribution to inference about borrower creditworthiness (as measured by ex-post default) from three sets of information: standard financial variables, soft/nonstandard variables (which are coded), and uncoded listing content (which reflects soft/nonstandard information that cannot be easily coded). We define the marginal contribution to inference as the extent to which inference will improve if lenders are able to observe one additional set of information, conditional on already observing the other two sets of information. For example, Column (1) measures the increase in total inference if investors are able to observe standard financial variables, conditional on already observing soft/nonstandard variables and uncoded listing content. To estimate the marginal inference from standard financial information, we estimate a first-stage regression of interest rate on the other two sets of information (in this case, soft/nonstandard variables and uncoded listing content, which is estimated as the residual of $1/(1+r)$ regressed on all coded information). We then estimate an ROC curve using the residual from this first stage estimation. The ROC curve's corresponding AUC represents the marginal inference from standard financial variables. Alternatively, we can measure marginal inference using r-squareds instead of the AUC. In a second-stage specification, we regress default on the residual from the first stage estimation. The r-squared from the second-stage estimation represents the marginal contribution to inference from standard financial variables. Panel B repeats the exercise using fraction repaid as the outcome variable. Note that because fraction repaid is a continuous variable, we present the r-squared as our measure of inference instead of the AUC. We control for coded standard and non-standard variables as quadratics, with Amount Delinquent and Revolving Credit Balance measured in log form. We also include dummy variables for each of the following variables taking on a value of zero: Number of Current Delinquencies, Number of Delinquencies in Last 7 Years, Number of Public Record Requests in Last 10 Years, Number of Public Records in Last 12 Months, Revolving Credit Balance, Amount Delinquent, and Revolving Credit Balance. We further include dummy variables for Amount Delinquent and Revolving Credit Balance less than 100 USD. All other variables are as described in Table 2.

Table 6: Inferring Creditworthiness along the Credit-score Dimension

Dependent Variable: Interest Rate	(1)		(2)	
	OLS		Censored Regression	
Estimate	Coefficient	(S.E.)	Coefficient	(S.E.)
Combined γ : Inference			0.330 ***	(0.033)
Regression Coefficients				
Credit score/40	-0.0055 ***	(0.0008)		
Credit category	-0.0218 ***	(0.0009)		
α_2 : Change between Categories HR and E			-0.038 ***	(0.005)
α_3 : Change between Categories E and D			-0.059 ***	(0.005)
α_4 : Change between Categories D and C			-0.049 ***	(0.004)
α_5 : Change between Categories C and B			-0.051 ***	(0.005)
α_6 : Change between Categories B and A			-0.031 ***	(0.005)
α_7 : Change between Categories A and AA			-0.042 ***	(0.005)
β_1 : Change within Category HR			-0.011 *	(0.006)
β_2 : Change within Category E			-0.011 *	(0.007)
β_3 : Change within Category D			-0.027 ***	(0.005)
β_4 : Change within Category C			0.000	(0.005)
β_5 : Change within Category B			-0.014 **	(0.006)
β_6 : Change within Category A			-0.005	(0.007)
β_7 : Change within Category AA			-0.052 ***	(0.008)
N	17,212		194,033	
R ²	0.492		0.431	
Implied Coefficients and Tests				
$\gamma_1 = \beta_1/\delta_1$: Inference in Credit Category HR			0.229 *	(0.120)
$\gamma_2 = \beta_2/\delta_2$: Inference in Credit Category E			0.189 *	(0.099)
$\gamma_3 = \beta_3/\delta_3$: Inference in Credit Category D			0.332 ***	(0.056)
$\gamma_4 = \beta_4/\delta_4$: Inference in Credit Category C			-0.006	(0.107)
$\gamma_5 = \beta_5/\delta_5$: Inference in Credit Category B			0.253 ***	(0.092)
$\gamma_6 = \beta_6/\delta_6$: Inference in Credit Category A			0.165	(0.192)
$\gamma_7 = \beta_7/\delta_7$: Inference in Credit Category AA			0.450 ***	(0.055)
p-value: $\gamma_i=\gamma$			0.002	
p-value: $\gamma_i=0$			0.000	


This table examines the ability of lenders to infer borrower creditworthiness along the credit-score dimension. Column (1) takes a simple approach and asks whether, conditional on the observable credit category, credit score predicts the interest rate (measured as the one-year interest rate). It estimates an OLS specification in which the sample is restricted to funded listings. Column (2) implements a more flexible specification described in Equation (B.1) in Appendix B, and estimates the extent of inference that takes place using the full baseline sample, including unfunded listings. In Column (2), all estimates are based on a censored normal regression with the interest rate as the dependent variable. Standard errors are allowed to be clustered by borrower (some borrowers apply for more than one loan) and are in parentheses with * significant at 10%; ** significant at 5%; and *** significant at 1%.

Table 7: Decomposing Inference along the Credit-score Dimension

	(1)	(2)	(3)	(4)
	All Credit Categories	Low Credit Categories (HR - C)	High Credit Categories (B - AA)	Low = High p-value
All Listing Content (γ)	0.328*** (0.027)	0.244*** (0.044)	0.417*** (0.028)	0.001
Decomposition of γ				
1. Standard Financial Variables	0.312*** (0.020)	0.210*** (0.020)	0.421*** (0.034)	0.000
1.1 Number of Current Delinquencies	0.079 *** (0.006)	0.110 *** (0.010)	0.045 *** (0.007)	0.000
1.2 Number of Credit Inquiries, Last 6 months	0.054 *** (0.003)	0.073 *** (0.004)	0.034 *** (0.003)	0.000
1.3 Amount Delinquent	0.051 *** (0.006)	0.085 *** (0.010)	0.015 *** (0.006)	0.000
1.4 Debt-to-Income Ratio	0.048 *** (0.007)	0.001 (0.008)	0.099 *** (0.011)	0.000
1.5 Amount Requested	-0.005 (0.005)	-0.124 *** (0.006)	0.122 *** (0.009)	0.000
1.6 All Other Standard Financial Variables	0.085 *** (0.016)	0.065 *** (0.017)	0.106 *** (0.028)	0.226
2. Soft/Nonstandard Information	0.016 (0.032)	0.034 (0.045)	-0.004 (0.044)	0.557
2.1 Borrower Maximum Rate	0.064 *** (0.004)	0.083 *** (0.005)	0.043 *** (0.007)	0.000
2.2 Listing Category	-0.026 *** (0.003)	-0.048 *** (0.005)	-0.002 (0.005)	0.000
2.3 Member of Group	-0.016 *** (0.002)	-0.028 *** (0.004)	-0.003 *** (0.001)	0.000
2.4 Group Leader Reward Rate	-0.015 *** (0.002)	-0.028 *** (0.004)	-0.002 (0.002)	0.000
2.5 All Other Nonstandard Variables	-0.031 *** (0.005)	-0.042 *** (0.008)	-0.019 *** (0.006)	0.025
2.6 Uncoded Listing Content	0.040 (0.032)	0.096 ** (0.045)	-0.020 (0.044)	0.066


This table decomposes our estimate of inference of creditworthiness (along the dimension captured by credit score) presented in Table 6 Column (2) into sources of inference. The decomposition is based upon the baseline censored normal specification with the addition of 216 control variables, each interacted with seven credit-category dummies, such that the coefficient on each control variable is allowed to vary by credit category. Column (1) presents results for the entire sample, while the next two columns, (2)-(3), present the combined gamma separately for the lower credit categories (C, D, E, and HR) and the higher credit categories (AA, A, and B). Column (4) presents the p-value from a test of whether the estimates for the lower and higher credit categories are equal. The top row presents our estimate of gamma. The rows below decompose the gamma in the top row into two main groups: 1. standard financial variables and 2. soft/nonstandard information. Each of these two main groups are broken down further into subgroups 1.1 - 1.6 and 2.1 - 2.6, respectively. Please refer to Appendix C Table C.1 for the full decomposition results and variable definitions. Standard errors are allowed to be clustered by borrower (some borrowers apply for more than one loan) and are in parentheses with * significant at 10%; ** significant at 5%; and *** significant at 1%.

Figure 1: Stylized Receiver Operator Curve (ROC)




This figure shows a stylized receiver operator curve which we use to summarize the predictive power of various credit screening measures, such as the credit score or the interest rate. It plots the fraction of true positives out of the positives (TPR = true positive rate) vs. the fraction of false positives out of the negatives (FPR = false positive rate), at various threshold settings. ROC curves with greater area under the curve (AUC) represent superior predictors for default. See Section III for more details.

Figure 2: Stylized Relationship between Interest Rate and Credit Score




This figure shows the stylized hypothesized relationship between a borrower's credit score and the market interest rate on her (funded) loan.

Figure 3: ROC Curves -- Full Sample




Panels (A) and (B) present the ROC curves for the interest rate and credit score, respectively. Panel (C) presents the ROC curve using the best possible score that an econometrician would construct if he used all observable standard financial variables including exact credit score. Panel (D) presents the ROC curve using the best possible score than an econometrician would construct if he used all observable standard financial variables as well as coded non-standard or softer variables. See Section III.A for a detailed description of the creation of the curves in Panels (C) and (D). We can reject equality between the interest rate AUC in Panel (A) and each of the AUCs in Panels (B), (C), and (D) with p-values of less than 0.001.

Figure 4: ROC Curves -- Low Credit Categories




This figure plots the same ROC curves as described in Figure 3, but with the sample restricted to credit categories HR-C (lower-quality borrowers). We can reject equality between the interest rate AUC in Panel (A) and each of the AUCs in Panels (B), (C), and (D) with p-values of less than 0.001.

Figure 5: ROC Curves -- High Credit Categories



This figure plots the same ROC curves as described in Figure 3, but with the sample restricted to credit categories B-AA (higher-quality borrowers). We can reject equality between the interest rate AUC in Panel (A) and each of the AUCs in Panels (B), (C), and (D) with p-values of less than 0.001.

Figure 6: Market Interest Rate and Credit Scores



This figure shows the "raw" relationship between a borrower's credit score and the one-year interest rate on her funded loan. Each point in the graph plots the average interest rate over an eight-point range in credit scores. Solid lines separate the seven credit categories. Starting from left to right, the categories are: HR, E, D, C, B, A, AA. Lenders observe the borrower's credit category but, do not observe the borrower's exact credit score.

Appendix C, Table C.1: Decomposing Inference, Part I (Standard Financial Variables)

	(1)			(2)			(3)			(4)
	All Credit Categories			Low Credit Categories (HR - C)			High Credit Categories (B - AA)			Low = High p-value
Standard Financial Variables										
No. of Current Delinquencies	0.079	(0.006)	***	0.110	(0.010)	***	0.045	(0.007)	***	0.000
No. of Credit Inquiries, Last 6 Months	0.054	(0.003)	***	0.073	(0.004)	***	0.034	(0.003)	***	0.000
Amount Delinquent	0.051	(0.006)	***	0.085	(0.010)	***	0.015	(0.006)	***	0.000
Debt-to-Income Ratio	0.048	(0.007)	***	0.001	(0.008)		0.099	(0.011)	***	0.000
Amount Requested	-0.005	(0.005)		-0.124	(0.006)	***	0.122	(0.009)	***	0.000
No. of Delinquencies, Last 7 Years	0.033	(0.004)	***	0.043	(0.006)	***	0.023	(0.005)	***	0.006
No. of Public Records, Last 10 Years	0.023	(0.002)	***	0.018	(0.004)	***	0.028	(0.003)	***	0.056
Total No. of Credit Lines	-0.004	(0.005)		-0.008	(0.009)		0.001	(0.005)		0.391
Bank Card Utilization Ratio	-0.003	(0.011)		0.008	(0.006)		-0.015	(0.021)		0.290
No. of Public Records, Last 12 Months	0.000	(0.002)		-0.001	(0.002)		0.000	(0.003)		0.896
No. of Current Credit Lines	0.004	(0.008)		0.006	(0.015)		0.002	(0.006)		0.807
No. of Open Credit Lines	-0.002	(0.008)		-0.001	(0.014)		-0.002	(0.006)		0.945
Revolving Credit Balance	-0.011	(0.007)		-0.025	(0.010)	***	0.005	(0.010)		0.028
Homeownership Dummy	0.024	(0.006)	***	0.011	(0.005)	**	0.039	(0.010)	***	0.013
Credit History Age	0.007	(0.005)		0.010	(0.007)		0.004	(0.007)		0.558
State of Residency (52 Dummies)	-0.013	(0.005)	***	-0.024	(0.007)	***	-0.002	(0.006)		0.024
Employment Status (5 Dummies)	0.002	(0.002)		0.007	(0.004)	*	-0.004	(0.001)	**	0.009
Length of Current Employment Status	-0.003	(0.001)	**	-0.005	(0.002)	**	-0.001	(0.001)		0.059
Personal Annual Income (7 Dummies)	0.014	(0.005)	***	0.012	(0.006)	**	0.016	(0.009)	*	0.711
Borrower Occupation (62 Dummies)	0.011	(0.006)	**	0.011	(0.008)		0.011	(0.007)		0.990
Missing Data (2 Dummies)	0.001	(0.002)		0.003	(0.002)		0.000	(0.003)		0.464


This table shows the decomposition of our estimate of gamma presented in Table 6, Column (2). The decomposition results are divided into standard financial variables, presented here, and soft/nonstandard variables, presented in the next page. The decomposition is based upon the baseline censored normal specification with the addition of 216 control variables, each interacted with seven credit category dummies, such that the coefficient on each control variable is allowed to vary by credit category. All controls except for dummy variables are entered as quadratics. *Amount delinquent* and *revolving credit balance* are introduced as logs with dummies for values equal to zero and values less than or equal to 100. *Missing Data* consists of two dummies equal to one when subsets of the standard financial variables are missing in the data (observations with missing standard financial variables account for less than one percent of our sample). Column (1) presents results for the entire sample, while the next two columns, (2)-(3), present the combined gamma separately for the lower credit categories (C, D, E, and HR) and the higher credit categories (AA, A, and B). Column (4) presents the p-value from a test of whether the estimates for the lower and higher credit categories are equal. Standard errors are allowed to be clustered by borrower (some borrowers apply for more than one loan) and are in brackets with * significant at 10%; ** significant at 5%; and *** significant at 1%.

Appendix C, Table C.1: Decomposing Inference, Part II (Nonstandard Variables)

	(1)			(2)			(3)			(4)
	All Credit Categories			Low Credit Categories (HR - C)			High Credit Categories (B - AA)			Low = High p-value
Non-Standard Variables										
Borrower Maximum Interest Rate	0.064	(0.004)	***	0.083	(0.005)	***	0.043	(0.007)	***	0.000
Listing Category (8 Dummies)	-0.026	(0.003)	***	-0.048	(0.005)	***	-0.002	(0.005)		0.000
Member of Group Dummy	-0.016	(0.002)	***	-0.028	(0.004)	***	-0.003	(0.001)	***	0.000
Group Leader Reward Rate (9 Dummies)	-0.015	(0.002)	***	-0.028	(0.004)	***	-0.002	(0.002)		0.000
Duration of Loan Listing (4 Dummies)	-0.011	(0.002)	***	-0.009	(0.003)	***	-0.012	(0.003)	***	0.447
Bank Draft Annual Fee Dummy	0.000	(0.001)		0.000	(0.001)		0.000	(0.001)		0.924
Borrower Lists City Dummy	-0.001	(0.001)		-0.001	(0.002)		0.000	(0.000)		0.623
Borrower Provides Image Dummy	-0.002	(0.001)	**	-0.004	(0.001)	***	0.000	(0.001)		0.044
HTML Character No.	0.000	(0.002)		-0.001	(0.002)		0.001	(0.002)		0.542
Text Character No.	-0.005	(0.002)	***	-0.006	(0.004)		-0.005	(0.001)	***	0.808
Average Word Length	0.002	(0.001)		0.004	(0.003)		-0.001	(0.001)		0.075
Average Sentence Length	-0.003	(0.001)	**	-0.007	(0.002)	***	0.002	(0.001)	**	0.001
No. of Numerics	-0.003	(0.004)		0.000	(0.003)		-0.006	(0.008)		0.510
Percent Misspelled	-0.001	(0.001)		-0.001	(0.002)		0.000	(0.001)		0.502
No. of Dollar Signs	-0.003	(0.004)		-0.003	(0.003)		-0.003	(0.008)		0.983
Percent of Listing as Signs	0.003	(0.002)	**	0.004	(0.003)		0.002	(0.001)		0.570
No. of Characters in Listing Title	-0.001	(0.001)		-0.002	(0.001)		0.000	(0.001)		0.292
No. of Friend Endorsements	-0.007	(0.002)	***	-0.016	(0.003)	***	0.003	(0.003)		0.000
Uncoded Listing Content	0.040	(0.032)		0.096	(0.045)	**	-0.020	(0.044)		0.066

This table shows the decomposition of our estimate of gamma presented in Table 6, Column (2). The decomposition results are divided into standard financial variables, presented in the previous page, and soft/nonstandard variables, presented here. The decomposition is based upon the baseline censored normal specification with the addition of 216 control variables, each interacted with seven credit category dummies, such that the coefficient on each control variable is allowed to vary by credit category. We control for coded standard and non-standard variables as quadratics, with Amount Delinquent and Revolving Credit Balance measured in log form. We also include dummy variables for each of the following variables taking on a value of zero: Number of Current Delinquencies, Number of Delinquencies in Last 7 Years, Number of Public Record Requests in Last 10 Years, Number of Public Records in Last 12 Months, Revolving Credit Balance, and Amount Delinquent. We further include dummy variables for Amount Delinquent and Revolving Credit Balance less than \$USD100. Missing Data consists of two dummies equal to one when subsets of the standard financial variables are missing in the data (observations with missing standard financial variables account for less than one percent of our sample). Column (1) presents results for the entire sample, while the next two columns, (2)-(3), present the combined gamma separately for the lower credit categories (C, D, E, and HR) and the higher credit categories (AA, A, and B). Column (4) presents the p-value from a test of whether the estimates for the lower and higher credit categories are equal. Standard errors are allowed to be clustered by borrower (some borrowers apply for more than one loan) and are in brackets with * significant at 10%; ** significant at 5%; and *** significant at 1%.

Appendix C, Figure C.1: Relationship between Interest Rate and Credit Score



This figure shows a more realistic hypothesized relationship between a borrower's credit score and the market interest rate on her (funded) loan.

Appendix D: Sample Listing

PROSPER Let's bring on each other

Home
Get a Loan
Bid on Loans
Community
My Account
Help

Search Listings
Portfolio Plans
Advanced Search
About Lending
Rates
Performance
Watch List

help me pay off credit cards and propose to my girlfriend

(Listing #000104) [Back to search results](#)

LISTING SUMMARY



\$8,881.00 @ 8.94%

Bid down from 13.92%

Bid Now
Bidding has ended

Funding: 100% funded

Bid: 321 bids
Ended
Listing became a loan

Borrower APR: 8.94%

Mo. payment: \$255.60 (by auto)

BORROWER INFO

#1402
 CAMDEN, MA 

Member and Friends of the
 Boston Area College
 Community

4-10-10-11 (13)

[2 friend bids](#)
[2 questions & answers](#)
[2 friends, 2 verified](#)
[1 loan total, 1 active](#)

FORECAST



[View](#)
[Error](#)
[Report this listing](#)

CREDIT PROFILE

A credit grade Help		10% debt to income ratio	
<small>Homeownership: not verified</small>			
Now delinquent:	0	First credit line:	May-2001
Amount delinquent:	\$0	Current / open credit lines:	3 / 8
Delinquencies in last 7y:	0	Total credit lines:	4
Public records last 12m / 18y:	0 / 0	Revolving credit balance:	\$3,331
Inquiries last 6m:	1	Bankcard utilization:	25%
<small>Credit and home ownership information provided by Experian.</small>		<small>Employment and income provided by borrower.</small>	

DESCRIPTION

Purpose of loan:
 Am using this loan to pay off my \$8,881 credit card bill currently at 14% at a lower interest rate and to buy my girlfriend, Jennifer, an engagement ring costing approx. \$4,500. Up until about 3 months ago, I would receive all my purchases through my credit card and I ended up letting it get slightly away from me. As a result, I've devised a plan to pay off as much debt as possible per month (currently, I pay \$500 to my credit card company) and live on a necessity only budget. The next phase of my plan after eliminating my credit card debt was to immediately go back and almost as much debt as I have now to buy Jennifer a ring. Then along came prosper. With your help, to be able to ask Jennifer to marry me sooner than expected and maybe not even be in debt when I do!!

My financial situation:
 Currently, I work as a software engineer in Waltham, MA. I make a pretty good living and enjoy what I do. The people I work with like and respect me and I feel my job is very secure and also portable (i.e. I can work from anywhere with an internet connection) should I need to move. Jennifer is in her 4th year of med school and is looking at residencies. Below, you can see my monthly expenses which will be going down some May/June since Jennifer and I will be moving in together. I invest in the stock market and I am also using Prosper on the lender's side. I have a little bit of cash set aside for a rainy day and a bit more available (though not as quickly attainable) in case of a financial hiccup. I also put away 12% of my gross pay into a 401k which my company contributes to with profit sharing.

Appendix: Sample Listing - Continued

Monthly net income: \$ 2004

Monthly expenses: \$ 1705

Housing: \$ 520

Insurance: \$ 200

Car expenses: \$ 120

Utilities: \$ 40

Phone, cable, internet: \$ 65

Food, entertainment: \$ 400

Clothing, household expenses: \$ 50

Credit cards and other loans: being paid with this loan

Other expenses: \$ 0

Prosper Loan: \$270

FRIENDS AND FAMILY WINNING BIDS



This member has no winning bids from friends and family

QUESTIONS & ANSWERS

This borrower has not publicly answered any questions from lenders

BID HISTORY

Legend: = In group = Friend = Winning = Partially winning = Outbid



Bidder / Relationship	Rate	Amount Bid	Winnig	Status	Bid Date (PT)
anjha29	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 8:28 AM
user13	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 9:18 AM
JLamar	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 9:11 AM
stamm2004	8.99%	\$100.00	\$100.00	✓	Oct-09-2007 8:50 AM
janda1983	8.99%	\$100.00	\$100.00	✓	Oct-09-2007 7:55 AM
Proctor2	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 7:50 AM
omroady	8.99%	\$100.00	\$100.00	✓	Oct-09-2007 7:47 AM
omroady	8.99%	\$100.00	\$100.00	✓	Oct-09-2007 7:40 AM
jocelyn	8.99%	\$75.00	\$75.00	✓	Oct-09-2007 7:35 AM
awwawake	8.99%	\$100.00	\$100.00	✓	Oct-09-2007 7:28 AM
DGS_Castel	8.99%	\$51.43	\$51.43	✓	Oct-09-2007 7:24 AM
oxos_tecnat	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 7:15 AM
Grisan2007	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 6:54 AM
ek3	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 6:50 AM
LoanCherri	8.99%	\$100.00	\$100.00	✓	Oct-09-2007 6:29 AM
Grisia11	8.99%	\$200.00	\$200.00	✓	Oct-09-2007 5:01 AM
Budget1	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 5:03 AM
Gocritings2nu	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 5:28 AM
stevens444	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 5:13 AM
Deaf_Flow	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 5:05 AM
Curtisgnat	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 4:59 AM
Eliar	8.99%	\$50.00	\$50.00	✓	Oct-09-2007 4:49 AM