Screening in New Credit Markets: Can Individual Lenders Infer Borrower Creditworthiness in Peer-to-Peer Lending?

Faculty Research Working Papers Series

Screening in New Credit Markets
Can Individual Lenders Infer Borrower Creditworthiness in Peer-to-Peer Lending?

Rajkamal Iyer
MIT Sloan School of Management

Asim Ijaz Khwaja
John F. Kennedy School of Government – Harvard University
National Bureau of Economic Research

Erzo F. P. Luttmer
John F. Kennedy School of Government – Harvard University
National Bureau of Economic Research

Kelly Shue
Department of Economics, Harvard University

September 2009
RWP09-031

Screening in New Credit Markets

Can Individual Lenders Infer Borrower Creditworthiness in Peer-to-Peer Lending?

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

August 2009

Abstract

The current banking crisis highlights the challenges faced in the traditional lending model, particularly in terms of screening smaller borrowers. The recent growth in online peer-to-peer lending marketplaces offers opportunities to examine different lending models that rely on screening by multiple peers. This paper evaluates the screening ability of lenders in such peer-to-peer markets. Our methodology takes advantage of the fact that lenders do not observe a borrower's true credit score but only see an aggregate credit category. We find that lenders are able to use available information to infer a third of the variation in creditworthiness that is captured by a borrower's credit score. This inference is economically significant and allows lenders to lend at a 140-basis-points lower rate for borrowers with (unobserved to lenders) better credit scores within a credit category. While lenders infer the most from standard banking "hard" information, they also use non-standard (subjective) information. Our methodology shows, without needing to code subjective information that lenders learn even from such "softer" information, particularly when it is likely to provide credible signals regarding borrower creditworthiness. Our findings highlight the screening ability of peer-to-peer markets and suggest that these emerging markets may provide a viable complement to traditional lending markets, especially for smaller borrowers.

JEL codes: *D53*, *D8, G21, L81*
Keywords: *Peer-to-peer credit markets, Market-based Lending, Screening, Market Inference, Information and Hierarchies.*

---

* Iyer: MIT Sloan School of Management riyer@mit.edu; Khwaja: Harvard Kennedy School and NBER akhwaja@ksg.harvard.edu; Luttmer: Harvard Kennedy School and NBER erzo_luttmer@harvard.edu; Shue: Harvard University shue@fas.harvard.edu. We are extremely grateful to Prosper.com for being so generous with their time in answering our queries. We thank Raymond Fisman, Lawrence Katz, Atif Mian, Enrichetta Ravina, David Scharfstein, Jeremy Stein, and seminar participants at Harvard and the NBER finance meetings 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 based on borrowers' creditworthiness. Traditionally, banks have played the dominant role in allocating credit partly because they are attributed to have the financial expertise to evaluate borrowers and effectively intermediate capital (Diamond, 1984). While there is a broad consensus on the importance of banks in financial intermediation, the recent banking crisis has highlighted short-comings in the traditional lending models, particularly in allocating credit to smaller borrowers.[1] While there is increasing debate in how these short-comings can be addressed, a variety of new lending models offer potentially valuable insights. 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. While such markets may be better at utilizing non-standard/"softer" information, the (peer) lenders typically lack the financial and screening expertise of traditional banks. In this paper, we evaluate whether such lending platforms are able to effectively screen for borrower creditworthiness. Thus, we examine the viability of such lending platforms in improving small borrowers' credit access, in turn complementing traditional lending models.

Web-based peer-to-peer lending markets, such as Prosper, Zopa, Kiva, Myc4, Lending Club, Pertuity Direct, and Fynanz, have grown dramatically both in number and size. Prosper has funded over $178 million in loans and currently has 830,000 members. These markets are quickly gaining popularity in lending to smaller-scale borrowers such as individuals and small firms, both in developed and developing economies.[2] The uncollateralized nature of lending in these online markets makes it particularly attractive for small borrowers who might otherwise turn to payday lenders or credit card debt, often at exorbitant rates (Adams, Einav, and Levin, 2009). However, as non-financial experts dominate peer-to-peer markets, their ability to judge financial risk and information is key to the viability of these markets. While there is some evidence from other contexts, such as prediction markets, that non-experts can extract information effectively (Wolfers

[1] Moreover, there are both theoretical arguments and empirical evidence that banks do very little screening for small borrowers and rely excessively on collateral, thereby preventing some creditworthy borrowers from obtaining loans (Stiglitz and Weiss, 1981; Ang et al., 1995; Avery et al., 1998; Manove et al., 2001).

[2] While micro credit institutions have improved financial access for small borrowers in many economies, they primarily rely on group lending principles, which can sometimes make it difficult for individual borrowers to access credit. Equity and corporate debt markets also provide financing but they are typically limited to large-scale, mature borrowers.

and Zitzewitz, 2004), there is scant direct evidence on whether these peer-to-peer markets can effectively screen borrowers and allocate credit.[3]

This paper uses unique loan-level data from an online peer-to-peer market, Prosper.com, to examine the extent to which multiple lenders can collectively infer borrowers' underlying creditworthiness by exploiting the potentially rich information setting that peer-to-peer lending websites allow. We propose a methodology that takes advantage of the fact that we as econometricians observe a borrower's exact credit score but lenders on Prosper.com only see an aggregated credit category. Thus, if lenders offer loans at lower interest rates to borrowers who have better credit scores *within* a given credit category, lenders must have correctly inferred that these borrowers are more creditworthy than others in the same credit category.[4] Our methodology quantifies lenders' inference of creditworthiness 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. Our methodology also allows us to decompose the magnitude of inference associated with different types of information available to the lender.[5]

We find that, within a given credit category, lenders are able to infer one-third of the differences in creditworthiness that are captured by a borrower's exact credit score. This is an economically significant effect because inference allows lending at a 140-basis-points lower rate for borrowers at the top of a typical credit category relative to borrowers at the bottom of that category. Our results show that lenders exhibit greater inference for borrowers in higher credit categories. In addition, this inference is mostly based on hard, verified financial information that is also normally used by banks to screen borrowers (henceforth referred to as "standard banking" variables). Within such types of information, the greatest inference is derived from variables such as a borrower's number of current delinquencies, debt-to-income ratio, amount delinquent, and the number of credit inquiries in the last six months, although there is variation in these variables' relative importance across credit categories. For example, delinquencies (amount and number) are more

[3] For example, small election markets, like the Iowa electronic markets, and event markets, like TradeSports, that rely on aggregating information from a relatively small number of (non-expert) individuals seem to provide reasonably accurate predictions.

[4] The final interest rate for a funded loan is determined through sequential bidding and reflects the lenders' overall perception of the quality and, hence, the creditworthiness of the borrower. The loan is funded only if the total bids equal or exceed the amount requested by the borrower and the final interest rate is determined by the highest reservation rate among the set of lenders that successfully bids.

[5] The borrower listing contains hard, verified information obtained from the credit rating agency (past defaults, number of credit lines, etc.) and soft, subjective, non-verified information (picture, description, etc.) that borrowers voluntarily provide. Lenders do not have access to any additional information about borrowers apart from the listing.

informative in lower credit categories while the debt-to-income ratio is more salient for the higher credit categories.

Lenders also learn from other softer, subjective (non-verified) information that is voluntarily posted by borrowers (henceforth referred to as "non-standard" variables), particularly in the lower credit categories. Of the non-standard variables, we find that lenders draw the most inference from the maximum interest rate that the borrower posts that she is willing to pay for the loan. This rate is likely to serve as a credible and costly signal since borrowers posting too low a rate risk not having the loan funded, and this signal may be costlier for lower-quality borrowers with fewer alternate funding options. Our results suggest that, as one would expect, lenders pay greater attention to the most credible signals that borrowers can send.[6] We also find, especially among the lower credit categories, a high degree of inference from the non-coded component of the listing.

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 information without explicitly coding it since this inference is computed as a "residual," i.e. the variation of interest rates with the exact credit score that remains after controlling for coded information.

A concern with the interpretation of our findings may be that 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. However, this possibility is very unlikely. Prosper strongly discourages borrowers from revealing detailed personal information (like credit score or personal contact information) and a text search through all listing text indeed does not find any self-reported credit scores. While we do not have access to the "questions and answers" data to conduct a similar check, even if such information was reported it would not be credible as every borrower has an incentive to report the highest score in her credit category. Not surprisingly, restricting our sample to the period before the introduction of the question-and-answer feature provides similar results. We also show the robustness of our results to various policy changes introduced by Prosper, differences in usury laws, and group affiliation of borrowers.

[6] 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.

An important caveat of our approach is that we use credit score as a plausible but imperfect proxy for true creditworthiness. An ideal inference test requires a measure of true creditworthiness, defined as the *ex-ante* probability that a given loan will default. Such *ex-ante* probabilities are, by definition, not observable. *Ex-post*, only a realization that depends on the true probability is observed, and only for the subsample of listings that are funded. The true ex-ante default probability depends not only on the borrower's attributes (both observed and unobserved) but also on the characteristics of the specific loan, including the amount borrowed, loan terms, and other elements of the loan. The credit score provides an estimate of the true default probability, but it is only based on a subset of predictors.[7] Despite this limitation, the credit score is the best available measure of the ex-ante default probability – credit bureaus have strong incentives to construct credit scores that are accurate predictors of future default likelihoods.[8] Nevertheless, we interpret our results with care. There may be aspects of creditworthiness that are not fully captured by credit score. Lenders may also use listing content to infer along these alternative dimensions of creditworthiness. This affects our results in two ways. First, while our results suggest that lenders infer a third of the variation in creditworthiness that is captured by credit score, lender inference could be higher or lower for other dimensions of creditworthiness. Second, our decomposition of sources of inference shows what types of information are useful in terms of inferring credit score. While this is likely an estimate of the source of information's overall contribution to inference of creditworthiness, we may overestimate or underestimate its value for inference of dimensions of creditworthiness not captured through credit score.

Our paper adds to the recent literature that examines peer-to-peer credit markets. Recent work on Prosper in particular shows that these markets display discrimination based on personal attributes like race and physical appearance (Pope and Sydnor, 2008; Ravina, 2008; Theseira, 2008). If such discrimination is taste-based, it brings into question the ability of these markets to distribute credit based upon borrower creditworthiness. Our paper complements this literature by focusing instead on whether there is any direct evidence of lenders' inferring borrower creditworthiness in such markets. We do find evidence of such inference, and we examine the types of information that

[7] Specifically, the credit score disregards certain codable observable characteristics (such as race or location) because of legal restrictions, and by definition, it is not directly based on borrower characteristics that are not observable to the rating agency. Moreover, a credit score is person-specific rather than specific to a person and a loan, so it ignores the loan characteristics. Finally, it is not practically feasible to condition the score on qualitative information, which is hard or impossible to code.

[8] While credit score is ostensibly based on "hard" factors like past borrower behavior and default history, these factors are also likely correlated with "softer" attributes such as a borrower's personal description and narrative.

lenders use to infer underlying creditworthiness. To the extent that credit score does not capture all aspects of creditworthiness, a generalization of our inference results suggests that even if these markets started reporting the exact credit score of borrowers, one would expect that the listing content could still help lenders improve the accuracy of their estimates of borrower creditworthiness.

Our paper also highlights how new lending platforms, such as peer-to-peer markets, may complement traditional lending models. One can think of these contributions in terms of (i) incentives and (ii) ability to screen. With regard to incentives, the setup of peer-to-peer markets is inherently competitive, with multiple lenders competing for the same borrower (see also Boot and Thakor, 1997).[9] Peer-to-peer lenders may also provide stronger screening incentives because they lack access to securitization markets, which may in turn lead to lax screening (Keys et al., 2010). In addition, the hierarchy in a peer-to-peer structure is completely flat, thus reducing the impediments of using and transmitting "soft" information that could help evaluate creditworthiness of small borrowers (Aghion and Tirole, 1997; Stein, 2002; Liberti and Mian, 2009).[10]

In terms of ability to screen, a concern is that peer-to-peer markets are likely to have lower individual financial expertise and experience to judge borrower creditworthiness. However, these markets may have participants who are skilled at judging particular aspects of the borrower that banks are unable to gauge . For example, a lender who works in the sector where the borrower proposes an entrepreneurial business idea may better assess the viability of the proposal. Peer-to-peer markets may also make better use of social network information. While Freedman and Jin (2008) find evidence of adverse selection due to informational problems faced by lenders in Prosper, they also find that social networks (endorsements by friends) may help alleviate these problems. In a similar spirit, Lin et al. (2009) find that stronger and more verifiable relational networks help reduce the adverse selection problems in Prosper. Thus, the possibility of adverse selection further enhances the value of screening borrowers in these markets. Finally, the aggregation of information in these markets could lead to better judgment of the creditworthiness of borrowers. A large set of

---

[9] Typically, small borrowers cannot simultaneously apply to a large number of banks. However, they can apply to a large set of lenders with a single loan application in a peer-to-peer market place. In addition, banks do not use multiple credit officers to screen small borrowers while peer-to-peer markets allow all lenders to potentially screen each borrower.

[10] This literature offers both theoretical and empirical evidence that, with respect to screening small borrowers, the organizational structure of banks may cause impediments in using subjective information to evaluate creditworthiness. A related literature on relationship banking documents the importance of soft information in allocating credit to small businesses and argues that flatter organizational structures increase the use of soft information (Berger and Udell, 2002; Santikian 2009).

individuals may collectively be better able to judge parts of the borrowers' information, particularly the non-standard or subjective aspects.[11]

Our results suggest that, despite not being financial experts, individual lenders in peer-to-peer markets can partly infer underlying borrower creditworthiness. Given these markets' ability to make valuable inferences and their non-collateral-based lending structure, peer-to-peer markets can be particularly helpful for small borrowers who may otherwise be limited to costly sources of finance like payday lenders and credit-card debt. In addition, increasing the ability of borrowers to credibly signal their quality can help improve the screening function of these markets. However, the inference of lenders is incomplete, and combined with evidence of possible discrimination, it is clear that these markets have their shortcomings as well. In sum, our results suggest that new models, such as peer-to-peer lending can indeed complement existing lending models and improve access to credit, particularly for small individual borrowers.

II. Context and Data

A. Context

The marketplace model of peer-to-peer lending on the internet enables individual lenders to locate individual borrowers and vice-versa. There has been an explosive growth in the online peer-to-peer market across the world. In the U.S alone, there are around twelve active online peer-to-peer lending sites. Furthermore, in Europe and Asia, online peer-to-peer lending markets are on the increase.[12] 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 US clients and intermediates capital mostly between individual lenders and small borrowers. Prosper has funded over $178 million in loans and currently has 830,000 members.

All Prosper loans are personal, three-year fixed-rate, unsecured loans. Borrowers request loans by creating a public listing on the Prosper.com website, and they can 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 a good credit risk, and their income and expenditure

[11] The success of micro-credit is partly attributed to the importance of joint liability for screening (Stiglitz, 1990; Ghatak, 2000). However, in the case of micro-credit, it is the other group members (who are jointly liable for the loan) that provide the screening rather than a larger set of individual lenders (typically unconnected to the borrower) who collectively screen the borrower in a more market-based setup.

[12] See http://en.wikipedia.org/wiki/Peer-to-peer_lending.

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, public records, and income. The Appendix provides one such sample listing.

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 upon 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.[13]

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

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 exceeds the loan amount requested by the borrower. If the total amount bid by lenders is greater than the amount requested by the

[13] 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 above. For consistency of results, we restrict our sample to the post February 12, 2007 period. However, results are robust to using the pre February 12, 2007 sample (see Table 4).

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 bids 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 according 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 more 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.[14] Our sample contains all listings posted between February 12, 2007 and October 2008.[15] Our sample covers 194,033 listings, of which 17,212 were funded.

Table 1 provides summary statistics of 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 banking variables and non-standard variables. The standard banking variables include hard, verified financial 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). Among the universe of listings, the average loan amount requested is $8015. The average maximum interest rate borrowers are willing to pay is 21%. Bad credit categories and high 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

---

[14] Note that even borrowers do not have access to the exact credit score obtained from the credit rating agency. 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.

[15] 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 in February 12, 2007. See Section 2, Part A for more details.

representing individuals with better credit variables are much more likely to be funded. For example, the debt-to-income ratio among funded listings is significantly lower at 33%.

The set of non-standard variables includes borrower choice variables that are unique to the Prosper marketplace, as well as basic coded information drawn from soft/qualitative listing content (pictures, text descriptions, friend endorsements, etc.). 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 soft 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 of borrower credit score. 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 residual uncoded sources of listing content.

III. Methodology

Our empirical strategy exploits the fact that credit scores are only reported as categorical variables to Prosper lenders. Thus, if we find that the interest rate at which lenders are willing to lend decreases with the exact credit score *within* a credit category, it must be that lenders are able to infer differences in creditworthiness across borrowers in the same credit category from other information provided on the website.[16] Moreover, given that lenders do observe credit categories, we can quantify lenders' inference of creditworthiness 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. 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.

As we detail below, our strategy also sheds light on the extent to which lenders rely on different types of information to make their inference about creditworthiness. While it may seem challenging to quantify or code qualitative data (such as pictures and other personal details), an advantage of our strategy is that we can still derive the contribution of such information: the

[16] Even if lenders are not consciously doing so, they act as if they are discerning between shades of creditworthiness since they are bidding on interest rates based upon their inferred potential returns to an investment.

contribution of non-quantified 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* quantified information.

The data section already described the listing information available to lenders. As noted previously, we categorize the information provided by borrowers based on whether it consists of standard banking variables (typically hard and verified financial information) or non-standard banking variables (typically soft and unverified information such as pictures or textual descriptions). The idea behind this classification is to distinguish between the information traditional lenders like banks typically use and the more subjective, softer, and non-verified information that is commonly available in peer-to-peer markets.

A. Estimating Overall Inference

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 1 plots the borrower's exact credit score, which is a proxy for 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). In this stylized figure, we assume that this hypothetical relationship is linear. We denote the credit score at the border between category k-1 and category k by c_k, and in this stylized figure, we assume that all credit categories are of equal size. 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 the borrowers' creditworthiness from information *other* than the categorical credit-score variable (as illustrated by the discontinuous downward sloping red line).

The degree to which lenders are able to 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 discontinuously drops by an amount α at each credit-score boundary. Hence, the total drop over one credit category (including one boundary) equals $\alpha + \beta$. We denote this total drop by $\delta \equiv \alpha + \beta$. 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 $\gamma \equiv \beta / \delta = \beta / (\alpha + \beta)$, and refer to γ as the amount of "inference" made by lenders.

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

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

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. As 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. This means that $CreditScore_i$ / $CatSize$ increases by exactly one if we move from the starting point of a credit category to the ending point. 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 $CreditScore_i$ / $CatSize$ increase 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, while 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.

The benefit of this stylized setup and the corresponding regression is that it is simple. 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. Figure 2 depicts this more realistic situation. The dashed blue line shows the underlying relationship between interest rate and exact credit score for the hypothetical scenario that exact credit score were 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 $\delta_k = (1-\lambda_k)\alpha_k + \lambda_{k+1}\alpha_{k+1} + \beta_k$.[17] 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:

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

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 is constant everywhere else. The coefficient α_k measures the jump in interest rate at the credit-score boundary between

[17] By definition, we cannot 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 $(1-\lambda_7)\alpha_7$.

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:

$$I_k^{Cum}(CreditScore_i) = \begin{cases} 0 & \text{if} \quad CreditScore_i < c_k \\ 1 & \text{if} \quad CreditScore_i \geq c_k \end{cases}, \tag{3}$$

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:

$$FracGap_k(CreditScore_i) = \begin{cases} 0 & \text{if} \quad CreditScore_i \leq c_k \\ \dfrac{CreditScore_i - c_k}{c_{k+1} - c_k} & \text{if} \quad c_k < CreditScore_i \leq c_{k+1} \\ 1 & \text{if} \quad c_{k+1} < CreditScore_i \end{cases}, \tag{4}$$

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 (2), 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 αs are jointly 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 equations (1) or (2), 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.[18] This censoring problem would bias our estimates of inference if we estimate equations (1) or (2) using ordinary least squares. Instead, we estimate equations (1) and (2) as censored regressions with the censoring occurring at the borrower maximum rate specified by each listing. The censored regressions, which are a generalization of the Tobit model, rest 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, the estimates from the censored regressions will yield consistent estimates of the parameters determining the interest rates that lenders require to fund a listing.

We use a modified version of equation (2) to test whether the exact credit score is predictive of default. In particular, we use an indicator for whether the loan defaulted as the dependent variable (rather than the interest rate). In this case, the β_k measures the predictive power of the exact credit score for default while the α_k measures whether the probability of default jumps at the credit-category boundaries.

B. Decomposing Inference by Source of Information

So far, 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 (2) controls for all the quantified listing variables:

$$InterestRate_i = \mu + \sum_{k=2}^{N} \alpha_k I_k^{Cum}(CreditScore_i) + \sum_{k=1}^{N} \beta_k^{Resid} FracGap_k(CreditScore_i) + \sum_{m=1}^{M} x_i^m \varphi^m + \varepsilon_i \,, \qquad (5)$$

[18] 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.

where x_i^m denotes the m^{th} quantitative variable in the listing of borrower i and φ^m denotes the corresponding regression coefficient.[19] In regression (5), 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 listing content outside the set of controls x_i^m. Since we measure exact credit scores 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 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: $\theta_k^m \varphi^m \equiv \beta_k^m$.

We decompose our original estimate β_k from the regression without the controls for quantified listing characteristics (regression 2) as follows:[20]

$$\beta_k = \beta_k^{Resid} + \sum_{m=1}^{M} \theta_k^m \varphi^m \equiv \beta_k^{Resid} + \sum_{m=1}^{M} \beta_k^m . \tag{6}$$

In equation (6), $\sum_{m=1}^{M} \beta_k^m$ is the part of the within-category drop in interest rates that can be attributed to quantified information, while the remainder is explained by non-coded information. Thus, rather

[19] 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.

[20] This is an application of the standard omitted variable bias formula. For a derivation and explanation of the omitted variable bias formula, see for example pages 245-246 of Greene (1993). 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.

than attempting to quantify 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 scores within credit-score categories after controlling for all quantitative information. To ensure that β_k^{Resid} reflects qualitative 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. Instead of reporting each single β_k^m, we report a sum of the βs that correspond to standard banking variables and a sum of the βs that correspond to non-standard variables. We also include β_k^{Resid}, which measures the contribution of non-coded information, with the non-standard 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).

IV. Results

We now present the results. We first show that credit score is indeed a proxy for creditworthiness. We then examine whether, and to what extent, lenders can infer the dimension of creditworthiness captured by credit score and explore what information they use to do so.

A. Does Credit Score Matter?

Table 2 first examines whether credit score is indeed related to underlying creditworthiness. While almost all credit scoring models use credit score as a predictor of creditworthiness and recent research supports the usefulness of credit score in mitigating adverse selection (e.g. Adams, Einav, Levin, 2009), we provide direct support for this by examining whether it predicts actual borrower behavior in our sample, such as the likelihood that a borrower will default on the loan. We take a conservative approach and classify a loan as in default if it is two or more months late.

We indeed find that credit score predicts the likelihood of default. Column (1) shows that, for every 40-point increase in underlying credit score, the likelihood of default decreases by 1.0 percentage point. We use 40-point intervals for ease of comparison given that Prosper defines categories based on 40-point intervals. The regression fit is not very strong in these regressions but this is to be expected since our outcome variable consists of default realizations on loans that have by and large not reached maturity while the exact credit score predicts default rate probabilities over the long term.

Since lenders observe credit categories, a related question is whether credit score is still predictive of default *conditional* on credit categories. Column (2) examines this possibility and shows that variation in credit scores within categories is indeed important in predicting borrower default. Within each credit category an increase of 40 points in credit scores implies a 1.2 percentage point lower default rate. While this measure is more relevant when we look at the interest rate as an outcome variable, we also provide the combined "gamma" value for this regression, i.e. the fraction of the underlying relationship between credit score and default rate that is captured *within* each credit category. Since the outcome in question is default rate, a factor likely based mostly on borrower behavior rather than lender inferences, and because default probability should be a continuous decreasing function of credit score, one would expect gamma to be close to one here. Column (2) shows that gamma is 1.16. We cannot reject that gamma is significantly different from one (p-value: 0.67). Column (3) implements a more flexible specification that is the equivalent of equation (2) in the methodology section but where we use default rate as the outcome variable. Here, the betas estimate how much within-category credit-score variation impacts the default rate, and the alphas capture the additional impact of each credit category. In addition to reporting the gamma for each credit category, we also calculate a combined gamma that, as described in the methodology section, is the weighted average of category-specific gammas. The combined gamma estimate is 1.04, and we cannot reject that it is different from one (the same holds for the individual gammas as well). The combined gamma reported in both Column (2) and Column (3) being close to one is reassuring and offers an informal check on our methodology because in cases where the outcome variable is a direct outcome of credit score (in other words, it is not inferred by lenders), one would expect that gamma would be close to one.

Column (4) performs a robustness check on our definition of default and shows that the same results hold when we replicate Column (3) but define a loan to be in default if it is more than one month late.

B. Can Lenders Infer Creditworthiness?

Having provided evidence that credit score captures a dimension of creditworthiness (since it predicts default behavior), we now turn to the main question of the paper: Are lenders are able to infer this dimension of creditworthiness from information provided in the listing?

Before turning to our regressions, we present the empirical analogue to Figures 1 and 2. In Figure 3, 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 credit score from other listing information. In addition, there are discrete jumps in interest rates at the credit-category boundaries, which shows that lenders exhibit imperfect inference of the full information content of 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 3 presents this regression. The coefficient on credit score/40 shows that the interest rate falls by 0.54 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 infer variation in creditworthiness within credit categories from other information in the listing. The coefficient on 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 × 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 from Figure 1 and the first regression in Table 3 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 the 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 (2). 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 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, where the censoring takes place at the listing-specific borrower maximum rate.

Column (2) of Table 3 implements our preferred approach (equation (2) in the methodology section) and estimates directly the extent of inference that takes places. 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. The large magnitude of our estimate of combined gamma suggests that, despite not being financial experts, lenders are collectively able to exploit other information provided on the Prosper site in order to infer creditworthiness.

To understand the economic significance of this result, note that the αs and βs sum to 39 percentage points. In order words, the mean offer 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*40/(900-520)) for a typical 40-point credit category. This decline in the offer rate is greater than the decline in the market interest rate because the 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. The inference estimate of 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, which implies that they are willing to offer an interest rate that is 137 (= 0.330 × 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 category, despite not observing exact credit score.

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, but 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 to therefore compare high and low credit categories by grouping individual ones, and we will do so later.

The fact that inference is incomplete (γ <1) implies that borrowers just below a category boundary pay a significantly higher interest rate than borrowers just above the boundary. One may therefore expect that Prosper disproportionally attracts listings by individuals with credit scores in the lower ranges of each category, and Freedman and Jin (2008) 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.

C. Robustness of Lender Inference:

While the results in Table 3 suggest that lenders are able to infer a part of borrower creditworthiness (proxied by the credit score), one may raise the concern that this finding does not reflect inference but rather direct communication of the exact credit score by the borrowers to the lenders. We do not think such concerns are valid in practice for several reasons. First, Prosper prohibits any direct contact between borrower and lenders. While it does allow borrowers to post information in the listing and also has a facility for questions and answers (intermediated via Prosper), this information is unverified. 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. Information such as credit score or social security numbers would be strictly unacceptable, and efforts were taken to ensure no such information was posted or seen. Nevertheless, as a robustness check, we also estimate lender inference in the sample period (prior to

February 12, 2007) when there was no facility for question and answers (Table 4, row 2).[21] As the results show, we find that even in this sample period, the inference parameter gamma is 0.46. This confirms that our estimate of inference is not a result of direct communication but indeed due to inference by lenders.

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 these changes are not a concern in practice. In Table 4, rows 3 and 4, we estimate the gamma for the sample before and after these changes. We find that the combined gamma is similar both in the pre- and post change period. Another concern could be that borrowers in some states are subject to usury laws (Rigbi, 2009). These laws may create an artificial ceiling on the interest rates and impact the extent of inference. As a robustness check, we also estimate the gamma for the period without usury law restrictions, and we again find a gamma of 0.32 (row 5). We also carry out several other robustness checks. To address the concern that some borrowers are affiliated with groups where group members might know each other and share personal information, we also estimate the gamma for a sample restricted to borrowers that are not affiliated with any groups and find similar results (row 6). 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 gamma for a sample restricted to first-time loans and listings (row 7) and to first-time loans (row 8). We again find similar results. Since our methodology relies on taking advantage of boundaries between credit categories and because the two extreme categories do not have boundaries on both sides, we also estimated the gamma excluding the top and the bottom credit categories and find that our results remain robust (row 9).

The estimate of inference in our baseline specification draws both on the observed interest rate for the subsample of funded listings and the information contained in whether a listing is funded or not. In a final pair of robustness tests, we estimate inference if we only use one of these two sources of information. In row 10, we ignore information contained in the observed interest rate by estimating a censored probit of a dummy for whether the listing is funded on the same

[21] For documentation of this implementation date, see http://www.prosper.com/help/topics/whats_new.aspx.

explanatory variables as in our baseline regression. In row 11, we ignore information contained in the funding decision by running a truncated regression on the subsample of funded listings. In both specifications, we estimate a statistically significant gamma that is similar in magnitude to our baseline estimate.

In row 12, we estimate our baseline model using OLS for the subsample of funded listings. This is not strictly speaking a robustness test since the OLS regression does not properly account for censoring. We find a gamma of 0.39, suggesting that our estimate of overall inference would not be severely biased if we failed to correct for censoring on the borrower maximum rate. However, as discussed later, correcting for censoring turns out to be important in order to correctly decompose inference of credit score from different sources of information.

Figure 4 presents an illustration of how combined gamma varies over time. We divide the data up into bi-monthly time periods and plot the gamma for each period. While, as expected, there is some variation given the sample periods, sizes, and policy changes, by and large, the inference in each period is substantial, and differences over time are within the margin of error shown by the 95-percent confidence intervals.

D. What Information Do Lenders Use to Infer Creditworthiness?

While it is remarkable that lenders in a peer-to-peer market are able to infer a third of the variation in creditworthiness captured by credit score, what sources of information allow them to do so? As detailed in the methodology section, we can decompose our "inference parameter" gamma into the separate gammas for each of the variables that the borrower observes. We group information into two broad categories of interest: standard banking variables (variables generally used by banks) and non-standard variables (variables chosen by the borrower). Generally speaking, standard banking variables are more likely to be hard, verifiable, "screening type" variables, while non-standard variables are likely to be subjective, non-financial, potentially harder to verify, and more likely to behave like "signals."

The standard banking variables are readily coded, and we provided the details and summary statistics of variables included in this category in Table 1. Non-standard variables – the various "softer" pieces of information such as pictures, individual background, description, and online exchanges – while readily identified, are 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 not clear to which attributes a particular lender may react.

However, a key strength of our strategy is that, provided we appropriately control for all the hard information, we do not need to code or specify the soft information if we are only interested in understanding how much a lender is able to infer from such information. The idea is that the residual "gamma" (γ^{Resid}) will reflect the inference contribution from all such variables. Before presenting the results, we should note two caveats that we discussed previously. First, the decomposition presented is for inference drawn for the dimension of creditworthiness that is captured by credit score. Thus, the contribution of a variable in drawing inference along the credit-score dimension of creditworthiness need *not* equal its contribution to inference along a dimension of creditworthiness that is not captured by credit score. Second, if a particular variable impacts both the credit-score dimension of creditworthiness and another dimension, this may bias our estimate of the variable's contribution to the credit-score dimension. We will overestimate its contribution if the variable impacts the other dimension of creditworthiness in the same direction as the dimension captured by credit score (since part of the inference which we attribute to the credit-score dimension is really due to the other dimension) and underestimate it otherwise.

Table 5 presents the result of our decomposition. For the sake of brevity, we only present the combined inference parameter, gamma, in Table 5. The first column presents the results from a single regression (equation (5) in the methodology section) that decomposes the total combined gamma into components that are explained by specific variables in the listing. The next two columns present this decomposition separately for the low credit categories (HR, E, D, and C) and for the high credit categories (B, A, and AA).[22] 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 3 (Column (2)). As before, the total combined gamma is 0.33.[23] We find that the gamma for the lower credit categories is 0.25, while the gamma for the high credit categories is 0.41. An F-test rejects equality of estimates between the high and the low credit categories, suggesting that there is differential inference across credit categories. The next rows present the contributions that the standard banking and non-standard banking 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-

[22] We chose this categorization as it roughly provides us with an equal number of loans in both categories.

[23] In the first line of Table 5, we report the sum of all the components of γ. As noted in the methodology section, the decomposition of gamma into its components only holds in expectation in the case of a censored regression. As a result, the estimate of the sum of the components, 0.328 from equations (5) and (6), is close but not identical to the direct estimate of gamma, 0.330 from equation (2), that we presented in Table 3.

category that show the largest (in magnitude) inference. The Appendix Table presents the individual gammas for all the variables separately.

Reading down the first column in Table 5, we see that of the total gamma of 0.328, most of the inference comes from standard banking variables (0.312 or 95%). The inference from the non-standard banking variables is only 0.016 or 5%. However, part of the reason there is less inference drawn from non-standard banking variables is that gamma is negative for some of these variables, masking the positive contribution to inference of other non-standard banking variables. We revisit this issue of negative contributions to inference later in this section. 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, in general, lenders learn more from standard banking variables, which are more 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 borrower's creditworthiness but also are verified and therefore less subject to the possible "cheap talk" concerns of voluntarily posted and unverified information. Moreover, it is possible that the standard banking variables are more closely associated with the dimension of credit score captured by creditworthiness, although credit score is likely to be influenced by "softer" borrower attributes, as well.

Second, in examining which variables are used by lenders to draw inferences among the standard banking variables, we find that most of the inference is driven by 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 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.

To provide some insight into such differences in relative inference, we offer a mechanical explanation of why the magnitude of the inference changes for a given variable across the high and low credit categories. In the methodology section, we explained how 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 borrowers' 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 increase. 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.

The third main finding from the decomposition exercise is that inference from non-standard banking variables is relatively more important for lower credit categories, especially when we consider some of the specific variables (such as borrower maximum rate) in this category.[24] This may not be surprising if one believes that (variation in) financial information is less revealing to distinguish between low-quality borrowers (e.g., differences between someone being delinquent ten times versus twelve may be less revealing than zero versus two times) This leaves more room in the lower credit categories to rely on non-traditional methods of screening. However, as evidenced by several variables that show negative inference, this also leaves more room for incorrect inferences being drawn by lenders.

Among the coded non-standard 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) - the average inference is 0.064 (or 19% of total inference) across all credit categories and is greater for lower (33.9%) than higher credit categories (10.2%). The fact that the borrower maximum rate generates much more inference than other information in the non-standard variables group is not surprising for two reasons. First, this information is verified. Second, and perhaps more importantly, it is likely to serve as a credible signal of creditworthiness. As one would expect, borrowers that post 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 (since exact credit scores are observable by banks), it is less costly for them, relative to less creditworthy borrowers, to post a lower borrower maximum rate.

[24] Note that since credit score is likely to be more directly influenced by hard information and standard banking variables, our estimate on importance of soft information likely represents a lower bound, and soft information may be more valuable along dimensions of creditworthiness other than credit score.

While establishing this as a separating equilibrium requires 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. Examining the results in more detail shows that there is greater inference for borrower maximum rate in lower credit categories because these categories show a higher sensitivity of the interest rate to this borrower choice variable.

The fourth main finding from Table 5 concerns the importance of inference from uncoded soft information (the "residual" inference). While the residual gamma is insignificant for the whole sample, we estimate a statistically significant residual gamma of 0.096 (39% of total inference) from uncoded sources in the lower credit categories. This suggests that, in the lower credit categories, lenders draw inferences from subjective listing content that we did not code. We find similar results when we estimate the "residual" inference using specifications where we use linear controls or cubic controls for all of our x variables (results not reported), suggesting that this estimate is robust to the form of the specification.

We further note that not all measured inference is positive. For some variables, like amount requested, this negative inference likely reflects inference along other dimensions of creditworthiness since it is plausible that, holding credit score constant, larger loan amounts increase default likelihood.[25] For other variables (to the extent that we believe lenders are driven by profit motives), this negative inference may be indicative of mistakes lenders make. An alternate interpretation could be that lenders do know that a borrower is more likely to default but still offer her a better interest rate due to charitable motives. Whether such incorrect or non-profit maximizing inference can be sustained in equilibrium is a more complicated question. However, it does suggest that there may be pitfalls and challenges to inference, particularly from (non-verified) information that borrowers choose to post.

[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, they incorrectly believe that a variable representing a negative borrower attribute is positively correlated with credit score and mistakenly offer lower interest rates for higher values of that variable), in the case of amount requested, we believe that this is due to the concern regarding our decomposition exercise, namely 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, 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. Hence, we are likely to underestimate the degree of inference about creditworthiness from amount requested. In our discussions we therefore deemphasize amount requested, focusing instead on variables for which the inference estimate is less likely to be biased.

V. Conclusion

Our results show that lenders in peer-to-peer markets are able to partly infer borrowers' creditworthiness using the rich information set that these markets provide. Moreover, while lenders in these markets mostly rely on standard banking variables to draw inferences on creditworthiness, they also use non-standard or soft sources of information in their screening process, especially in the lower credit categories. In addition, the use of credible signals (like borrower maximum rate) in screening suggests that enhancing the opportunity for borrowers to post credible signals can further help in facilitating the screening process. While this finding is reassuring in that it suggests that these markets are not entirely influenced by "cheap" talk, there is the caveat that lenders in these markets may sometimes make incorrect inferences.

The broader question, though, is to what extent peer-to-peer markets can complement traditional lenders such as banks. In a very narrow sense, one may argue that if the only thing these markets can infer is the credit score, then revealing the score would take away the need to make such inference. However, it is implausible to think that creditworthiness is fully captured by an individual's credit score. We focus on credit score only because it provides us with a strategy to identify how much lenders can infer about a factor that reflects creditworthiness. To the extent that such inference is similar for other dimensions (besides credit score) that reflect creditworthiness but which may be much harder to quantify or verify, this paper suggests that peer-to-peer markets hold significant promise. Moreover, our results show greater lender inference from credible information, which suggests that modifications to the design of these markets (by facilitating such credible signals) may further improve screening from subjective information. The uncollateralized nature of lending and the ability of lenders to partly screen suggests 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 costly sources of finance like payday lenders and credit-card debt. How best to design these markets to further enhance their role in allocating credit is a promising direction for future enquiry. The current financial crises may provide the additional interest and impetus to do so.

References:

Adams, William, Liran Einav, and Jonathan Levin, 2009, "Liquidity Constraints and Imperfect Information in Subprime Lending," *American Economic Review*, 99(1): 49-84.

Aghion, Philippe, and Jean Tirole, 1997, "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105(1): 1-29.

Altonji, Joseph G., and Charles R. Pierret, 2001, "Employer Learning and Statistical Discrimination," *Quarterly Journal of Economics*, 116(1): 313-350.

Ang, James S., James W. Lyn, and Floyd Tyler, 1995, "Evidence on the Lack of Separation Between Business and Personal Risk Among Small Businesses," *Journal of Small Business Finance*, 4(2): 197-210.

Avery, Robert B., Raphael W. Bostic, and Katherine A. Samolyk, 1998, "The Role of Personal Wealth in Small Business Finance," *Journal of Banking and Finance*, 22(6-8): 1019-1061.

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.

Boot, Arnoud W.A., and Anjan V. Thakor, 1997, "Financial System Architecture," *Review of Financial Studies*, 10(3): 693-733.

Diamond, Douglas W., 1984, "Financial Intermediation and Delegated Monitoring," *Review of Economic Studies*, 51(3): 393-414.

Farber, Henry S. and Robert Gibbons, 1996, "Learning and Wage Dynamics," *Quarterly Journal of Economics*, 111(4): 1007-1047.

Freedman, Seth, and Ginger Z. Jin, 2008, "Do Social Networks Solve the Information Problems for Peer-to-Peer Lending? Evidence from Prosper.com," Working paper, University of Maryland.

Ghatak, Maitreesh, 2000, "Screening by the Company You Keep: Joint Liability Lending and the Peer Selection Effect," *Economic Journal*, 110(465): 601-631.

Greene, William H., 1993, "Econometric Analysis," 2nd ed. Macmillan, New York.

Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, 2010 (forthcoming), "Did Securitization Lead to Lax Screening? Evidence From Subprime Loans," *Quarterly Journal of Economics*, 125(1).

Liberti, Jose M., and Atif R. Mian, 2009 (forthcoming), "Estimating the Effect of Hierarchies on Information Use," *Review of Financial Studies*.

Lin, Mingfeng, Siva Viswanathan, and N. R. Prabhala, 2009, "Judging Borrowers by the Company they Keep: Social Networks and Adverse Selection in Online Peer-to-Peer Lending," Working paper, University of Maryland.

Manove, Michael, A. Jorge Padilla, and Marco Pagano, 2001, "Collateral versus Project Screening: A Model of Lazy Banks," *RAND Journal of Economics*, 32(4): 726-744.

Pope, Devin G., and Justin Sydnor, 2008, "What's in a Picture? Evidence of Discrimination from Prosper.com," Working paper, Wharton.

Ravina, Enrichetta, 2008, "Love & Loans: The Effect of Beauty and Personal Characteristics in Credit Markets," Working Paper, Columbia GSB.

Rigbi, Oren, 2009, "The Effect of Usury Laws: Evidence from the Online Loan Market," Working Paper, Stanford GSB.

Santikian, Lori, 2009, "Dissecting Bank Relationships: The Role of Cross-Selling and Connections in Commercial Lending to Small Businesses," Working paper, Harvard University.

Stein, Jeremy, 2002, "Information Production and Capital Allocation: Decentralized vs. Hierarchical Firms," *Journal of Finance*, 57(5): 1891-1921.

Stiglitz, Joseph E., 1990, "Peer Monitoring and Credit Markets," *World Bank Economic Review,* 4(3): 351-366.

Stiglitz, Joseph E., and Andrew M. Weiss, 1981, "Credit Rationing in Markets with Imperfect Information," *American Economic Review*, 71(3): 393-410.

Theseira, Walter, 2008, "Competition to Default: Racial Discrimination in the Market for Online Peer-to-Peer Lending," Working paper, Wharton.

Wolfers, Justin, and Eric Zitzewitz, 2004, "Prediction Markets," *Journal of Economic Perspectives*, 18(2): 107-126.

Figure 1: Stylized Relationship between Interest Rate and Credit Score

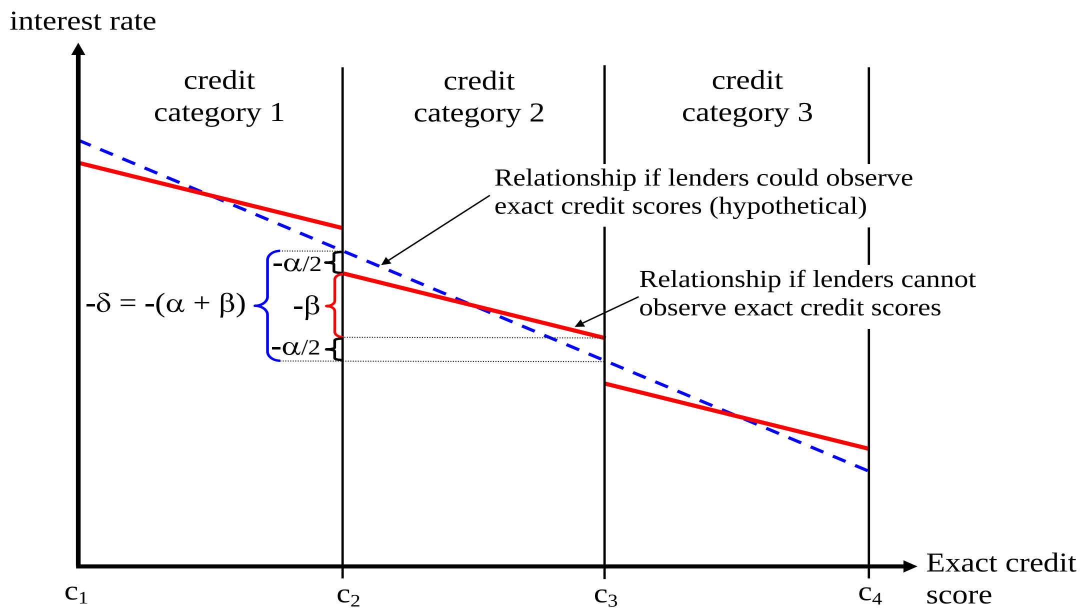

This figure shows the stylized hypothesized relationship between a borrower's credit score and the market interest rate on her (funded) loan.

Figure 2: Relationship between Interest Rate and Credit Score

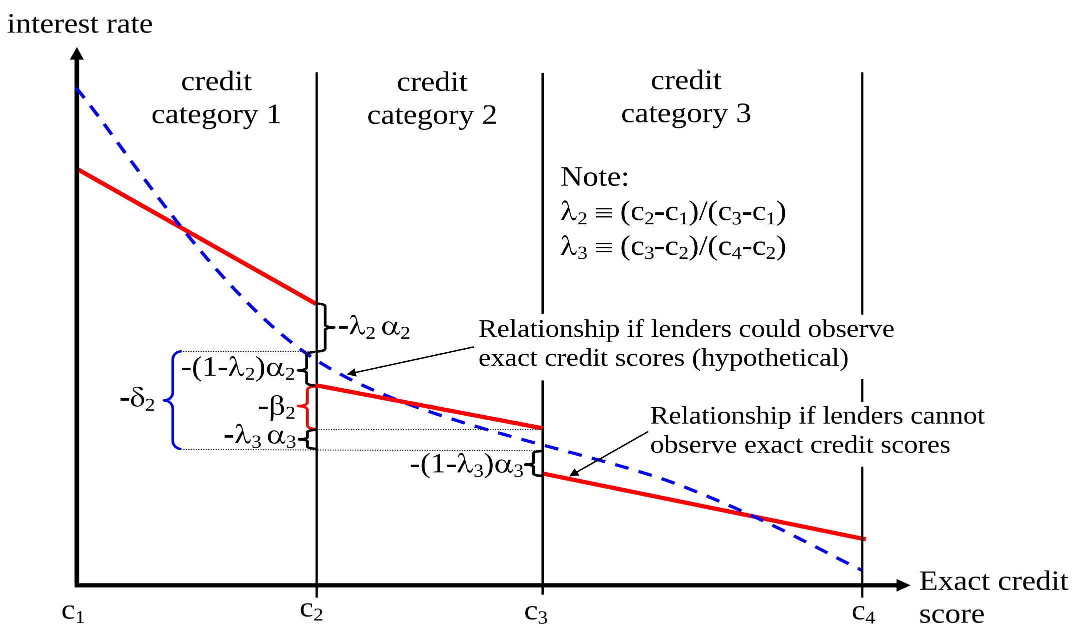

This figure shows a more realistic hypothesized relationship between a borrower's credit score and the market interest rate on her (funded) loan.

Figure 3: Market Interest Rate and Credit Scores

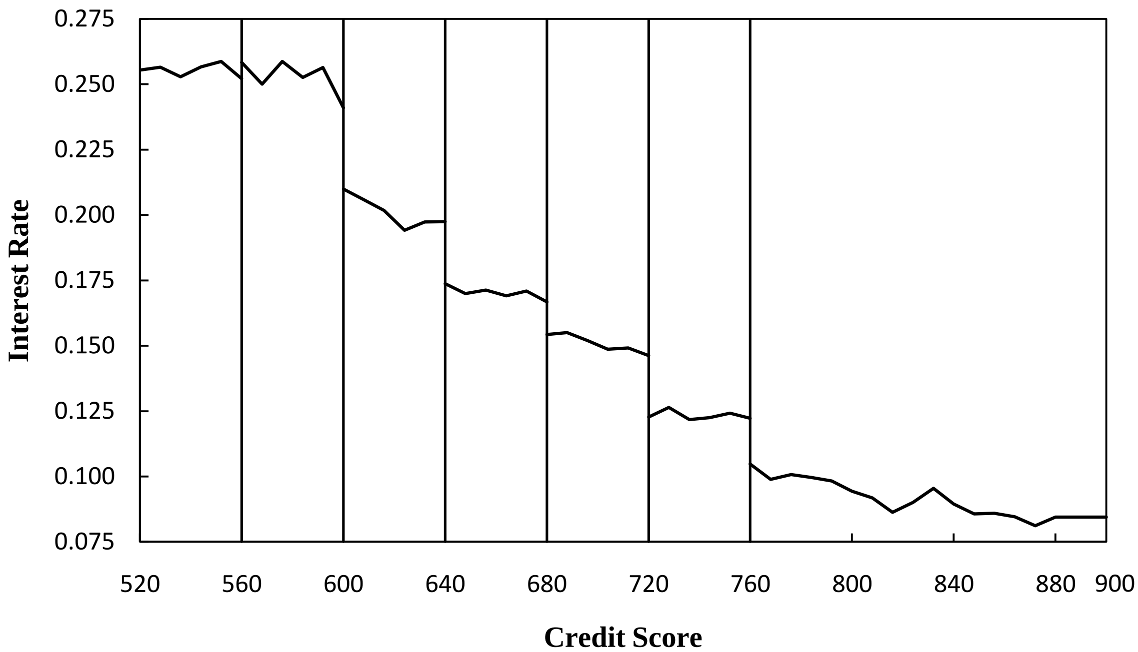

This figure shows the "raw" relationship between a borrower's credit score and the interest rate on her funded loan. Each point in the graph plots the average interest rate over an eight-point range in credit scores. Solid lines separate the seven credit categories. Starting from left to right, the categories 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.

Figure 4: Inference Over Time

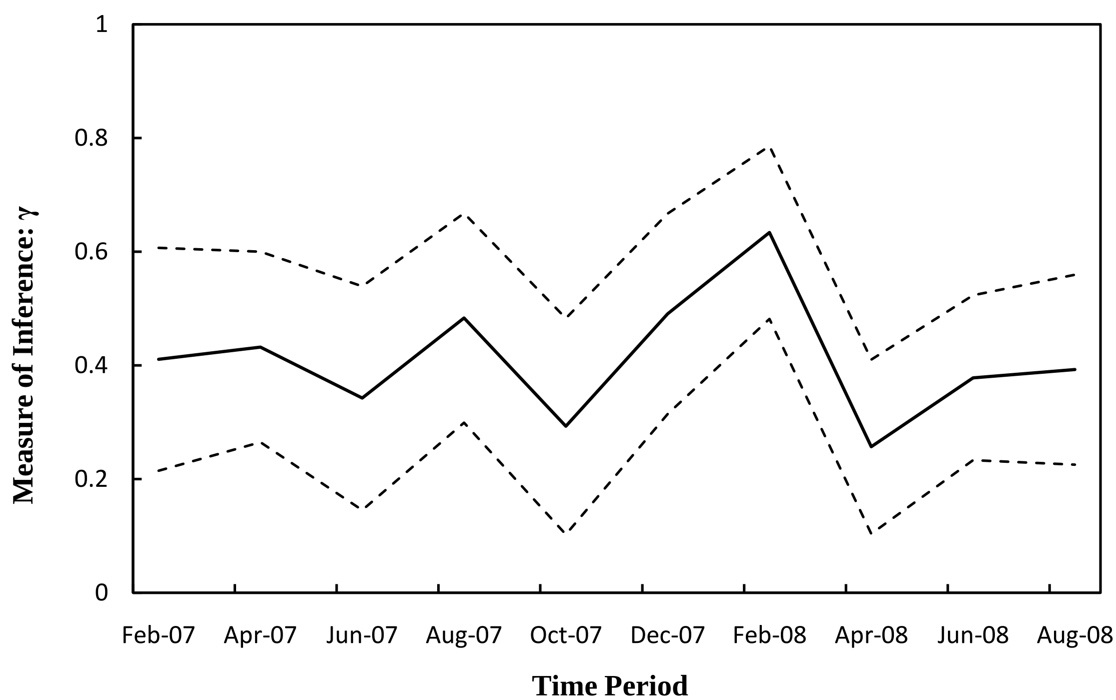

This figure shows our measure of inference, γ, for each two-month window from February 2007 to September 2008. Dotted lines represent 95% confidence intervals for each two month γ estimate.

Table 1: Summary Statistics

	All Listings		Funded Listings	
	Mean	S.D.	Mean	S.D.
General				
Credit Score	609.5	73.8	676.0	74.5
Credit Category Dummies				
Credit Category HR	0.343		0.068	
Credit Category E	0.164		0.074	
Credit Category D	0.178		0.173	
Credit Category C	0.136		0.211	
Credit Category B	0.082		0.183	
Credit Category A	0.055		0.140	
Credit Category AA	0.044		0.152	
Loan Outcomes				
Annual Lender Interest Rate			0.166	0.068
Fraction 1 or more months late			0.063	
Fraction 2 or more months late			0.044	
Fraction 3 or more months late			0.031	
Fraction of Listings Funded	0.089			
Standard Banking 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
Employment Status Dummies				
Full-Time	0.812		0.859	
Part-Time	0.041		0.040	
Self-Employed	0.096		0.074	
Retired	0.028		0.020	
Not Employed	0.023		0.008	
Length of Current Employment Status (months)	20.91	51.90	22.73	53.52
Personal Annual Income Dummies				
N/A or Unable to Verify	0.053		0.025	
Not Employed	0.021		0.007	
$1- $24,999	0.163		0.120	
$25,000 - $49,999	0.402		0.372	
$50,000 - $74,999	0.211		0.253	
$75,000 - $99,999	0.078		0.117	
$100,000+	0.064		0.101	

Table 1 - Continued: Summary Statistics

	All Listings		Funded Listings	
	Mean	**S.D.**	**Mean**	**S.D.**
Non-Standard Variables				
Borrower Maximum Rate	0.21	0.09	0.21	0.08
Duration of Loan Listing Dummies				
3 Days	0.044		0.037	
5 Days	0.046		0.055	
7 Days	0.693		0.661	
10 Days	0.218		0.247	
Listing Category Dummies				
Not Available	0.386		0.380	
Debt Consolidation	0.281		0.262	
Home Improvement Loan	0.024		0.033	
Business Loan	0.098		0.100	
Personal Loan	0.114		0.121	
Student Loan	0.025		0.024	
Auto Loan	0.017		0.017	
Other	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				
HTML Character Number	283	271	309	350
Text Character Number	963	716	1106	806
Average Word Length	4.63	0.58	4.59	0.55
Average Sentence Length	122.75	97.14	106.96	68.62
Number of Numerics	13.03	11.31	14.49	14.32
Percent of Words Misspelled	0.03%	0.03%	0.03%	0.04%
Number of Dollar Signs	8.98	5.78	8.49	7.25
Percent of Listing as Signs	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				
0%	0.916		0.867	
0.25%	0.002		0.010	
0.50%	0.015		0.046	
0.75%	0.001		0.002	
1.00%	0.034		0.047	
1.50%	0.004		0.007	
2.00%	0.019		0.017	
3.00%	0.006		0.003	
4.00%	0.003		0.001	
Number of Friend Endorsements	0.324	0.769	0.519	0.973
Observations	**194033**		**17212**	

For the sake of brevity, we do not provide summary statistics of 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 specifications in Table 5 and in the Appendix tables. Definitions of variables that may not be self-explanatory are as follows: *Percent of Listings as Signs* refers to the percentage of the listing text that is composed of non alpha-numeric signs, e.g. $/.,{}(). *HTML Character Number* refers to the number of characters in the listing text used to specify html formatting and reflects the extent to which borrowers formatted the text of their listings. *Public Records* includes information like bankruptcies, judgments, tax liens, state, and country court records, and, in some states, overdue child support, found in the borrowers' credit reports. *Bank Draft Annual Fee Dummy* equals one if the borrower elected to pay a 1% annual fee charged for not using the electronic funds transfer option.

Table 2: Default Rates

	(1)		(2)		(3)		(4)	
Dependent Variable:	**Default = Loan is 2 or more months late**		**Default = Loan is 2 or more months late**		**Default = Loan is 2 or more months late**		**Default = 1 or more months late**	
Estimate	**Coefficient**	**(S.E.)**	**Coefficient**	**(S.E.)**	**Coefficient**	**(S.E.)**	**Coefficient**	**(S.E.)**
Combined γ			1.160 ***	(0.378)	1.036 ***	(0.145)	0.927 ***	(0.154)
Regression Coefficients								
Credit score/40	-0.010 ***	(0.001)	-0.012 ***	(0.004)				
Credit category			0.002	(0.004)				
α_2: Change between HR and E					0.003	(0.012)	0.007	(0.014)
α_3: Change between E and D					-0.006	(0.010)	-0.009	(0.012)
α_4: Change between D and C					0.002	(0.008)	-0.005	(0.009)
α_5: Change between C and B					0.011	(0.009)	0.002	(0.010)
α_6: Change between B and A					0.000	(0.011)	0.001	(0.012)
α_7: Change between A and AA					-0.012	(0.012)	-0.010	(0.015)
β_1: Change within HR					-0.017	(0.015)	-0.022	(0.018)
β_2: Change within E					0.007	(0.012)	-0.004	(0.015)
β_3: Change within D					-0.009	(0.010)	-0.001	(0.012)
β_4: Change within C					-0.025 **	(0.010)	-0.017	(0.012)
β_5: Change within B					-0.012	(0.011)	-0.011	(0.013)
β_6: Change within A					-0.011	(0.016)	-0.021	(0.018)
β_7: Change within AA					-0.017	(0.023)	-0.037	(0.028)
N	17212		17212		17212		17212	
R^2	0.077		0.077		0.079		0.071	
Implied Coefficients and Tests								
$\gamma_1 = \beta_1/\delta_1$					1.190	(0.824)	1.488	(1.090)
$\gamma_2 = \beta_2/\delta_2$					1.266	(1.456)	0.796	(2.313)
$\gamma_3 = \beta_3/\delta_3$					0.842	(0.638)	0.163	(1.334)
$\gamma_4 = \beta_4/\delta_4$					1.382 ***	(0.363)	0.936 **	(0.424)
$\gamma_5 = \beta_5/\delta_5$					1.951	(1.675)	1.162	(0.919)
$\gamma_6 = \beta_6/\delta_6$					0.800	(0.629)	0.908 **	(0.375)
$\gamma_7 = \beta_7/\delta_7$					0.501	(0.511)	0.704 *	(0.390)
δ_1: Overall Change for HR					-0.015	(0.010)	-0.015	(0.012)
δ_2: Overall Change for E					0.006	(0.008)	-0.005	(0.010)
δ_3: Overall Change for D					-0.011	(0.007)	-0.008	(0.008)
δ_4: Overall Change for E					-0.018 ***	(0.007)	-0.019 **	(0.008)
δ_5: Overall Change for B					-0.006	(0.008)	-0.010	(0.009)
δ_6: Overall Change for A					-0.013	(0.011)	-0.023 *	(0.013)
δ_7: Overall Change for AA					-0.035	(0.022)	-0.053 *	(0.027)
p-value: α_i=0					0.813		0.938	
p-value: γ_i=1			0.671		0.832		0.976	

This table examines whether credit score predicts creditworthiness as represented by default rates. Each specification includes listing month fixed effects to control for listing age. Column (1) shows marginal effects from a probit regression of default (defined as two or more months late) on credit score divided by average credit category size. Column (2) examines whether credit score is predictive of default after conditioning on credit categories. Column (3) implements a more flexible specification that is the equivalent of Equation (2), Section 3 except with default rate as the dependent variable. Column (4) shows that similar results hold when we replicate Column (3) but define a loan to be in default if it is one or more months late. Results are also robust when default is defined as three or more months late. 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: Inferring Creditworthiness

	(1)			(2)		
Dependent Variable: Interest Rate	**OLS**			**Censored Regression**		
Estimate	**Coefficient**		**(S.E.)**	**Coefficient**		**(S.E.)**
Combined γ: Inference				0.330	***	(0.033)
Regression Coefficients						
Credit score/40	-0.005	***	(0.001)			
Credit category	-0.022	***	(0.001)			
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	17212			194033		
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)
δ_1: Overall Change for Credit Category HR				-0.049	***	(0.005)
δ_2: Overall Change for Credit Category E				-0.060	***	(0.004)
δ_3: Overall Change for Credit Category D				-0.081	***	(0.004)
δ_4: Overall Change for Credit Category E				-0.050	***	(0.003)
δ_5: Overall Change for Credit Category B				-0.055	***	(0.004)
δ_6: Overall Change for Credit Category A				-0.031	***	(0.005)
δ_7: Overall Change for Credit Category AA				-0.115	***	(0.008)
p-value: γ_i=γ				0.002		
p-value: γ_i=0				0.000		

This table examines the ability of lenders to infer borrower credit score. Column (1) takes a simple approach and asks whether, conditional on the observable credit category, credit score predicts the interest rate. It estimates an OLS specification in which the sample is restricted to funded listings. Column (2) implements our baseline specification described in Equation (2), Section 3 and estimates the extent of inference that takes place using the full baseline sample, including unfunded listings. In Column (2) and all tables hereafter unless otherwise noted, all coefficient, combined, and implied estimates are based upon censored normal regressions with 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 4: Robustness of Measure of Inference

Estimations using:		Combined γ Coefficient		(S.E.)	N
(1)	Baseline sample (All listings 2/12/2007 - 10/16/2008)	0.330	***	(0.033)	194033
(2)	Period without question and answers (Pre 2/12/2007)	0.461	***	(0.115)	5933
(3)	Period before suggested borrower maximum rate and portfolio plans (Pre 10/30/2007)	0.352	***	(0.051)	64485
(4)	Period after suggested borrower maximum rate and portfolio plans (Post 10/30/2007)	0.343	**	(0.038)	129548
(5)	Period without usury law restrictions on interest rates (Post 4/15/2008, excluding Texas and South Dakota)	0.317	***	(0.050)	68658
(6)	Sample restricted to listings with no group affiliation	0.351	***	(0.036)	159359
(7)	Sample restricted to first time listings	0.419	***	(0.044)	93117
(8)	Sample restricted to first time loans	0.355	***	(0.034)	183455
(9)	Baseline sample, measure of inference (γ) calculated excluding top and bottom credit categories	0.250	***	(0.042)	194033
(10)	Censored probit specification, dependent variable: funded dummy	0.287	***	(0.045)	194033
(11)	Truncated regression, sample restricted to funded listings	0.385	***	(0.077)	17212
(12)	OLS specification, sample restricted to funded listings	0.390	***	(0.033)	17212

This table supports the robustness of our inference estimates from Table 3. Combined gammas are calculated according to Equation (2), Section 3. Row (1) shows estimates from Column (2) of Table 3 based upon our baseline specification. Row (2) 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 (3) and (4) 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 (5) 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 (6) restricts the sample to listings posted by borrowers with no group affiliations. Rows (7) and (8) restrict the sample to listings that represent the first listing or first funded listing (loan) for borrowers, 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. Row (9) uses the full sample, but presents a combined gamma that excludes the lowest and highest credit categories, HR and AA. Row (10) shows the results from a censored probit specification with the dummy variable for whether the listing is funded as the dependent variable. Row (11) estimates a truncated regression using the funded listings sample, i.e. the sample where interest rate is not censored by the borrower maximum rate. Row (12) shows the results from an OLS specification with interest rate as the dependent variable, restricted to the funded listings sample. OLS does not account for the censoring of interest rates in unfunded listings by the borrower maximum rate. 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: Decomposing Inference

	(1)	(2)	(3)	(4)
	Gamma	**Gamma Low Cat (1-4)**	**Gamma High Cat (5-7)**	**Low = High p-value**
All Listing Content (γ)	0.328*** ***	0.244*** ***	0.417*** ***	0.001
	(0.027)	(0.044)	(0.028)	
Decomposition of γ				
1. Standard Banking Variables	0.312***	0.210***	0.421***	0.000
	(0.020)	(0.020)	(0.034)	
1.1 Number of Current Delinquencies	0.079 ***	0.110 ***	0.045 ***	0.000
	(0.006)	(0.010)	(0.007)	
1.2 Number of Credit Inquiries, Last 6 months	0.054 ***	0.073 ***	0.034 ***	0.000
	(0.003)	(0.004)	(0.003)	
1.3 Amount Delinquent	0.051 ***	0.085 ***	0.015 ***	0.000
	(0.006)	(0.010)	(0.006)	
1.4 Debt to Income Ratio	0.048 ***	0.001	0.099 ***	0.000
	(0.007)	(0.008)	(0.011)	
1.5 Amount Requested	-0.005	-0.124 ***	0.122 ***	0.000
	(0.005)	(0.006)	(0.009)	
1.6 All Other Standard Banking Variables	0.085 ***	0.065 ***	0.106 ***	0.226
	(0.016)	(0.017)	(0.028)	
2. Non Standard Variables	0.016	0.034	-0.004	0.557
	(0.032)	(0.045)	(0.044)	
2.1 Borrower Maximum Rate	0.064 ***	0.083 ***	0.043 ***	0.000
	(0.004)	(0.005)	(0.007)	
2.2 Listing Category	-0.026 ***	-0.048 ***	-0.002	0.000
	(0.003)	(0.005)	(0.005)	
2.3 Member of Group	-0.016 ***	-0.028 ***	-0.003 ***	0.000
	(0.002)	(0.004)	(0.001)	
2.4 Group Leader Reward Rate	-0.015 ***	-0.028 ***	-0.002	0.000
	(0.002)	(0.004)	(0.002)	
2.5 All Other Non Standard Variables	-0.031 ***	-0.042 ***	-0.019 ***	0.025
	(0.005)	(0.008)	(0.006)	
2.6 Other (Residual) Inference	0.040	0.096 **	-0.020	0.066
	(0.032)	(0.045)	(0.044)	

This table decomposes our estimate of inference presented in Table 3, 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. For the sake of brevity, we only present the estimate of inference parameter, gamma, and its decomposition. Column (1) presents the overall combined gamma, 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 combined gammas 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 groups: 1. standard banking variables and 2. nonstandard variables, and further break those down into subgroups 1.1 - 1.6 and 2.1 - 2.6. Please refer to the Appendix for the full decomposition results. 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%.

Appendix: Decomposing Inference, Part I (Standard Banking Variables)

	(1) Gamma			(2) Gamma Low Cat (1-4)			(3) Gamma High Cat (5-7)			(4) Low = High p-value
Standard Banking 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 III, Column (2). The decomposition results are divided into standard banking variables, presented here, and non-standard 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 banking variables are missing in the data (observations with missing standard banking variables account for less than one percent of our sample). For the sake of brevity, we only present the estimate of inference parameter, gamma, and its decomposition. Column (1) presents the overall combined gamma, 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 combined gamma for the lower and higher credit categories is 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: Decomposing Inference, Part II (Non-Standard Variables)

	(1) Gamma			(2) Gamma Low Cat (1-4)			(3) Gamma High Cat (5-7)			(4) Low = High p-value
Non-Standard Variables										
Borrower Maximum 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
Other Residual Inference	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 III, Column (2). The decomposition results are divided into standard banking variables, presented in the previous page, and non-standard 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. All controls except for dummy variables are entered as quadratics. For the sake of brevity, we only present the estimate of inference parameter, gamma, and its decomposition. Column (1) presents the overall combined gamma, 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 combined gamma for the lower and higher credit categories is 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: Sample Listing

PROSPER® Let's bank 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 #208364)

« Back to search results

LISTING SUMMARY

Help

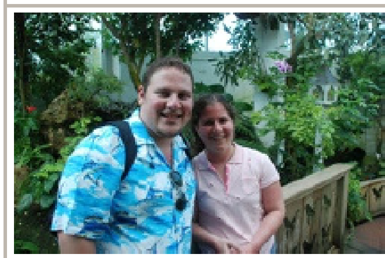

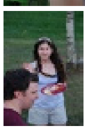

$8,081.00 @ 8.90%

Bid down from 13.99%

Bid Now

(Bidding has ended)

Funding:	100% funded
Bids:	321 bids Ended Listing became a loan
Borrower APR:	9.59%
Mo. payment:	$256.60 (3y loan)

Watch | Email | Report this listing

BORROWER INFO

Help

hs4g2

CANTON, MA

Members and Friends of the Boston Area College Community

(33)

0 friend bids

0 questions & answers

0 friends, 0 verified

1 loan total, 1 active

FORECAST | COMPARE | Help

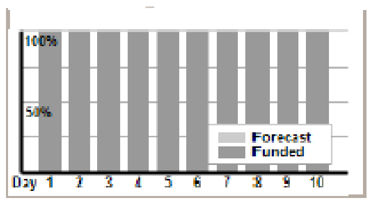

CREDIT PROFILE

Help

A credit grade — Homeownership not verified

10% debt to income ratio

Now delinquent:	0	**First credit line:**	Mar-2001
Amount delinquent:	$0	**Current / open credit lines:**	3 / 3
Delinquencies in last 7y:	0	**Total credit lines:**	4
Public records last 12m / 10y:	0 / 0	**Revolving credit balance:**	$3,531
Inquiries last 6m:	1	**Bankcard utilization:**	29%

Credit and homeownership information provided by Experian.

Employment status:	Full-time employee
Length of status:	1y 0m
Stated income:	$25,000-$49,999
Occupation:	Computer Programmer

Employment and income provided by borrower.

DESCRIPTION

Purpose of loan:

I'm using this loan to pay off my $3,531 credit card bill currently at 14% at a lower interest rate and to buy my girlfriend, Jennifer, an engagement ring costing approx. $4,550. Up until about 3 months ago, I would revolve 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 $550 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 into almost as much debt as I have now to buy Jenn a ring. Then along came prosper. With your help, I'll be able to ask Jenn to marry me sooner than expected and maybe not even be in debt when I do it!

My financial situation:

Currently, I work as a software engineer in Wellesley, 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 (Jenn 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 come May/June since Jenn 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 hurricane. 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: $ 2084

Monthly expenses: $ 1705
Housing: $ 535
Insurance: $ 200
Car expenses: $ 125
Utilities: $ 40
Phone, cable, internet: $ 85
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

Help

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

Help

Bidder / Relationship	Rate	Amount Bid	Winning	Status ▲	Bid Date (PT)
wolfpac79	8.90%	$50.00	$50.00	✓	Oct-09-2007 8:26 AM
user13	8.90%	$50.00	$50.00	✓	Oct-09-2007 8:13 AM
JDLanier	8.90%	$50.00	$50.00	✓	Oct-09-2007 8:11 AM
steamboatgal	8.90%	$100.00	$100.00	✓	Oct-09-2007 8:00 AM
lender1853	8.90%	$100.00	$100.00	✓	Oct-09-2007 7:56 AM
Porsche2	8.90%	$50.00	$50.00	✓	Oct-09-2007 7:52 AM
mmmoney	8.90%	$100.00	$100.00	✓	Oct-09-2007 7:47 AM
mmmoney	8.90%	$100.00	$100.00	✓	Oct-09-2007 7:40 AM
universe	8.90%	$75.00	$75.00	✓	Oct-09-2007 7:35 AM
swissbanker	8.90%	$100.00	$100.00	✓	Oct-09-2007 7:28 AM
OGS_Capital	8.90%	$51.42	$51.42	✓	Oct-09-2007 7:24 AM
moose_spencer	8.90%	$50.00	$50.00	✓	Oct-09-2007 7:16 AM
Orphan2007	8.90%	$50.00	$50.00	✓	Oct-09-2007 6:54 AM
wkk	8.90%	$50.00	$50.00	✓	Oct-09-2007 6:50 AM
LoanChimp	8.90%	$100.00	$100.00	✓	Oct-09-2007 6:29 AM
Gromila18	8.90%	$200.00	$200.00	✓	Oct-09-2007 6:01 AM
Badger1	8.90%	$50.00	$50.00	✓	Oct-09-2007 5:53 AM
Goodthings2you	8.90%	$50.00	$50.00	✓	Oct-09-2007 5:26 AM
stevedavis444	8.90%	$50.00	$50.00	✓	Oct-09-2007 5:13 AM
Deal_Flow	8.90%	$50.00	$50.00	✓	Oct-09-2007 5:06 AM
Curlingman	8.90%	$50.00	$50.00	✓	Oct-09-2007 4:59 AM
5Star	8.90%	$50.00	$50.00	✓	Oct-09-2007 4:49 AM