Inferring Asset Quality:

Determining Borrower Creditworthiness in Peer-to-Peer Lending Markets

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

June 2011

Abstract

To what extent can non-expert market participants collectively infer the underlying quality of an asset? We answer this question by examining a novel peer-to-peer lending market that allows us to estimate both the magnitude of inference and the degree to which it arises from different sources of information. Our methodology takes advantage of the fact that while we as econometricians observe a borrower's exact credit score, market participants (lenders) only see an aggregate credit category. We find that individual lenders are able to infer a third of the variation in creditworthiness that is captured by a borrower's credit score. This represents over two-thirds of what market participants could have possibly inferred given the information available to them. This inference is economically significant and allows lenders to offer a 140-basis-point lower interest rate for borrowers with (unobserved to lenders) better credit scores within a credit category (spanning 40 points in credit score). Lenders infer the most from verifiable financial information. In addition, our methodology shows that lenders also learn from non-standard, subjective, "softer" information, particularly when it is likely to provide credible signals regarding borrower creditworthiness.

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

---

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

I. Introduction

An important role of markets is the aggregation of information that is dispersedly held by its participants (Hayek, 1945). The ability of market prices to aggregate information and reflect the fundamental values of assets is key to market efficiency. A central question is the extent to which market participants are able to use information to correctly infer the fundamental value of an asset. Beyond understanding the degree of such inference, another important question pertains to the weights agents place on different sources of information. For instance, how much do agents rely on verifiable and non-verifiable information? What are the mistakes agents make in inference? Understanding these questions is important in examining the functioning of markets and in designing regulatory policies to address their shortcomings.

However, estimating the magnitude of inference by market participants regarding underlying asset quality is challenging. Decomposing inference from different sources of information is even harder. This is because, in real-world settings, one generally does not have a clean measure of underlying asset quality (fundamental value) and one rarely observes the information set available to each market participant.

This paper uses a unique setting that provides an ideal laboratory to quantify collective inference by individual, non-expert, market participants. It also allows us to estimate the magnitude of inference arising from different sources of information. The setting used is that of an online peer-to-peer lending market, Prosper.com, where borrowers post loan listings and multiple individual lenders bid to fund a portion of the loan at a desired rate. The unique feature of the setting that we exploit is that lenders only see the borrower's aggregate credit category but we as econometricians observe a borrower's exact credit score – a much finer measure of the borrowers' underlying creditworthiness (asset quality). 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.[1]

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. Another novel aspect of the setting is that we observe the entire set of information that is available to lenders as documented in the borrowers'

[1] 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 amount bid equals or exceeds 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 bid.

listings.[2] This allows us to decompose the magnitude of inference of borrower creditworthiness that arises from different sources of information.

We find that, within a given credit category (spanning 40 points in credit score), lenders are able to infer a third of the difference in creditworthiness that is captured by a borrower's exact credit score. This is an economically significant effect because this degree of inference allows lenders to offer a rate that is 140 basis points lower for borrowers at the top of a typical credit category than for borrowers at the bottom of that category. Given that credit score is computed based on proprietary formulas developed by credit bureaus and not all variables that go into the computation are available to lenders, it is by no means obvious that lenders can piece together the information provided in the listing and infer a third of the true credit score.[3] In fact, we estimate that lenders infer as much as 69% of what they could have potentially extracted from the information provided on the Prosper website.

Our results show that lenders exhibit greater inference for borrowers in higher credit categories. The majority of inference is based on hard, verified financial information (henceforth referred to as "standard financial" variables). We also find, especially among the lower credit categories, a high degree of inference from softer, subjective, and non-verified information that is voluntarily posted by borrowers (pictures, listing text, etc.), henceforth referred to as "non-standard" variables. 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," that is, the variation of interest rates with the exact credit score that remains after controlling for a very flexible functional form of coded information.

Interestingly, 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

---

[2] The borrower listing contains hard, verified information obtained from the credit rating agency (past defaults, number of credit lines, etc.), as well as soft, subjective, non-verified information (picture, description, etc.) that borrowers voluntarily provide.

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

alternate funding options. Our results suggest that, consistent with the models of cheap talk, individuals pay greater attention to the more credible signals sent.[4]

Given our particular setting, a potential concern is 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 private portion of "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 across states, and group affiliation of borrowers.

While we estimate the lenders' inference of creditworthiness using credit score, there may be dimensions of creditworthiness that are not captured by credit score, i.e. credit score is not the only measure of underlying asset quality. Moreover, in decomposing inference along various sources of information, since credit score is ostensibly based on "hard" factors like past borrower behaviour and default history, we may underestimate the importance of soft information.

Although we cannot apply our precise credit-score methodology to dimensions of creditworthiness that are not observed, we can shed light on inference along other dimensions directly by examining the relationship between the interest set by lenders and a borrower's *ex-post* likelihood of default. The idea behind this test is that if lenders are indeed able to infer a borrower's underlying (yet unobserved) creditworthiness, the interest rate they set should predict future default.[5] Moreover, if such a prediction were based on dimensions other than credit score, the interest rate should predict default even if we directly control for the borrower's actual (and unobserved to the lender) credit score. We indeed find that interest rates strongly predict future borrower default and

[4] 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.

[5] One may be concerned that interest rates could directly influence the likelihood of default (Stiglitz and Weiss, 1981). To address this concern, we examine whether interest rate has a causal effect on default using credit-category borders as the exogenous instruments (there is a large shift in interest rates for nearly identical borrowers at the credit-category boundaries). We do not find a significant increase in likelihood of default across credit-category boundaries, suggesting that defaults are unlikely to be caused by higher interest rates set by lenders.

that this relationship remains as strong even when we control for a flexible functional form of the borrower's credit score. Moreover, the relationship also holds if we control for hard and codified soft information, implying that lenders infer creditworthiness even from the residual, uncoded, soft information.

Our results confirm that lenders also infer borrower creditworthiness along dimensions of creditworthiness beyond credit score, and that soft information is an important source of this inference. We also find that the controls for hard and codified soft information are jointly significant in predicting default even after we control for the interest rate. This implies that lenders did not incorporate all available information on the likelihood of default in the interest rate they set. Consistent with our previous results, this is evidence (assuming lenders try to maximize their expected returns) that, while lenders make substantial inference of the underlying asset quality (borrower creditworthiness), their inference is by no means complete.

Our paper contributes to the literature that examines information aggregation, inference and learning in markets. There are several theoretical papers that focus on information aggregation through prices (Grossman, 1976; Townsend, 1978; Grossman and Stiglitz, 1980; and Vives, 1993, 1995). Another strand of literature focuses on learning in decentralized markets (Wolinsky, 1990; Duffie and Manso, 2007; and Duffie, Giroux, and Manso, 2009). On the empirical front, Biais, Hillion, and Spatt (1999) and Davies (2003) examine learning in the pre-opening period in equity markets. There are also several experimental papers that examine price formation in asset markets (Plott and Sunder, 1988; Forsythe and Lundholm, 1990; Bronfman et al., 1996; Cao, Ghysels, and Hatheway, 1998; and Hanson et al., 2006). Our paper adds to this literature by quantifying the magnitude and sources of inference in a market composed primarily of non-expert or unsophisticated participants.

In addition, our paper examines the extent and nature of *mistakes* in inference. For example, we find that lenders correctly infer the most from aspects of non-standard borrower information that are likely to serve as credible signals of creditworthiness. However, lenders also offer lower interest rates to borrowers who post personal pictures and friend endorsements, even though pictures and endorsements are empirically uncorrelated with underlying borrower creditworthiness. Our identification of mistakes in inference contributes to the literature on the possibility of manipulation in asset markets, e.g. through the provision of false unverifiable soft information by borrowers (Camerer, 1998 and Strumpf and Rhode, 2003). Furthermore, by decomposing the extent of inference from different sources of information, the paper sheds light on the relevance of different types of information in markets (Crawford and Sobel, 1982; Farrell and Rabin, 1996; Forsythe et al., 1999; Berger et al., 2002; and Petersen, 2004).

A related strand of literature examines prediction markets. Small election markets, like the Iowa electronic markets, and event markets, which rely on aggregating information from a relatively small number of non-expert individuals, seem to provide reasonably accurate predictions (Wolfers and Zitzewitz, 2004). Our results show that, despite not being financial experts, individual lenders in peer-to-peer markets can infer a significant fraction of underlying borrower creditworthiness using both soft and hard information. Moreover, given these markets' ability to make valuable inferences and their non-collateral-based lending structure, they may offer a potential capital source for small borrowers who may otherwise be limited to costly sources of finance such as payday lenders and credit-card debt.

Our work also complements the recent literature that specifically examines lending in peer-to-peer markets. Pope and Sydnor (2010), Ravina (2008), and Theseira (2008) examine whether these markets display discrimination based on personal attributes such as race and physical appearance. Peer-to-peer markets may also make better use of social network information. While Freedman and Jin (2008) find evidence of adverse selection due to informational problems faced by lenders in Prosper, they 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. In contrast to these papers that carefully document lending behavior in peer-to-peer markets, our focus is instead on estimating the magnitude of inference by market participants regarding underlying asset quality and decomposing the extent of inference along different information sources. We are able to do so by developing a methodology that exploits our unique proprietary data containing exact credit scores.

While we focus on one particular market, it provides a compelling setting in which to study inference more broadly. Our results suggest that even in markets with non-expert participants, while individuals make mistakes, the magnitude of inference from hard and soft information regarding asset quality is high and a significant fraction of the fundamental value is incorporated into prices.

II. Context and Data

A. Context

A wide range of markets (stock markets, prediction markets and more recently peer-to-peer markets) perform the role of aggregating information using non-expert individuals. We examine inference of the underlying value of assets, i.e. loans, in peer-to-peer markets. The marketplace model of online peer-to-peer lending enables individual lenders to locate individual borrowers and vice-versa. There has been rapid growth in online peer-to-peer markets 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.[6] 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 $192 million in loans and currently has 940,000 members.

All Prosper loans are personal, three-year fixed-rate, unsecured loans. Borrowers request loans by creating public listings on the Prosper.com website. They 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 good credit risks, and their income and expenditure flows. Some borrowers also post optional pictures of themselves or of themes related to their loan purpose. The third listing component, credit information, contains verified hard information obtained by Prosper through a credit check. The credit information section contains information on each borrower's delinquencies, credit lines, home ownership status, debt, inquiries, public records, and income. A sample listing is provided in the Appendix.

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.[7]

---

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

[7] 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 below. For consistency of results, we restrict

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 borrower, the interest rate is bid down. Lenders with lower reservation interest rates are given priority in the bidding hierarchy. The final interest rate is determined by the highest reservation interest rate among the set of lenders that successfully bid for the loan.

After the listing is funded and approved by the borrower, the borrower begins to make monthly payments that are divided across lenders 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

our sample to the period after February 12, 2007. However, results are robust to using the sample from before February 12, 2007 (see Table 4).

borrower.[8] Our sample contains all listings posted between February 12, 2007 and October 16, 2008.[9] 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 financial variables and non-standard variables. The standard financial variables include hard, verified financial information from the borrower's credit report. As expected, funded listings tend to have borrowers with better credit scores – in particular, funded listings tend to have far fewer "high risk" borrowers (those in the lowest credit categories). Still, 14.3% of the funded listings default within one year, with the one-year default rate ranging from 5.4% in credit category AA to 30.6% in credit category HR. Among the universe of listings, the average loan amount requested is $8015. The average maximum interest rate borrowers are willing to pay is 21%. Lower 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 representing individuals with better credit variables are much more likely to be funded. 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

---

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

[9] Prosper entered a "quiet period" in October 2008, in which it ceased making new loans in anticipation of an SEC cease-and-desist procedure. Prosper emerged from the quiet period in July 2009 using a new system of classifying prospective borrowers into credit categories. We therefore cannot use data on loans originating after October 2008. We also use data from May 2006 to February 12, 2007 as part of a robustness check. However, we exclude data from this period in our baseline sample because the credit-category boundaries changed on February 12, 2007. See Section 2, Part A for more details.

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

The underlying asset in the context of our setting is the individual loan. Ideally one would like to know the true creditworthiness of the individual borrower i.e., ex-ante probability of default. However since we can't observe the true creditworthiness of an individual we will use the credit score as a measure of individual creditworthiness. We will return later to a discussion of other dimensions of underlying asset quality (creditworthiness). Our empirical strategy exploits the fact that a proxy for underlying asset quality -credit score- is 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 collectively able to infer differences in asset quality/creditworthiness across borrowers (in the same credit category) from other information provided on the website.[10] Moreover, given that lenders do observe credit categories, we can quantify the magnitude of inference 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 without coding it. The 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.

A. Estimating Overall Inference

[10] 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.

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* increases by exactly one within each credit category). The fraction of this total change that lenders infer from information other than the credit-score categories is given by $\gamma = \beta / (\alpha + \beta)$. Thus, a γ of zero means that lenders are not at all able to infer creditworthiness from information other than the credit-score categories, 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 a possible 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 $\delta_k = (1-\lambda_k)\alpha_k + \lambda_{k+1}\alpha_{k+1} + \beta_k$.[11] 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:

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

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 that is constant everywhere else. The coefficient α_k measures the jump in interest rate at the credit-score boundary between credit categories k-1 and k, the coefficient β_k measures the change in interest rate within category k, and ε_i is the error term. Formally, we define $I_k^{Cum}(CreditScore_i)$ as an indicator variable that equals one if borrower i is in credit category k or higher:

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

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:

[11] 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 $(1-\lambda_7)\alpha_7$.

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

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 statistically significant.

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

When we estimate equation (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.[12] This censoring problem would bias our estimates of inference if we estimate equation (2) using ordinary least squares. Instead, we estimate equation (2) as a censored regression with the censoring occurring at the borrower maximum rate specified by each listing. A censored regression, which is a generalization of the Tobit model, rests on the

[12] 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.

implicit assumption that listings that were not funded would have been funded at some interest rate larger than the observed borrower maximum rate. If the error term has a homoskedastic and normal distribution, estimates from a censored regression will yield consistent estimates of the parameters determining the interest rates that lenders require to fund a listing.

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

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:

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

where x_i^m denotes the m^{th} quantitative variable in the listing of borrower i and φ^m denotes the corresponding regression coefficient.[13] In regression (5), the fitted interest rate can change with

[13] 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

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 score within credit categories by *FracGap*, this residual correlation is measured by β_k^{Resid}. Second, the fitted interest rate may vary within a credit category because (i) listings with higher values of *FracGap* may have different observable characteristics and (ii) the interest rate responds to these characteristics. We measure component (i) – the degree to which observable characteristic x^m varies with *FracGap* – by running a regression of the observations of x^m within category k on $FracGap_k$ and a constant term. We denote the coefficient on $FracGap_k$ in this bivariate regression by θ_k^m. We measure component (ii) – the degree to which the interest rate responds to characteristic x^m – by the regression coefficient φ^m. The total contribution of variable x^m to the relationship between *FracGap* and interest rate within category k is given by the product of these two components: $\beta_k^m \equiv \theta_k^m \varphi^m$.

We decompose our original estimate β_k from the regression without the controls for quantified listing characteristics (regression 2) as follows:[14]

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

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 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 score within credit-score categories after controlling for all quantitative information. To ensure that β_k^{Resid} reflects qualitative information rather than

---

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.

[14] 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.

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 financial 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).

C. Examining Default

While the above methodology provides us with both the magnitude and source of inference, it does so along the credit-score dimension of creditworthiness. While credit score is a plausible proxy for asset quality (creditworthiness) one may ask how inference looks along other dimensions of creditworthiness. However, this requires the econometrician observing a direct measure of creditworthiness (i.e., the *ex-ante* likelihood of default), and that is not possible. While one may be tempted to use actual *ex-post* realizations of default as a proxy in order to estimate the degree of inference, this is not feasible, as we explain below.

Nevertheless, we can still use default outcomes to show whether credit score is indeed a salient dimension of creditworthiness, and whether market participants infer dimensions of creditworthiness beyond those captured by credit score. In addition, we can use default outcomes to estimate the highest gamma that could have been achieved if lenders had perfectly used all hard and coded soft information available in the listing sheet.

To test whether exact credit score is predictive of default, we use a modified version of equation (2). 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.

To investigate whether lenders are able to infer creditworthiness along dimensions not captured by credit score, we note that examining the relationship between *ex-post* default rates (a direct function of *ex-ante* and unobserved creditworthiness) and the interest rate set by lenders provides evidence for whether there is such broader inference (see Appendix I for details). The idea behind this test is that if interest rates predict *ex-post* default, then lenders are indeed able to price creditworthiness. More importantly, if this relationship continues to hold even after we control for credit score, it shows that lenders also infer dimensions of creditworthiness beyond what is captured by credit score. To address a potential concern that this relationship may be partly due to better (worse) borrowers setting lower (higher) borrower maximum rates (which in turn bound the interest rate the loan is funded at), we control for borrower maximum rate in all the regressions. To address a second potential concern that interest rates may directly impact default, we use an instrumental variables strategy that is explained in detail in Section IV, Part E.

To estimate the highest gamma that could have been achieved if lenders had perfectly used all hard and coded soft information available in the listing sheet, we first regress default realizations on a very flexible functional form of all financial and non-standard variables.[15] We use this regression to predict a default probability for each loan. Because these predicted default probabilities are based on information that was observable to the lenders, lenders could also have made these predictions themselves (if they had used all coded listing content optimally). Next, we rerun our baseline specification (equation 2) using the predicted default probability as the dependent variable instead of the interest rate. The gamma from this baseline specification with predicted default as the dependent variable tells us what fraction of predictable default along the credit score dimension occurs *within* credit categories. Hence, we consider it a benchmark for our baseline estimates because it tells us what gamma the lenders could have achieved if they had adjusted their interest rate for the default risk that they could have predicted based upon coded listing content.

We see the analysis based on *ex-post* default behavior as complementary to our main analysis based on the relationship between interest rates and exact credit score within credit categories. However, as we show in Appendix I, the main drawback of the analysis based on *ex-post* default behavior is that it does not allow us to quantify the degree of inference using an intuitive metric

[15] Specifically, we use a flexible set of 215 controls for standard financial and non-standard variables as described in the Appendix, which are further interacted with the seven credit category dummies.

(such as the γ measure in our main analysis). A second drawback is that we can only measure default status at 12 months after loan origination for our whole sample rather than default status at maturity because most of the loans in our sample have not yet reached their maturity date. Hence, the results rest on the assumption that default hazards in the second and third year of the loan respond to the same variables as the one-year default rate. Nevertheless, we present results from this analysis because, even though it cannot provide the *magnitude* of inference along non-credit score dimensions, it does show whether such broader inference *exists*.

IV. Results

We now present the results. We first show that credit score is indeed a proxy for underlying asset quality (creditworthiness). Next, we 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. We finish by exploring inference of creditworthiness along dimensions that are not captured by credit score.

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 the one-year default rate in our sample.[16]

We find that credit score indeed predicts the likelihood of default. Column (1) shows that, for every 40-point increase in the credit score, the likelihood of default within the first year decreases by 3.3 percentage points. We use 40-point intervals for ease of comparison given that Prosper defines categories based on 40-point intervals. Since lenders observe credit categories, a related

[16] We measure one-year defaults by the loan being four or more months late by the sixteenth month since origination of the loan (one year plus four months). We use four months late as our cutoff because Prosper starts to sell loans off to collection agencies after they are four months late. We restrict our attention to the one-year default rate because this measure is available for our entire sample of loans. In contrast, two-year and three-year default rates can only be observed for a fraction of the observations because most loans in our sample have not yet reached maturity. For the 42% of loans that were originated sufficiently early that we can both observe their one-year and two-year default rate, we find that the two-year default rate (26.0%) is roughly double that of the one-year default rate (14.6%), indicating that the default hazard is approximately constant overall. However, this average masks that the default hazard is falling over time for lower credit categories and rising over time for the better credit categories.

question is whether credit score is still predictive of the one-year 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 3.0 percentage point lower one-year 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, that is, 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 0.931. We cannot reject that gamma is 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 the one-year 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 0.841, 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.

In Column (4), we examine whether credit score is a sufficient statistic for default behavior by adding controls for our standard financial variables and our non-standard variables. As we do throughout the paper, we use a flexible functional form for these controls by entering each of these variables as a quadratic that is interacted with a full set of credit-category dummies. We find that both the standard financial variables and the non-standard variables are highly predictive of the one-year default rate, and that, as a result, the adjusted R^2 of the regression more than triples compared to Column (3). This finding implies that credit score does not capture all dimensions of creditworthiness, which will serve as a caveat to our later results on the degree of inference, γ. In addition, the inclusion of these flexible controls does not strongly impact the β coefficients (they remain jointly significant at 5.6%). This also supports our later results showing that a perfect inference gamma value of one is not attainable (i.e., the hard and coded soft information provided on Prosper is not sufficient to provide all the relevant information in exact credit score).

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: To what extent are lenders 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 reasonably 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 website in order to infer creditworthiness. A benchmark for our inference estimate of 0.33 is the amount of inference that could have been attained if lenders had optimally used all coded information from the listing content. This benchmark, estimated using the method described in Section III.C, is 0.42. Because the benchmark is only for coded listing content and our inference estimate may also be partly based on non-quantifiable listing content, a fairer comparison is to relate the inference based solely on coded content with our benchmark. As we will see later (Table 5), we estimate an inference of 0.29 from coded content only. Thus, lenders were able to infer 0.29/0.42 = 69% of what was attainable given the information provided on the Prosper website, which strikes us as a remarkable achievement.

To understand the economic significance of this inference result, note that the αs and βs sum to 39 percentage points. In other 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.[17] 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 (correctly) 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. The (partial) inference is economically quite meaningful.

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 therefore to 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 (2010) present evidence consistent with such adverse selection. Adverse selection, however, does not bias our estimates since we observe exact credit score and our estimator does not depend on the density of observations by credit score within a category.

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

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

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).[18] 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 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

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

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 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 the 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?

We find that lenders are able to infer about a one-third of the relevant information content of credit score. What sources of information are lenders using for inference? 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 financial variables and non-standard variables (variables chosen by the borrower). Generally speaking, standard financial 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 financial variables are readily coded, and we provided the details and summary statistics of variables included in this category in Table 1.[19] 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).[20] The last column presents the p-value from a test of whether the combined gamma is equal across the low and high categories.

[19] For the sake of brevity, Table 1 does 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 relevant specifications in Table 2, Table 5, Table 6, and the Appendix tables.
[20] We chose this categorization as it roughly provides us with an equal number of loans in both 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.[21] 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 financial and non-standard financial variables make to the total combined gamma. We report both the aggregate gammas for these sub-categories and the gammas for the variables within each sub-category that show the largest (in magnitude) inference. The Appendix Table presents the individual gammas for all the variables separately.

Reading down the first column in Table 5, a first impression is that most of the inference appears to come from standard financial variables. However, part of the reason there is less aggregate inference drawn from non-standard financial variables is that gamma is negative for some of these variables and thus masks the positive contribution to inference of other non-standard financial variables. We revisit this issue of negative contributions to inference later in this section. For now, we want to caution against concluding that only hard information matters for inference. In fact, as we note below, close to half of the inference in lower credit categories comes from *uncoded* subjective listing content.

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 the aggregate, lenders do seem to learn more from standard financial variables, which are verifiable 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 less subject to the possible "cheap talk" concerns of voluntarily posted and unverified information. Moreover, it is possible that the standard financial 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. However, given that credit score is computed by credit bureaus using proprietary technology and also that all the variables that go into computing the credit score are not available to the lenders, it is interesting that lenders can infer a significant fraction of the credit score.

[21] 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.

Second, in examining which variables are used by lenders to draw inference from the standard financial variables, we find that most of the inference is indeed 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, consider 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 borrower attributes/choices are related to their credit score). Thus, inference may increase for a variable across credit categories if either (or both) of the coefficients 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 variables is relatively more important for lower credit categories (14% of overall inference), especially when we consider some of the specific variables (such as borrower maximum rate) in this category.[22] This may not be surprising if one believes that variation in standard financial information is less revealing when attempting to distinguish between low-quality borrowers (e.g., differences between someone being delinquent ten times versus twelve may be less revealing than zero versus

[22] Note that since credit score is likely to be more directly influenced by hard information and standard financial 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.

two times). This leaves more room in the lower quality assets (low credit categories) to rely on non-traditional methods of inference. However, as evidenced by several variables that show negative inference, this also leaves more room for incorrect inferences being drawn by market participants.

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 credit categories (33.9%) than for 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. The borrower maximum rate is likely to serve as a credible signal of creditworthiness because lender bids of the loan interest rate cannot start above the borrower maximum rate. As one would expect, borrowers that choose a lower borrower maximum rate have a lower probability of their listing being funded, even conditional on credit score (results not reported). Since more creditworthy borrowers likely have better "outside" borrowing options, it is less costly for them, relative to less creditworthy borrowers, to post a lower borrower maximum rate. While establishing this as a separating equilibrium requires 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. This non-standard variable is important for inference even when looking across all (standard and non-standard) variables; as Column (1) shows, it is the second most important inference variable among the 40 that we examined.

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 also find similar results (not reported) when we estimate residual inference for sub-samples where we expect softer information to be of more importance: listings with images; listings where the borrower has at least one delinquency recorded in the last seven years; and listings where the number of characters in the listing text exceeds 900 (the median among funded loans). We also 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 functional form of the control variables. This suggests that subjective listing content plays an important role for inference especially for weaker borrowers (i.e., low quality assets).

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. This would be the case if, for a given credit score, larger loan amounts increase default likelihood.[23] 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 challenges in estimating inference and potential biases towards underestimating it, particularly from the non-standard information in loan listings.

E. Lender Inference along Dimensions of Creditworthiness not captured by Credit Score

While it is encouraging that lenders are able to infer a substantial part of the dimension of creditworthiness that is captured by credit score, a skeptic may wonder whether such inference holds along dimensions of loan asset quality other than credit score.

Table 6 presents the results of the analysis based on *ex-post* default behavior. It establishes that lenders are able to make inferences on creditworthiness along dimensions not captured by credit score, that such inference is also based on non-coded qualitative information (not just hard information), and that, consistent with our previous results, the inference is incomplete. In all columns, the dependent variable is the one-year default status (either 0 or 1), and the independent

[23] 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, a higher amount requested likely predicts higher credit score because creditworthy individuals may believe that they can ask for larger amounts (which is generally the case in our data). On the other hand, all else equal, one expects that those who borrow more are more likely to default because they face larger repayment obligations. Thus, amount requested affects interest rates both through the credit-score dimension of creditworthiness and through the loan-size dimension. 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.

variable of interest is a functional form, $1/(1+r)$, of the interest rate (refer to Appendix I for the details of why we use this form).

Column (1) runs a baseline specification and shows that, as expected, there is a strong negative relationship between the likelihood of *ex-post* default and $1/(1+r)$. In effect, borrowers who are charged a higher interest rate by lenders are also more likely to default. We will show below, based on results from Table 7, that we can rule out reverse causality, i.e., that this coefficient is driven by the probability of default directly increasing in response to higher interest rates. Therefore, this result shows that lenders correctly charge a higher rate for borrowers with a higher *ex-ante* likelihood of default. We should caution here that the magnitude of the coefficient on interest rates *cannot* be interpreted as a measure of the degree of inference in the same way as our gamma parameter as described in Section B. Appendix I provides a detailed examination of this issue.[24]

Columns (2) and (3) then show that there is indeed inference on creditworthiness beyond the credit-score dimension. In other words, interest rate continues to predict default rate even when we flexibly control for credit score. We test this by including credit-category dummies as control variables (Column 2), or even more demandingly, by adding credit score dummies and splines in the exact credit score (i.e., the variables $FracGap_k$ in equation 2) to the regression (Column (3)). We obtain similar results if we control for fourth-order polynomials in credit score (regression not shown).

It is noteworthy that, in Columns (2) and (3), the credit variables (dummies and splines) are jointly significant. Both allude to and are consistent with our previous results on imperfect lender inference. The credit dummies (observed by lenders) being significant implies that the information of these dummies for the likelihood of default was not perfectly incorporated into the interest rate charged by lenders. The coefficients on the spline (unobserved by lenders) being significant implies that lenders did not fully infer the relevant information content of exact credit score from the information provided in the listing, i.e., $\gamma < 1$.

[24] The basic insight (see Appendix I for details) is that we observe only a realization of default rather than the true *ex-ante* default likelihood of a given loan (referred to as ω in the Appendix). Meanwhile, the magnitude of the coefficient of *ex-post* default on the interest rate depends both on the degree of inference (τ in the Appendix) and the relative signal-to-noise ratio in how interest rates are set. Without knowing the latter, we cannot measure the magnitude of the inference. Thus, while a positive and relatively large coefficient rejects the null of no inference, its magnitude cannot be readily translated into a degree of inference. If we assume perfect inference (τ=1), no noise in how interest rates are set ($\sigma_\eta^2=0$), a risk-free interest rate of 5%, and that 23% of the principal and interest due is recovered after default (based on our rough calculations of actual recovery rates), then equation 17 of Appendix I indicates that the coefficient on $1/(1+r)$ should be around 1.4. The fact that the coefficient is higher indicates that inference is incomplete (τ<1) and/or that lenders expected a higher recovery rate than actually occurred.

Column (4) goes a step further by adding extremely flexible controls for standard financial variables and coded non-standard variables. While the magnitude falls slightly,[25] the highly significant coefficient on the interest rate in Column (4) shows that lenders were able to predict default beyond what is possible based on the standard financial variables and non-standard variables. The fact that, even with controls for hard and coded soft information, the interest rate remains significant in explaining default implies that lenders use subjective residual (non-coded) listing content to assess borrower creditworthiness. Because a flexible functional form of credit score is included in the set of regressors, our results also show that subjective (non-coded) content is indeed very useful in inference of the dimensions of creditworthiness not captured by credit score.[26]

One concern is that higher interest rates by themselves could lead to borrower default by increasing the burden on borrowers (Stiglitz and Weiss, 1981). We address this by examining whether interest rate has a causal effect on default using credit-category borders as instruments. The intuition for the instrument is that, at the exogenously defined borders (AA, A, B, etc.), there is a sharp jump in interest rates (see Figure 3 for graphical evidence of the jump in interest rates at credit-category borders). Meanwhile, credit score is a continuous function of borrower quality, so borrowers with credit scores immediately above and below a credit-category border have very similar creditworthiness. In Table 7, Columns (1) and (2) show that interest rate when instrumented using credit-category dummies does not significantly affect the likelihood of default. This finding accords with the results in Table 2 and Table 3, where we find that, while interest rates change sharply as we jump across credit boundaries (alpha), there is no significant change in the default probability as we move across credit boundaries. The coefficient on $1/(1+r)$ in Column (2) of Table 7 is much smaller in magnitude than the point estimates in Table 6, and its 95-percent confidence interval excludes three of the four point estimates in Table 6, confirming that the results in Table 6 cannot be driven by the interest rate causing default behavior.

[25] The extra controls in Column (4) cause the magnitude of the coefficient on the functional form of interest rate to drop somewhat relative to Column (3). This is expected and has no implications for the degree of inference made by lenders. Adding controls decreases the signal-to-noise ratio (the term inside the square brackets in equation 17 in Appendix I) because controlling for factors that determine borrower quality (whether hard or soft) reduces overall variation that arises due to true (unobserved) borrower quality relative to that which arises due to noise.

[26] Both standard and non-standard variables are jointly statistically significant in Column (4) and substantially increase the predictability of the default rate, as evidenced by the near doubling of the adjusted R^2. These findings again confirm our results on imperfect inference, i.e., lenders fail to incorporate all relevant information in the listing into the interest rate they offer. Moreover, the joint insignificance of the spline in exact credit score confirms that exact credit score no longer significantly predicts the default probability once all the information from the listing is taken into account.

V. Conclusion

We measure the extent to which non-expert individuals can collectively infer asset quality using a unique natural experiment in which we, as econometricians, observe underlying proxies for asset quality as well as the full set of information available to market participants. Our methodology allows us to estimate both mistakes in inference and sources of inference (from both standard verified information as well as non-standard or softer information).

We find that market participants infer a third of the underlying asset quality (specifically, the variation in creditworthiness that is captured by a borrower's credit score). We estimate that this represents over two-thirds of what they could have possibly inferred given the information available to them. Inference allows lenders to offer an interest rate that is 140 basis points lower for borrowers at the top of a typical 40-point credit category than for borrowers at the bottom of that category. Our results also show that, consistent with models of cheap talk, market participants primarily rely on verifiable information and credible signals for inference. These results should however be tempered with the realization that lenders in these markets do make incorrect inferences as evidenced in our results both by imperfect and at times negative learning.

While our results highlight the ability of markets to aggregate information even when most of its participants are unsophisticated, it raises a range of further questions. How do the composition and characteristics of market participants affect information aggregation? Does the sequencing of when market participants engage matter? How do such (non-expert) decentralized markets compare with and complement smaller/expert-driven and more centralized markets? Addressing such questions is a promising direction for future enquiry.

References:

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

Berger, Allen N., Nathan H. Miller, Mitchell A. Petersen, Raghuram G. Rajan, and Jeremy C. Stein, 2002, "Does Function follow Organizational Form? Evidence from the Lending Practices of Large and Small Banks," *Journal of Financial Economics*, 76(2): 237-269.

Biais, Bruno, Pierre Hillion, and Chester Spatt, 1999, "Price Discovery and Learning During the Preopening Period in the Paris Bourse," *Journal of Political Economy*, 107(6): 1218-1248.

Bronfman, Corinne, Kevin McCabe, David Porter, Stephen Rassenti, and Vernon Smith, 1996, "An Experimental Examination of the Walrasian 'Tatonnement' Mechanism," *Rand Journal of Economics,* 27: 681–99.

Camerer, Colin, 1998, "Can Asset Markets Be Manipulated? A Field Experiment with Racetrack Betting," *Journal of Political Economy,* 106: 457–82.

Cao, Charles, Eric Ghysels, and Frank Hatheway, 1998, "Why Is the Bid Price Greater than the Ask? Price Discovery during the Nasdaq Pre-opening," Working paper, Pennsylvania State University.

Crawford, Vincent, and Joel Sobel, 1982, "Strategic Information Transmission," *Econometrica*, 50(6): 1431-1451.

Davies, Ryan J., 2003, "The Toronto Stock Exchange Preopening Session," *Journal of Financial Markets*, 6(4): 491-516.

Duffie, Darrell, and Gustavo Manso, 2007, "Information Percolation in Large Markets," *American Economic Review Papers and Proceedings*, 97: 203-209.

Duffie, Darrell, Semyon Malamud, and Gustavo Manso, 2009, "Information Percolation with Equilibrium Search Dynamics," Forthcoming *Econometrica.*

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

Farrell, Joseph, and Rabin Matthew, 1996, "Cheap Talk," *Journal of Economic Perspectives*, 10(3): 103-118.

Forsythe, Robert, and Russell Lundholm, 1990, "Information Aggregation in an Experimental Market," *Econometrica*, 58: 309–347.

Freedman, Seth, and Ginger Z. Jin, 2010, "Learning by Doing with Asymmetric Information: Evidence from Prosper.com," Working paper, University of Maryland.

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

Grossman, Sanford, 1976, "On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information," *Journal of Finance*, 31: 573-585.

Grossman, Sanford, and Joseph Stiglitz, 1980, "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70: 393-408.

Hanson, Robin, Ryan Oprea, and David Porter, 2006, "Information Aggregation and Manipulation in an Experimental Market," *Journal of Economic Behavior & Organization*, 60: 449–459.

Hayek, Friedrich, 1945, "The Use of Knowledge in Society," *American Economic Review*, 35: 519-530.

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.

Petersen, Mitchell, 2004, "Information: Hard and Soft," Working paper, Northwestern University.

Plott, Charles R., and Shyam Sunder, 1988 "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica* 56: 1085–1118.

Pope, Devin G., and Justin Sydnor, 2010, "What's in a Picture? Evidence of Discrimination from Prosper.com," *Journal of Human Resources*, forthcoming.

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.

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.

Strumpf, Koleman, and Paul Rhode, 2003, "Historical Presidential Betting Markets," *Journal of Economic Perspectives*, 18: 127–141.

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

Townsend, Robert M., 1978, "Market Anticipations, Rational Expectations, and Bayesian Analysis," *International Economic Review*, 19: 481-94.

Vives, Xavier, 1993, "How Fast do Rational Agents Learn?" *Review of Economic Studies*, 60: 329-347.

Vives, Xavier, 1995, "The Speed of Information Revelation in a Financial Market Mechanism," *Journal of Economic Theory,* 67: 178–204.

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

Wolinsky, Asher, 1990, "Information Revelation in a Market with Pairwise Meetings," *Econometrica*, 58: 1-23.

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

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

To establish the determinants of the slope of *ex-post* defaults with respect to interest rates, we develop a stylized model that takes lenders to be approximately risk neutral with respect to idiosyncratic default risks over small bets (typically around \$100) in the Prosper online market.[27] Lenders thus require an interest rate r such that their expected returns are equal to the "risk-free" rate of return, r_f:

$$(1 + r_f) = \mathrm{E}[(1+r)(1-(1-\kappa)W)], \tag{1}$$

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

$$\mathrm{E}[W \mid r] = \frac{1}{1-\kappa} - \frac{1+r_f}{1-\kappa}(1+r)^{-1}. \tag{2}$$

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

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

$s = \omega$ with probability π, and
$s \sim f(\omega)$ with probability $1-\pi$,

[27] The risk premium for *systematic* default risk is included in r_f.

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

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

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

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

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

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

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

$$s(r) = 1/(\pi\,(1-\kappa)) - (1 + r_f)\,/(\pi\,(1-\kappa)(1+r)) - \omega_m\,(1-\pi)/\pi. \tag{6}$$

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Substituting (11) into (12) yields:

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

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

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

$$(1+r)^{-1} = (1+r^*)^{-1} + \eta \tag{14}$$

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

$$(1+r^*)^{-1} = 1/(1 + r_f) - (\pi\,\tau\,s + (1-\pi\,\tau)\,\omega_m)(1-\kappa)/(1 + r_f) \tag{15}$$

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

$$\sigma^2_{(1+r^*)^{-1}} = \left(\pi\tau(1-\kappa)/(1+r_f)\right)^2 \sigma_\omega^2 \tag{16}$$

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

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

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

Figure 1: Stylized Relationship between Interest Rate and Credit Score

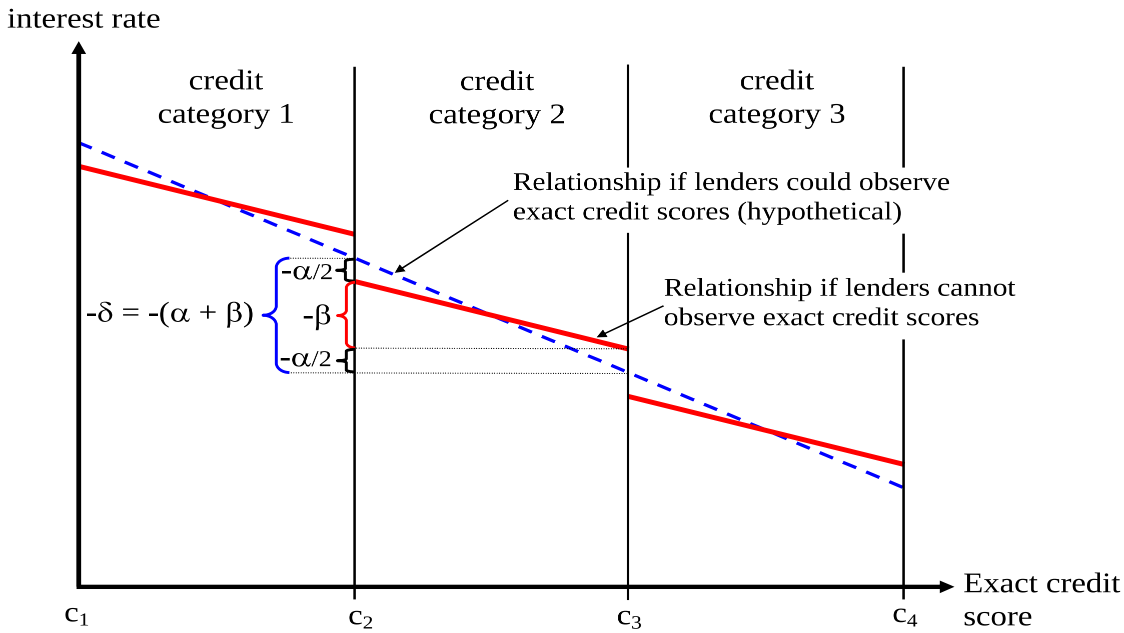

This figure shows the stylized hypothesized relationship between a borrower's credit score and the market interest rate on her (funded) loan.

Figure 2: Relationship between Interest Rate and Credit Score

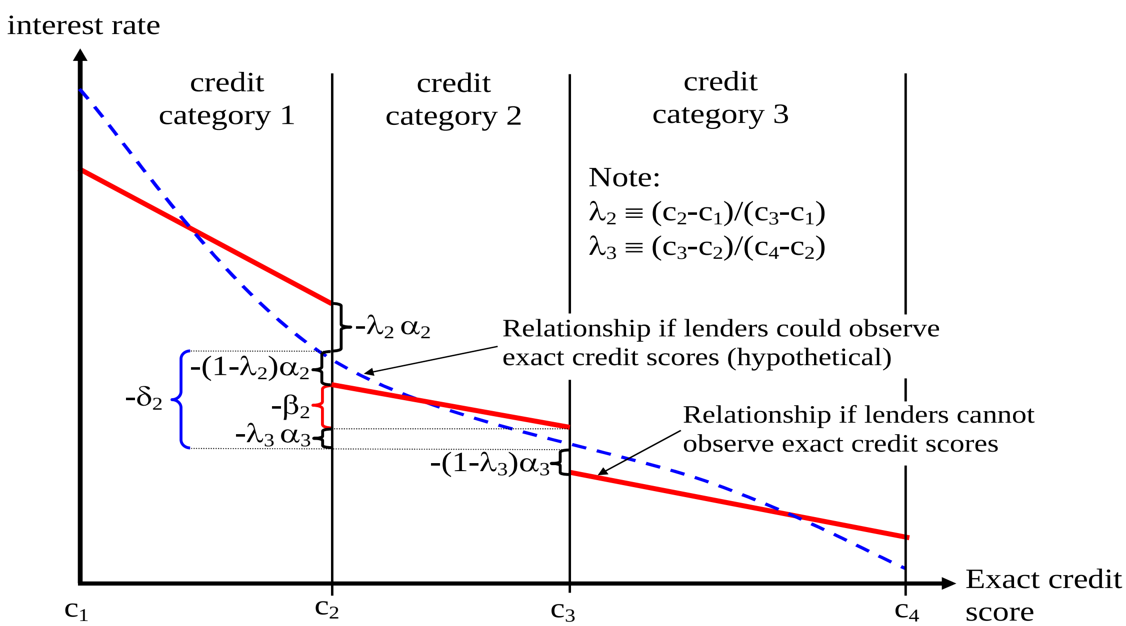

This figure shows a more realistic hypothesized relationship between a borrower's credit score and the market interest rate on her (funded) loan.

Figure 3: Market Interest Rate and Credit Scores

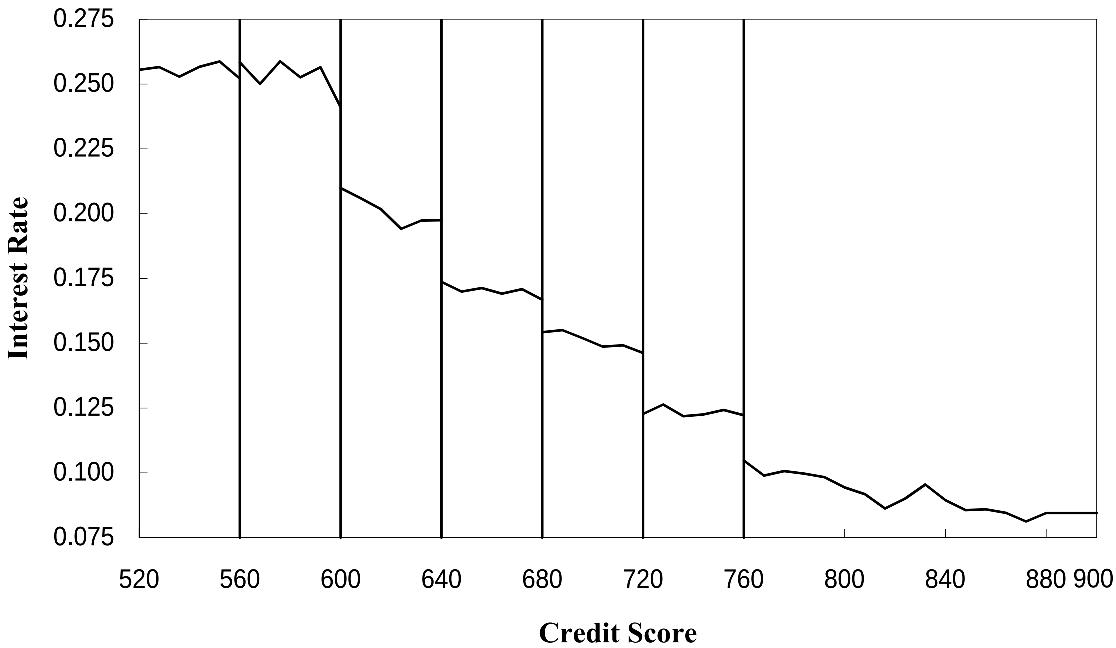

This figure shows the "raw" relationship between a borrower's credit score and the interest rate on her funded loan. Each point in the graph plots the average interest rate over an eight-point range in credit scores. Solid lines separate the seven credit categories. Starting from left to right, the categories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

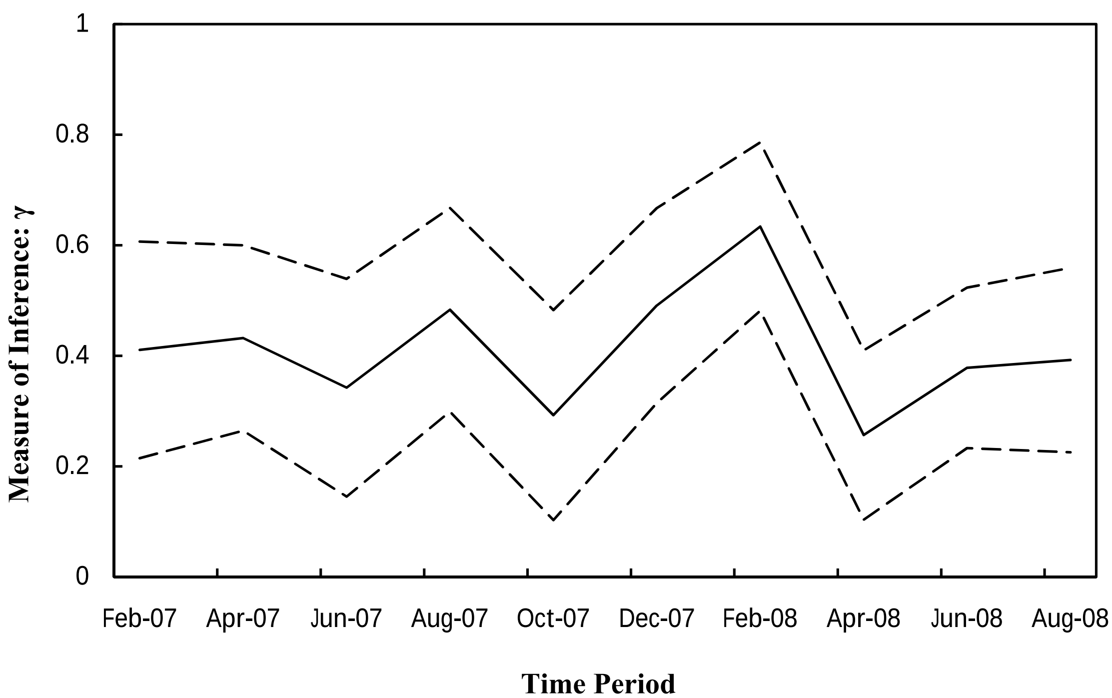

This figure shows our measure of inference, γ, for each two-month window from February 2007 to September 2008. Dotted lines represent 95% confidence intervals for each two-month γ estimate.

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
One-Year Default Dummy			0.143	
Credit Category HR			0.306	
Credit Category E			0.224	
Credit Category D			0.170	
Credit Category C			0.143	
Credit Category B			0.133	
Credit Category A			0.099	
Credit Category AA			0.054	
Standard Financial Variables				
Amount Requested ($)	8015	6577	6761	5788
Number of Current Delinquencies	2.89	4.54	0.77	2.28
Number of Delinquencies, Last 7 Years	9.68	15.78	4.30	10.52
Number of Public Record Requests, Last 10 Years	0.57	1.20	0.33	0.83
Total Number of Credit Lines	25.61	14.57	24.30	14.29
Number of Credit Score Inquiries, Last 6 Months	3.71	4.45	2.38	3.35
Amount Delinquent ($)	3191	12662	855	4504
Bank Card Utilization (total balances/total limits)	0.63	0.42	0.54	0.37
Number of Public Records, Last 12 Months	0.07	0.34	0.03	0.22
Number of Current Credit Lines	8.52	6.08	9.70	5.89
Number of Open Credit Lines	7.51	5.41	8.34	5.22
Revolving Credit Balance ($)	13446	33874	16773	38030
Debt-to-Income Ratio	0.54	1.37	0.33	0.90
Fraction Homeowners	0.37		0.48	
Credit History Age (years)	13.3	7.1	13.4	7.2
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 relevelent specifications in Table 2, Table 5, Table 6, and the Appendix tables. Definitions of variables that may not be self-explanatory are as follows: *One-Year Default Dummy* equals one if the loan is four or more months late by the sixteenth month of the loan (one year plus four months). *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

Dependent Variable: One-Year Default	(1)		(2)		(3)		(4)	
Estimate	**Coefficient**	**(S.E.)**	**Coefficient**	**(S.E.)**	**Coefficient**	**(S.E.)**	**Coefficient**	**(S.E.)**
Combined γ			0.931 ***	(0.169)	0.841 ***	(0.118)		
Regression Coefficients								
Credit score/40	-0.033 ***	(0.001)	-0.030 ***	(0.006)				
Credit category			-0.003	(0.006)				
α_2: Change between HR and E					0.015	(0.036)	absorbed	
α_3: Change between E and D					-0.002	(0.027)	absorbed	
α_4: Change between D and C					0.007	(0.017)	absorbed	
α_5: Change between C and B					0.036 **	(0.017)	absorbed	
α_6: Change between B and A					-0.007	(0.017)	absorbed	
α_7: Change between A and AA					-0.026 *	(0.015)	absorbed	
β_1: Change within HR					-0.111 **	(0.047)	-0.059	(0.053)
β_2: Change within E					-0.066 *	(0.038)	-0.067	(0.041)
β_3: Change within D					-0.034	(0.024)	-0.024	(0.025)
β_4: Change within C					-0.039 **	(0.020)	-0.040 **	(0.020)
β_5: Change within B					-0.049 *	(0.021)	-0.042 *	(0.022)
β_6: Change within A					-0.002	(0.022)	-0.001	(0.022)
β_7: Change within AA					-0.064 ***	(0.018)	-0.031	(0.027)
N	17212		17212		17212		17212	
Adjusted R^2	0.0311		0.0310		0.0333		0.1146	
Implied Coefficients and Tests								
$\gamma_1 = \beta_1/\delta_1$					1.155 ***	(0.387)		
$\gamma_2 = \beta_2/\delta_2$					1.105 ***	(0.413)		
$\gamma_3 = \beta_3/\delta_3$					1.080 *	(0.554)		
$\gamma_4 = \beta_4/\delta_4$					2.223 *	(1.140)		
$\gamma_5 = \beta_5/\delta_5$					1.410 ***	(0.385)		
$\gamma_6 = \beta_6/\delta_6$					0.198	(1.528)		
$\gamma_7 = \beta_7/\delta_7$					0.615 ***	(0.173)		
δ_1: Overall Change for HR					-0.096 ***	(0.033)		
δ_2: Overall Change for E					-0.060 **	(0.025)		
δ_3: Overall Change for D					-0.032 *	(0.018)		
δ_4: Overall Change for E					-0.018	(0.013)		
δ_5: Overall Change for B					-0.035 **	(0.014)		
δ_6: Overall Change for A					-0.012	(0.015)		
δ_7: Overall Change for AA					-0.104 ***	(0.023)		
p-value: α_i=0					0.252			
p-value: β_i=0					0.000		0.056	
p-value: γ_i=1			0.671		0.289			
p-value on Standard Financial Variables							0.0000	
p-value on Non-Standard Financial Variables							0.0000	

This table examines whether credit score predicts creditworthiness as represented by the one-year default dummy. A loan is in default if it is four or more months late by the sixteenth month of the loan (one year plus four months). Each specification includes listing month fixed effects to control for listing age. Column (1) shows results from an OLS regression of one-year default on credit score divided by 40, the size of the typical credit category. 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 the dependent variable. Column (4) shows default regressed on fracgap controls for each credit category and a flexible set of 215 controls for standard financial and non-standard variables as described in the Appendix, which are further interacted with the seven credit category dummies. 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	-0.0055 ***	(0.0008)		
Credit category	-0.0218 ***	(0.0009)		
α_2: Change between Categories HR and E			-0.038 ***	(0.005)
α_3: Change between Categories E and D			-0.059 ***	(0.005)
α_4: Change between Categories D and C			-0.049 ***	(0.004)
α_5: Change between Categories C and B			-0.051 ***	(0.005)
α_6: Change between Categories B and A			-0.031 ***	(0.005)
α_7: Change between Categories A and AA			-0.042 ***	(0.005)
β_1: Change within Category HR			-0.011 *	(0.006)
β_2: Change within Category E			-0.011 *	(0.007)
β_3: Change within Category D			-0.027 ***	(0.005)
β_4: Change within Category C			0.000	(0.005)
β_5: Change within Category B			-0.014 **	(0.006)
β_6: Change within Category A			-0.005	(0.007)
β_7: Change within Category AA			-0.052 ***	(0.008)
N	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: $\gamma_i=\gamma$			0.002	
p-value: $\gamma_i=0$			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 state 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 listings posted by borrowers with no previous Prosper listings	0.419	***	(0.044)	93117
(8)	Sample restricted to listings posted by borrowers with no previous Prosper 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 posted by borrowers with no previous Prosper listing or loan (funded listing), respectively. These tests confirm that our measurements of inference do not depend on information about the past repayment and listings history of borrowers who apply for more than one loan. 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 Category (HR -**	**Gamma High Category (B - AA)**	**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 Financial 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 Financial 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 financial variables and 2. non-standard 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%.

Table 6: Default and Interest Rates

Dependent Variable: One-Year Default	(1)	(2)	(3)	(4)
1/(1 + Interest Rate)	-1.333 ***	-1.417 ***	-1.409 ***	-1.043 ***
	(0.120)	(0.128)	(0.128)	(0.148)
P-values on Test of Joint Significance of:				
Credit Category Dummies	--	0.000	0.010	absorbed
Spline in Exact Credit Score	--	--	0.041	0.065
Standard Financial Variables	--	--	--	0.000
Non-Standard Variables	--	--	--	0.000
N	17212	17212	17212	17212
Adjusted R^2	0.0624	0.0663	0.0669	0.1182

This table examines whether interest rate predicts one-year default after controlling for other coded loan information. Each specification includes listing month fixed effects to control for loan age and a quadratic in borrower maximum rate. See Appendix I for a detailed derivation of the regression specifications. Column (1) shows results from an OLS estimation of one-year default on 1/(1+interest rate). Column (2) adds controls for credit category dummies. Column (3) further adds controls for a seven-part spline in exact credit score with kinks at the credit category boundaries. Column (4) adds controls for a flexible set of 215 controls for standard financial and non-standard variables as described in the Appendix, which are further interacted with the seven credit category dummies. 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 7: Causal Effect of Interest Rate on Default (IV)

Dependent Variable: One-Year Default	(1)		(2)	
Excluded Instruments: 7 Credit Category Dummies	**Coefficient**	**S.E.**	**Coefficient**	**S.E.**
interest rate	0.091	(0.413)		
1/(1+interest rate)			-0.177	(0.573)
Spline in exact credit score (kinks at credit category boundaries)	Y		Y	
N	17212		17212	
R^2	0.037		0.039	

This table examines whether interest rate has a causal effect on default using credit category borders as the exogenous excluded instruments. Each specification includes listing month fixed effects to control for listing age. Column (1) shows second stage results from a two stage least squares regression of one-year default on interest rate with controls for a spline in credit score (kinks in the spine are set at credit category boundaries). Interest rate is instrumented with credit category dummies (the excluded instruments). Column (2) shows a similar specification with 1/(1+interest rate) instead of interest rates so that coefficients can be compared to those in Table 6. First stage estimates are omitted for brevity and available upon request. 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%.

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

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

This table shows the decomposition of our estimate of gamma presented in Table 3, Column (2). The decomposition results are divided into standard financial 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 financial variables are missing in the data (observations with missing standard financial 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)			(2)			(3)			(4)
	Gamma			Gamma Low Categories (HR - C)			Gamma High Categories (B - AA)			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 3, Column (2). The decomposition results are divided into standard financial 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%.

Not For Publication

Appendix: Sample Listing

PROSPER® Let's bank on each other™

Out

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

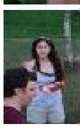

$8,081.00 @ 8.90%
Bid down from 13.99%

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

hs4p2
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

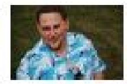

FORECAST | COMPARE | Help

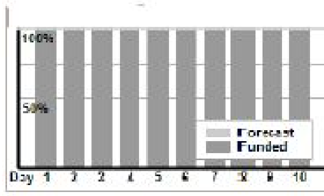

CREDIT PROFILE Help

A credit grade | Homeownership not verified | 10% debt to income ratio

Now delinquent:	0	**First credit line:**	Mar-2001	**Employment status:**	Full-time employee
Amount delinquent:	$0	**Current / open credit lines:**	3 / 3	**Length of status:**	1y 0m
Delinquencies in last 7y:	0	**Total credit lines:**	4	**Stated income:**	$25,000-$49,999
Public records last 12m / 10y:	0 / 0	**Revolving credit balance:**	$3,531	**Occupation:**	Computer Programmer
Inquiries last 6m:	1	**Bankcard utilization:**	29%		

Credit and homeownership information provided by Experian.

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	8.00%	$60.00	$60.00	✓	Oct 00 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
mrmoney	8.90%	$100.00	$100.00	✓	Oct-09-2007 7:47 AM
mrmoney	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
Gromila13	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
stevedavs444	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