Title
What's in a Label? A Series of Essays on Cognition in Markets and the Consequences of Formal Categorization

Author
Reschke, Brian Philip

Publication Date
2015

What's in a Label?

A Series of Essays on Cognition in Markets

and the Consequences of Formal Categorization

By

Brian Philip Reschke

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Toby E. Stuart, Chair
Professor Ming D. Leung
Professor Marti A. Hearst

Summer 2015

Abstract

What's in a Label?

A Series of Essays on Cognition in Markets

and the Consequences of Formal Categorization

by

Brian Philip Reschke

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor Toby E. Stuart, Chair

Markets are social systems. While price is preeminent, it is often insufficient for buyers to determine the value of a product or service. Because of difficult-to-discern heterogeneity in offerings, prices are offered in the context of other kinds of claims, such as the offering's type, quality, and the endorsements of previous buyers. This is especially the case for markets wherein the 'products' defy easy valuation, such as meals at restaurants, art, movies, lectures, and other goods or services where direct human skill is required.

A sizeable subfield of sociology has studied categorization in markets. A central premise of this literature is referred to as Zuckerman's (1999) 'categorical imperative': audiences seek to categorize candidates before considering their distinguishing features. Once categorized, the candidates are assessed relative to socially agreed upon criteria as to what is an acceptable candidate within a given category. A host of empirical work in this vein has demonstrated the hazard of failing to be categorized reliably, as well as defying social expectations of category membership. Almost universally, this work has equated category membership with having labels and has generally taken the existence of labels as a given.

In my dissertation, I consider the consequences of labeling, with a focus particularly on the interplay among labels and features. My dissertation investigates (a) the causal effect of labeling on the returns to coherent feature combinations, (b) the impact of labeling on similarity, and (c) the potentially positive results of novel label combination. Ultimately, I find evidence that labels and features provide different information depending on market conditions and the composition of audiences.

To Cami

Contents

List of Tables

List of Figures

Acknowledgements

I wish to thank the many individuals who have helped in the development of this dissertation.

The faculty in the Management of Organizations group and the MORS PhD students have provided a supportive community of scholars and friends. I especially thank Eliot Sherman and Weiyi Ng for our countless discussions of theory, code, and research design. Their engaged participation in several presentations of this work has heavily influenced my development as a researcher.

Tony Cricelli patiently accommodated my many, sometimes agitated requests for installation of open source software packages on the Haas research computing cluster. Without his cooperation, this work would have been prohibitively difficult to conduct.

Also, my many thanks to the faculty outside of Haas for accommodating a business PhD student in their seminars and office hours. Bin Yu, Chris Paciorek, John Canny, and Michael Anderson provided essential methodological instruction both within and outside the classroom. Lee Fleming and Jas Sekhon provided invaluable exposure to computational social science research through their seminar. As members of my orals exam committee, Lee Fleming and Bin Yu provided formative feedback on the early outlines of this work. As a member of my dissertation committee, Marti Hearst was crucial in helping me refine my research and compile it into a working whole. I have greatly benefitted from UC Berkeley's commitment to porous departmental boundaries.

I thank Ming Leung, whose willingness to bring me aboard as a fledgling graduate student gave me my start in the study of market categorization and exposed me to my research site. Chapter 4 of this dissertation contains one of our collaborations. I thank my advisor and committee chair, Toby Stuart, whose encouragement and thorough assessment of my research questions and design helped me persist in this project. Our innumerable conversations have greatly informed my thinking and the work that follows.

It would be unfortunate if the concentration of credit for this work centered on me and my professional circle alone. My family has played an essential supportive role. I thank my grandparents, parents, in-laws, and siblings, who continue to teach and inspire me. My daughters, Madelyn and Claire, have been ready reminders of what is most important in life. Their repeated, 'You can do it, Dad!' has helped me continue to the end.

Lastly, I express my deepest gratitude and appreciation for Cami. Her faith, resilience, and sacrifice in the care of our family have made this dissertation and the associated years of study possible.

1 Introduction

How do markets assign value? Economic sociology has nuanced the efficient markets hypothesis, in which disparities between price and value are removed by arbitrage, by identifying limits to the diffusion of market information and highlighting information's characteristically social nature. While price is preeminent, price is often insufficient for buyers to determine the value of a product or service. Often due to difficult-to-discern heterogeneity in offerings, prices are assessed vis-á-vis other information, such as claims as to product quality, affiliation with high-status actors, or geographic or social proximity. This is especially true of cultural markets, wherein the 'products' defy easy valuation—accordingly, much of the sociological study of markets has focused on restaurants, art, movies, and other goods or services directly involving human expertise or skill (DiMaggio 1987; Rao, Monin, and Durand 2003; Waguespack and Sorenson 2011). However, even markets for relatively homogeneous products have demonstrated reliance on social considerations, such as producers' compliance with localized norms of engagement (Ody-Brasier and Vermeulen 2014). Markets are decidedly social systems consisting of human actors subject to biases and cognitive constraints.

A sizeable subfield of this vein of sociology has studied market categories and categorization. An early premise of this literature is referred to as Zuckerman's (1999) 'categorical imperative': buyers (or, more generally, 'audiences') seek to categorize products or producers ('candidates') before considering their distinguishing attributes. This thesis suggests two, ideal-typical stages of valuation: categorization and differentiation. In the categorization stage, audiences seek to answer the question, 'what kind of a (product, person, organization) is this?' (Albert and Whetten 1985), mapping candidates to one or more concepts from a finite, generally agreed-upon set. Conditional on this categorization, audiences may then differentiate candidates by price, quality claims, social proximity, and other information, interpreting these signals in the context of that candidate's category. For example, a four-star product rating is meaningful only in view of what other similar products have been rated. Rather than assess all candidates and their attributes simultaneously, audiences categorize candidates to narrow their consideration sets to a (more) manageable size. Individual evaluators can then compare a focal candidate to others in their previous, direct experience.

Much of the research in market categories has focused on the constraints this categorical imperative presents to candidates. Candidates that fail to be identified with a recognized category (Zuckerman 1999) or that are identified with multiple categories (Hsu 2006; Hsu, Koçak, and Hannan 2009; Ruef and Patterson 2009) tend to be ignored or devalued relative to singly-categorized candidates. Progress in this vein has identified conditions in which the 'multi-category discount' is amplified or alleviated, as well as conditions in which category spanning commands a premium (Pontikes 2012; Smith 2011).

My dissertation focuses on an aspect of market categorization that scholars have tended the take for granted: the role of labels in conveying and conditioning category membership information. For the most part, research in market categorization has studied markets in which category membership is inferred from labels such as industry classification codes (Zuckerman 1999), IMDB movie genres (Hsu 2006), and Amazon product classifications (Hsu, Koçak, and Hannan 2009). Labels present a shortcut for audiences to infer what kind of candidate is under consideration. Yet labels are not the only means of conveying category membership. Membership may also be inferred from what a candidate says about themselves, the actions a candidate takes, or a candidate's affiliations. Words, actions, and network ties comprise a candidate's *features*.

Labels and features convey category membership information differently. Labels are especially potent signals of between-category membership because the mapping from the label to a concept is basically immediate. In contrast, features may have varied or ambiguous connotations and must often be taken in combination with one another to form a basis for meaning.

At same time, labels and features work together in important ways. Whereas labels by themselves convey discrete information of what categories candidates identify with, features can provide gradation in category membership, or qualify category membership implied by a label. For example, audiences may examine features in order to scrutinize candidates' claims made by labels (e.g., a customer may examine a grower's agricultural practices to assess whether the 'organic' label is warranted). Also, features may distinguish a candidate from others affixed with the label. For instance, Airbnb may signal their presence in the hotel and lodgings category through the appropriate SIC code, but through their statements in press releases, advertising, and social media communications, they may indicate that they are anything but a 'typical' lodging services firm.

In my dissertation, I study labels, features, and their interplay. My dissertation consists of three empirical studies. I examine how audiences rely on features to assign candidates to categories, even if labels conveying crisp category membership information are unavailable (Chapter 2). I also study how features that are typical of a category are alternately helpful or hazardous for candidates depending on whether labels are present (Chapter 3). Lastly, extending the focus on the desirability or danger of typicality, coauthor Ming Leung and I consider what audiences find atypicality most appealing (Chapter 4). Together, my research suggests that a reliance on labels alone to study categorization in markets omits important nuance in the sociological account of valuation.

While features could be a variety of candidate attributes (e.g., ingredients on restaurant menus—see Kovacs and Johnson 2014), I study the words that candidates use to describe themselves. In studying words (and word pairs), I take on a level of granularity seldom seen in market category research, but in doing so, I model a research strategy that could be readily redeployed at higher levels of analysis.

In each chapter that follows, I motive my focus on labels, features, and variation in the audiences that adjudicate them. These chapters share the empirical setting of Prosper.com, a peer-to-peer lending website that enables strangers to lend money to one another over the internet. Prospective borrowers created public profiles and loan requests in which they detailed their desired loan amounts, creditworthiness, and (critically) provided written descriptions of themselves and why they need a loan. In creating the marketplace, Prosper made loan request data and bidding data publically available for users and researchers to study.

As will be described further below, Prosper represents an ideal setting for studying labels, features, audiences, and their interplay. The centerpiece of my empirical strategy for Chapters 2 and 3 is a natural experiment in which a labeling system is introduced. From the beginning of Prosper's history, borrowers were asked to provide a written description of why they need a loan. Then, about two years after launching the online lending platform, Prosper gave borrowers the additional requirement of choosing a label that summarized the purpose of their loan. This sudden change on the platform and machine learning tools allow me to compare loan requests that are labeled and not labeled, and thus to assess how labels change the evaluation of two kinds of feature information: *coherence*, or the extent to which a lender's words identify them with one or many categories; and *typicality*, or the degree to which a lender's words are typical of lenders providing similar descriptions. In Chapter 4, coauthor Ming Leung and I exploit rich data on lender's bidding activity to assess how engaged a lender has been with the platform at the time of each bid. We examine how the value lenders place on atypical label combinations varies with lender engagement.

In Chapter 5, I provide general conclusions and outline paths for future research.

2 The Limit or License of Labels: How Formal Categorization Affects the Evaluation of Category Spanning

ABSTRACT

Research on market categorization has emphasized the constraints labels impose on producers and products. Drawing from recent developments in the economic sociology of markets and valuation, I investigate how labels impact the assessment of public accounts of activities and intentions. Specifically, I assess whether labels restrict or expand the scope of claims producers can credibly command in self-descriptions. Through machine learning, I assess the category memberships suggested by prospective borrowers' written descriptions of why they need money, and then I construct a continuous measure of descriptions' category spanning. Then, exploiting the introduction of a new labeling requirement, I assess the impact of formal classification on the valuation of self-description. I find that in the presence of labels, the respective benefits and hazards of focused and diffuse claims are attenuated.

INTRODUCTION

Classification systems reflect social boundaries that are not trivially crossed. A growing body of research in the sociology of markets has shown that categorization is inferred from labels affixed to candidates, such as patent classes or industry codes. Generally, the thrust of this literature has been to show the constraints that formalized classification has on candidates. Having a label, not having a label, or having multiple labels can be consequential. For instance, failing to have one's category membership sanctioned by appropriate critics (Zuckerman 1999) or exhibiting membership in multiple labeled categories (Hsu 2006; Hsu, Koçak, and Hannan 2009; Ruef and Patterson 2009) can result in audiences ignoring or devaluing candidates. Accordingly, producers that fail to be identified with recognized categories are ignored , and (inexperienced) producers identified with multiple categories are devalued (Hsu 2006; Hsu, Hannan, and Pólos 2011; Ruef and Patterson 2009; Zuckerman 2000; Zuckerman et al. 2003) . Market structure is reproduced both as audiences sanction conforming producers (White 1981; Zuckerman 1999).

Yet the pervasive assumption underlying this literature is that formal categorization systems need to be institutionalized to be 'imperative'—that is, there must be broad consensus concerning the attributes of actors that belong to various categories for complex identities to be devalued or ignored (Ruef and Patterson 2009). Implicit in this assumption is the notion that domains need labels to facilitate this consensus among audiences, and thus bring identity into consideration of value. Recent work in categorization has relaxed the reliance on labels, noting that identity can be inferred from features and actions. For example, Phillips, Turco, and Zuckerman 2013 study law firms' engagement in various areas of law, and Jensen and Kim (2014) describe opera companies' staging of conventional and unconventional operas. In addition to showing that not all category spanning is necessarily hazardous, these scholars also show that audiences are attentive to fine-grained features of actors in assessing identity, not just discrete signals such as labels..

My study is in this spirit of broadening the range of activities considered as mediums of identity claims, as well as revisiting the conditions in which the penalty for spanning multiple categories is operative. I focus on the features conveyed through *self-description*—the public accounts actors give of their actions and intentions. Descriptions are distinct from labels because they employ language much more expansively. Whereas a label typically maps onto one social category, description may lay claim to multiple, or even purport the creation of a new social category. In special cases, descriptions may present sequence of events, stories, or narratives to illustrate an actor's robustness . Across all of these uses is the effort to convey what or who a candidate is and how they are distinct from others.

Descriptions have been studied particularly in nascent markets, when formal categorization has yet to materialize. Here, actors almost of necessity resort to written or verbal descriptions of themselves to signal what kind of actor they are. Public discourse may shape formal market categories (Rosa et al. 1999).

The emphasis on when and how labels emerge is important, and yet could suggest that once formal classification systems have arrived, descriptions are either a redundant means of inferring identity, or worse, irrelevant. Yet investigation of existing theory suggests the answer is not obvious. The presence of labels may constrain claims in other channels as labels become a template for interpreting an actor's activities. But in summarizing an actors' claims, labels may actually give actors license to make varied identity claims in other channels without penalty. What is the effect of labels on the interpretation of identity claims in public discourse? How do descriptions and labels interact in the process of assigning value?

Investigations into the interplay among labels and discourse meet empirical challenges. The principal problem is assembling the counterfactual: how would actors present themselves without labels, and how would these claims be evaluated?

My research setting and empirical strategy address this challenge. I study Prosper.com, one of the earliest and largest crowdfunding markets in the United States. Prosper facilitates lending among strangers over the internet. Prospective borrowers create an online listing, including a written description of why they need money and why they are a good loan candidate. On December 5, 2007, Prosper imposed a new requirement on prospective borrowers. In addition to describing the purpose of their loan, borrowers now select one of seven labels summarizing the intended purpose of their loan: debt consolidation, home Improvement, business, personal, education, auto, and other. This change allows for comparison of labeled and unlabeled loan candidates and their proximity in discourse space.

I exploit the fact that all listings have text descriptions. Representing listings as vectors of word frequencies, I use supervised machine learning to model listing purpose as a function of description words. Next, through coarsened exact matching, I construct a matched sample of labeled and unlabeled listings that have similar distributions of observable characteristics. Using this sample, I characterize the extent to which borrower's written descriptions of why they need money reflect a coherent or diffuse claim, or the extent to which they suggests membership in one or more purpose categories. I also consider how imitating the claims of immediate predecessors is valued. Modeling the amount of funding a listing receives as a function of loan and borrower attributes, I investigate changes in the returns to coherence with the onset of labels.

Previous work has considered the role of labels in the multi-category discount (Leung and Sharkey 2014) and the contribution of qualitative information to funding (Herzenstein, Sonenshein, and Dholakia 2011); I extend this research and consider how the presence of labels alters the valuation of self-description. There are many aspects of self-description that could be interesting for study of valuation. Some recent work has examined linguistic styles of loan requests (e.g., readability, sentiment; Gao and Lin (2013)). The present study focuses instead on the content of user-submitted text and the category memberships it implies.

I contribute to the economic sociology of markets and organizations, as well as a budding literature studying crowdfunding. Extant work has provided rich descriptions of particular platforms (Mollick 2014) and the dynamics of the funding process. Others have investigated the viability of crowdfunding as an alternate or preliminary source of venture capital. Some have studied crowdfunding platforms to elucidate the collective action problem (Burtch, Ghose, and Wattal 2013) and have tested claims that crowdfunding (and online platforms generally) have removed biases in the allocation of capital (Agrawal, Catalini, and Goldfarb 2011). I approach crowdfunding platforms as 'markets in microcosm' that may yield insights for more established markets.

Below, I theorize the impact of labels on the evaluation of self-description, describe my research site, outline my methods, and present results.

THEORETICAL DEVELOPMENT

Research in market categories has been largely concerned with the relative performance of singly- and multiply-categorized candidates. Most of the focus has been on how candidates identified with multiple categories tend to be devalued. Accounts for this phenomenon can characterized as belonging to one of two camps: a normative account, in which multiple category membership is actively penalized because of social constructed norms of conduct (Carroll and Swaminathan 2000; Phillips, Turco, and Zuckerman 2013; Ruef and Patterson 2009); and a cognitive account, in which the prospective of evaluating multiple categories simultaneously confuses audiences or is cognitively difficult (Kovacs and Johnson 2014; Leung and Sharkey 2014). The general program of categories research—this dissertation included—has not sought so much to pit one set of mechanism against the other, to assert that only one account is valid, but to elucidate the conditions in which normative and cognitive mechanisms are at work, and when they may be working jointly.

To this point, most work in market categories have applied normative or cognitive accounts to generate expectations of how label combinations are evaluated. At first glance, my focus on features (and self-description particularly) could be seen as simply a generalization of the study of multi-category discount: features suggesting membership in multiple categories leads to lower valuations 'because of X'. If this was all, then the following would be more of a methodological contribution than a theoretical extension: that is, I do with features (words) what others have done using labels for methodological simplicity. But I argue the distinction between labels and features has more value than finding yet another way that multiply-categorized candidates are devalued. Ultimately, I am interested in what labels do to the interpretation of features: how the possession of a label limits or gives license to the scope of words, actions, or associations candidates can assume. This returns market categories research to some of its roots in labeling theory (Ashforth and Humphrey 1997).

Also, the question of labels' impact on features assembles comparison of normative and cognitive accounts of market categorization. These accounts supply conflicting expectations of how the absence or presence of labels impacts the evaluation of feature combinations. Under the normative account, labels provide a common reference for lay theories of a market's categorical structure, and this common reference is necessary for a multiple-category discount. Under the cognitive account, which was more center stage of the original formulation of the categorical imperative (Zuckerman 1999), feature combinations are especially consequential when labels are absent, and less so when labels are present.

Below, I describe these conflicting accounts and their respective expectations for how feature information is evaluated. In doing so, I necessarily focus on a particular type of feature information, since a variety of information could be considered—for instance, others have studied how grammatical errors or positivity are evaluated (Gao and Lin 2013). Since I am studying market categories, I focus on the category membership information features convey. In this chapter, I focus on *feature coherence*, or the extent to which feature combinations suggest membership in one or multiple categories. In the next chapter, I consider *typicality*, or the extent to which a candidate's features are representative of others identified with their associated categories.

Normative Account: Labels Necessary to Coordinate Audience Expectations Concerning Features

The normative account observes that categories are value-laden. Some categories may be considered 'better' than others. Some categories may be associated with conflicting logics such that particular combinations constitutes betrayal (Phillips, Turco, and Zuckerman 2013). Yet the normative account suggests the development of these values is not immediately evident and requires the process of institutionalization. Per Ruef and Patterson (2009), multiply-categorized candidates are penalized only when there is consensus regarding the content of categories. This could include expectations regarding what words are associated with being in a given category—what a 'typical' member of Type A says. When such consensus is strong, multiple-category claims invite assessment of one's ability to conformity to multiple rules simultaneously, raising questions of one's ability to allocate resources appropriately (Hsu 2006).

Without consensus in expectations, a candidate's claims as being in one category or another could not be systematically policed and penalized. Under this account, descriptions that make multiple identity claims in early periods of markets are less subject to socially constructed rules, because these rules are in flux. The presence of labels should make it easier for audiences to coordinate expectations, and they should make claims to multiple purposes more apparent.

In my empirical setting, an online market for loans, there is a natural, normative explanation for why borrowers identifying with multiple purposes would perform worse. Audiences may infer that such borrowers are either spread too thin financially to make payments (e.g., mired in education expenses as well as costly personal health challenges), or are irresponsible in their consumption and therefore undeserving (e.g., consolidating credit card debt and leasing a luxury vehicle). Also, even if the former case of category combination could elicit sympathy from some lenders, at some limit, the layering of personal woes invites suspicions of credibility. Readings of public Prosper forums suggest that lenders did not take all description text at face value, and outed imposters when possible.

Taken together, this suggests that feature coherence is positively (intentionally) valued under the normative account, and that is coherence is best appraised when labels are present.

Ceteris paribus, feature combinations that reflect multiple category membership receive lower valuations than candidates reflecting membership in one category, but only in the presence of labels.

Cognitive Account: Feature Coherence in Nascent Categorization

The cognitive account has its roots in the initial formulation of the categorical imperative (Zuckerman 1999): candidates are distinguished from one another, such as ranked-ordered in terms of quality (differentiation), but only after candidates can be assumed to be members of a similar type (categorization). Under this account, feature coherence informs the categorization stage first, and then possibly the differentiation stage.

When labels are absent, categorization proceeds 'in the rough', with individual audience members proceeding with their own lay theories of the categorical structure of the market. Evaluators compare candidates to others they have encountered in the past (Berger and Luckmann 1967). In these conditions, candidates whose features cohere around one type will more easily fit into audience members' individually developed schemata of what kinds of candidates are considerable. Thus, candidates with high feature coherence will be more readily evaluated.

With labels, categorization can proceed with much less cognitive effort. Feature coherence is less crucial to informing audiences' initial perceptions of candidate identity, and can enter more into the differentiation process.[1] If anything, the relationship between feature coherence and evaluation should abate in the presence of labels.

Ceteris paribus, feature combinations that reflect multiple category membership receive lower valuations than candidates reflecting membership in one category, regardless of the presence of labels.

DATA: PROSPER.COM

San Francisco-based Prosper Marketplace, Inc., began operations in 2005 and launched its online platform Prosper.com in February 2006. At the time of founding, some online lending sites were already in operations elsewhere (Britain-based Zopa.com) and more have emerged since (e.g., VirginMoney, LendingClub.com), but Prosper claims to be the first and largest peer-to-peer market of its kind in the U.S. Since 2006, the website has accumulated over 2 million registered members and funded more than $1 billion in funded loans (Prosper.com 2012).

Although the market has changed its policies several times, Prosper's loan creation process and auction approach have generally remained the same.[2] Prospective borrowers register with the website, authorize Prosper's access to their Experian credit history, and create a personal profile. To create a loan request (listing), borrowers construct a one-sentence title, describe the purpose of their loan, and indicate the desired amount and the maximum borrowing rate they are willing to endure. All loans are for a three-year term and are unsecured. In the period studied, requested amounts range from $1000 to $25,000, and borrowing rates range from 0% to 36%.

After a listing is approved by Prosper, it is posted to the website. Figure 1 contains a screenshot of the general browsing page from the early period of the market. Registered lenders may browse listings in chronological order or search listings by attributes. Major listing features (e.g., credit grade, amount requested) are visible on the main browsing page; once a lender clicks on a given listing, they can view more details, including descriptions. Figure 2 shows a screenshot of a listing representative of the analysis period.

[1]In this way, the cognitive account of categorization can accommodate the scrutiny of the normative account, but identifies feature coherence as more an interest of categorization than differentiation.

[2]Presently, interest rates are fixed by Prosper. Over the course of the period studied, interest rates are variable and are achieved through auction.

Figure 1: Sample Prosper Listings

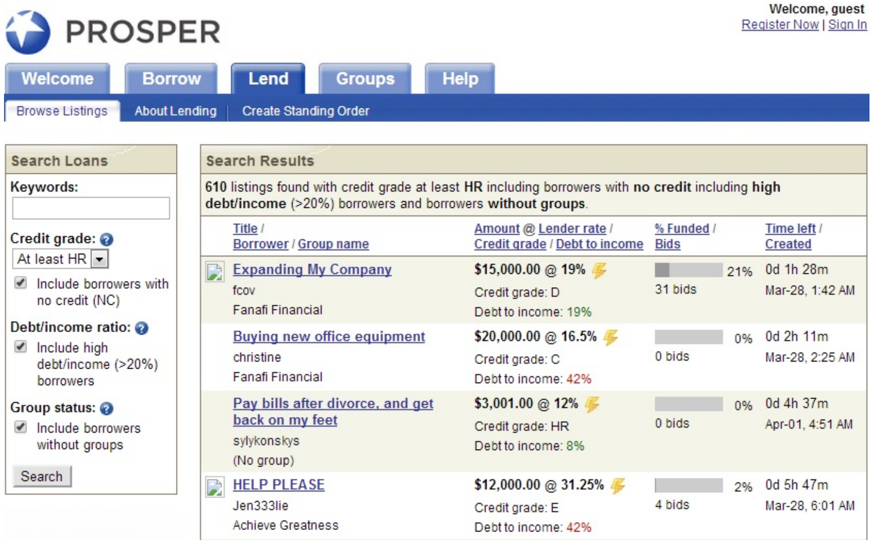

Note: A screenshot of Prosper - April 11, 2006.

Lenders 'bid' on selected listings by committing to fund a portion of the proposed loan at a specific rate. Bids are usually between \$50 and \$100, and listings usually require contributions from many lenders. Rates must match or be less than the maximum borrower rate listed by the borrower. Depending on the listing's funding scheme, more bids results in lower rates for borrowers. Under the 'open' funding option, the listing remains available for a full week, making it possible for lenders to compete in bids and thus lower the rate. In this case, lenders who bid at a higher interest rate than the realized interest rate are informed that they are 'outbid' through an email, and they have the opportunity to bid again until the 'open' listing closes. Lenders can make multiple bids on listings at a time. Listings with the 'closed' funding option become loans once the cumulative amount bid equals the amount requested. On average, listings in the analysis period receive bids totaling about 13% of the amount requested, with 7.4% of listings becoming loans.

Despite the term 'peer-to-peer', Prosper users do not lend money directly to each other. After a listing is fully funded, Prosper collects bid amounts from lenders, contracts with a bank to loan the amount to borrowers, and issues no-recourse promissory notes to lenders. Prosper collects monthly payments from borrowers and pays lenders. In the case of late payments, Prosper contracts collections services and notifies lenders and credit agencies. As a multi-sided platform, Prosper assesses fees from

Figure 2: Sample Listing: "Expanding My Company"

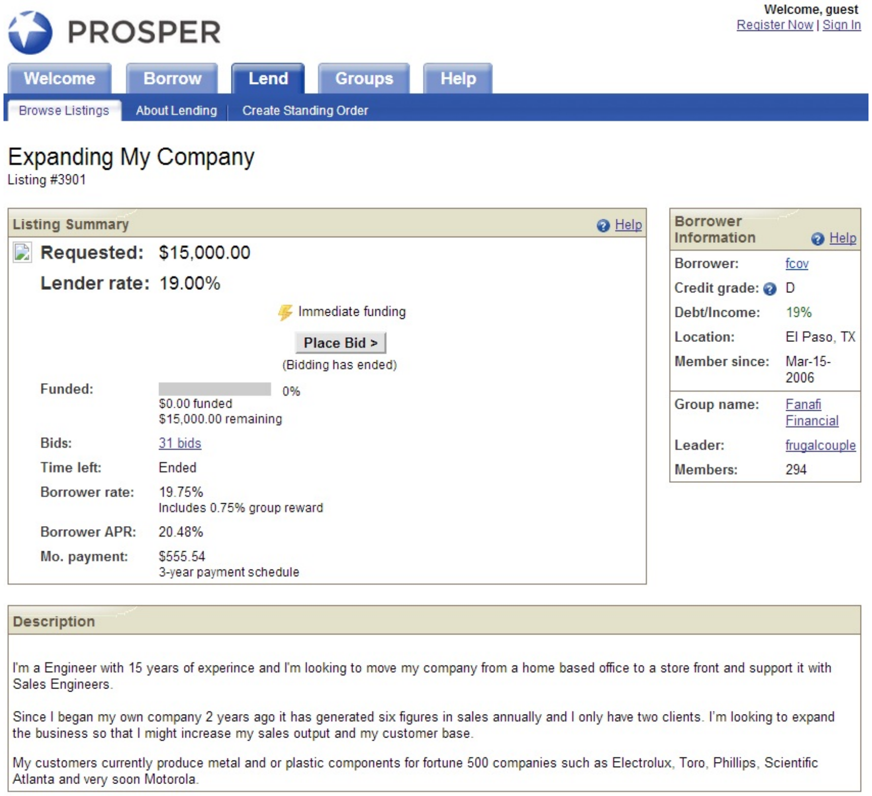

Note: A screenshot of a Prosper Listing - April 11, 2006.

both borrowers and lenders. In the period studied, borrowers pay a loan origination fee of zero to three percent, depending on creditworthiness, and lenders pay zero to one percent of the outstanding principal balance of the loan annually (SEC 2008). Prosper also collects fees for featured listings, which receive more prominence on the website.

Introduction of Purpose Labels

On December 5, 2007, Prosper adopted a simple taxonomy of loan purpose. Borrowers continue to provide written descriptions, and they select a label summarizing the intended use of the loan. During the period studied, the labels included the following: debt consolidation, home improvement, business, personal, education, auto, and other. Lenders could now search listings on the basis of categories, and in time could create standing orders to invest in listings of a given category (and other features) automatically.

While Prosper users actively discussed changes on the website in parallel discussion forums, the matter of purpose labels was relatively muted. On September 13, 2007—a few months before the labels were introduced—Prosper announced in its discussion forum that it would be introducing the feature, identifying debt consolidation, business financing, and student loan as possible options. The announcement drew little response from the Prosper lender community—most commenting on the announcement post were reacting to changes to Prosper groups—and those who did comment were brief and positive. One user observed, "I have certain types I like to avoid and certain types I prefer and it would be great to pre-filter those (of course I know borrowers could always lie, but at least it shrinks the results I have to go through)." . This comment suggests at least some lenders saw purpose labels as a way to incorporate preferences for certain kinds of borrowers, and that some acknowledged that purpose information was not necessarily factual.

METHODS

Below, I describe the general empirical strategy, detail the process of purpose label prediction, and present results.

Empirical Strategy

My empirical strategy has two main components. First, I need a way to assess whether candidate claims are focused (pertain to one category) or diffuse (pertain to multiple categories). Second, I need to construct proper comparisons of labeled and unlabeled candidates.

My setting features candidates that make public, written accounts of their intended use of resources, and I have a period of the market in which these accounts are associated with purpose labels. The loan candidates on Prosper could only choose one label, thus eliminating the simple strategy of comparing listings that have one label with those that have multiple labels. Yet I posit that although texts are associated with one category only, as suggested by a label, borrowers signal intentions that pertain to (a) multiple purposes, and (b) to varying extents. For example, a borrower that is labeled as consolidating their debt may describe their need to reduce their use of credit cards, but may also indicate their intent to remodel their home. Present research that equates multi-category membership with multiple labels ignores the potential heterogeneity in candidate claims through other symbolic action.

Furthermore, the fact that Prosper constrained borrowers to select one label only provides a methodological advantage. Maximum entropy classification is a process of predicting the classes (labels) of documents based on the observed words or other features. The result is a series of category weights—essentially, predicted probabilities that a document is associated with each of the possible labels. Although methods for multiply-labeled documents exist, having one label per document considerably reduces the computational task.

My ultimate interest is how labels impact the evaluation of focused versus diffuse identity claims. This requires me to characterize claims of labeled and unlabeled documents. Through supervised machine learning, I can use labeled documents to train a classification algorithm that I can then apply to labeled and unlabeled documents alike. However, the algorithm will have the highest predictive performance for documents used for training. This could potentially introduce bias between the 'untreated' and 'treated' documents that could obfuscate the impact of labels on claim evaluation. Accordingly, I take the approach of dividing my analysis period into three parts: an 'untreated' period, a 'treated' period, and an 'algorithm training period'.

Although several years of market data are available, to help ensure comparable market conditions, I select a tight observation window around the time of label adoption, and chose analysis periods so as to minimize global confounds. See Table A.1 in the Appendix for a list of changes to the Prosper platform before, during, and after the analysis period. The 'untreated' period of study is June 6, 2007 through December 6, 2007 (when purpose labels were adopted); the 'treated' period is December 7 through April 14, 2008; and the 'training' period is April 15, 2008 through June 27, 2008. There were 82,161 listings created and submitted for funding during the untreated period, and 49,275 and 35,546 listings submitted in the treated and training periods, respectively.

At the beginning of the untreated period, Prosper provided a template to prospective borrowers to structure their loan description. Borrowers were prompted to describe the purpose of their loan, their financial situation, and outline a monthly budget. This change reduced heterogeneity in loan descriptions and made borrowers more comparable. The end of the 'treated' period does co-occur with a change in the maximum borrower rate to 36% in almost every state, although inspection suggests this change doesn't drastically alter the composition of loan purposes, which is the sole input of interest in the training period. The training period ends when the Personal category was eliminated, which did impact the arrival rates of loan purposes.

The modal Prosper borrower rarely was funded on the first attempt during my analysis period. Borrowers made between 2-3 attempts on average. In robustness checks, I constrained listings in the untreated and treated periods to borrowers' first attempts only. Models estimated on this restricted sample yield results highly similar to those reported below.

Table 1 reports descriptive statistics for analysis listings, subdividing by 'pre-label' and 'labeled' periods. There are differences across these two market periods. For instance, Prosper credit grade (a site-specific binning of credit scores) shifts upward; the maximum borrowing rate, or the starting rate for the auction, increases. The change in amount requested is especially striking: borrowers requested $9,520.48 on average in the unlabeled period, compared to $8,605.13 in the labeled period.[3] Such differences between unlabeled and labeled epochs of the market motivate the coarsened exact matching approach described below.

Purpose Classification

What candidate attributes should be considered? The labels Prosper adopted distinguish among possible purposes for needing funding, as opposed to other borrower attributes, such as credit grade and social affiliations. While such affiliations could certainly be informative for lending decisions, I constrain my analysis to purpose claims for several reasons.

First, the central issue in the literature on category spanning is whether actors can present themselves as belonging to multiple types simultaneously. While the broader literature on boundary crossing has considered issues such as identification with multiple social categories, the emphasis in the sociology of markets has been more on nominal than on relational or interpersonal distinctions. Purpose is more of a nominal distinction. Second, traditional financial markets have characterized debt by the intended use of funds (e.g., education), especially when loans are secured by the target (e.g., auto loan). Third, while previous research has used human coding to count the number of social identity claims present in descriptions (Herzenstein, Sonenshein, and Dholakia 2011), it is difficult to set an a priori bound on the number of social claims to consider. Prosper's purpose labels provides a meaningful guide. I identify claims in descriptions corresponding to Prosper's purpose labels and construct a measure of how much a description reflects one or more of this limited set of purposes. If this is

[3]Further exploration failed to find the introduction of a new cap on amount requested for a given segment of the market: borrowers requested as much as $25,000 in every week of the data. It is possible that Prosper prompted users to ask for less, as borrowers that asked for a more manageable amount historically experienced a higher funding rate.

Table 1: Descriptive Statistics: All Listings

	Pre-Label Listings (N = 81,477)			Labeled Listings (N = 49,959)			All	
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.	Min.	Max.
Percent Funded	0.11	0.00	0.28	0.15	0.01	0.32	0.00	1.00
Nb. Words	197.91	161.00	130.74	179.27	142.00	123.51	1.00	770.00
Nb. Words (excluding budget)	159.12	127.00	118.03	147.33	113.00	114.06	1.00	770.00
Amount Requested	8404.27	6001.00	6328.02	7325.05	5000.00	6383.62	1000.00	25000.00
Max. Borrower Rate	0.18	0.18	0.07	0.19	0.18	0.09	0.00	0.36
Debt to Income Ratio[1]	0.53	0.27	1.26	0.43	0.26	0.98	0.00	10.01
Credit Grade: AA (760 and up)	0.02	0.00	0.15	0.04	0.00	0.20	0.00	1.00
Credit Grade: A (720-759)	0.03	0.00	0.17	0.05	0.00	0.22	0.00	1.00
Credit Grade: B (680-719)	0.05	0.00	0.22	0.08	0.00	0.27	0.00	1.00
Credit Grade: C (640-679)	0.10	0.00	0.30	0.14	0.00	0.34	0.00	1.00
Credit Grade: D (600-639)	0.16	0.00	0.37	0.18	0.00	0.38	0.00	1.00
Credit Grade: E (560-599)	0.18	0.00	0.39	0.18	0.00	0.38	0.00	1.00
Credit Grade: HR (520 to 559)	0.45	0.00	0.50	0.33	0.00	0.47	0.00	1.00
Duration	7.60	7.00	2.09	7.50	7.00	2.27	3.00	10.00
Homeowner	0.34	0.00	0.47	0.39	0.00	0.49	0.00	1.00
In Prosper Group	0.21	0.00	0.41	0.13	0.00	0.33	0.00	1.00
Automatic Funding	0.43	0.00	0.50	0.10	0.00	0.29	0.00	1.00
Nb. Attempt by Member	2.81	2.00	2.88	3.33	2.00	4.60	1.00	68.00
Number Live by Start	2617.84	2539.00	334.38	2372.85	2295.00	442.07	1684.00	3973.00
Percent Live in Same Category	0.36	0.33	0.15	0.32	0.27	0.16	0.01	0.52
Predicted Label: Debt	0.43	0.00	0.49	0.46	0.00	0.50	0.00	1.00
Predicted Label: Home	0.02	0.00	0.13	0.02	0.00	0.15	0.00	1.00
Predicted Label: Business	0.13	0.00	0.34	0.15	0.00	0.35	0.00	1.00
Predicted Label: Personal/Other	0.40	0.00	0.49	0.33	0.00	0.47	0.00	1.00
Predicted Label: Education	0.02	0.00	0.15	0.02	0.00	0.15	0.00	1.00
Predicted Label: Auto	0.01	0.00	0.11	0.02	0.00	0.12	0.00	1.00
Coherence	0.75	0.78	0.21	0.78	0.85	0.21	0.17	1.00
Second-Highest Probability	0.14	0.10	0.14	0.13	0.07	0.14	0.00	0.50
Typicality[2]	50.14	50.00	28.66	51.09	52.00	29.19	1.00	100.00

N Listings = 131,436. Note: Descriptive statistcs for listings that appeared on Prosper's website from June 6, 2007 through April 14, 2008.
[1] Missing values for 9236 observations.
[2] Missing values for 6 observations. Typicality is the subject of Chapter 3.

the 'wrong' partition of Prosper's identity claims—either, that purpose is not relevant in lender decisions, or that the purpose labels used are not the 'right' ones—then there should not be evidence of lender reliance on the focus of purpose claims.

Predicting the purpose of listings in the pre-label period is a key component of my empirical strategy. While human coding of listing text is possible, my interest in the coherence of borrower discourse requires (a) continuous measures of category membership, and (b) measures calibrated to local market conditions—that is, taking the coherence of other listings into account. The process described below satisfies these criteria.

I retrieved and cleaned text from listing titles and descriptions and prepared a sparse document-term matrix, wherein rows correspond to listings, columns correspond to features, and values of the matrix are the number of times a given listing uses a given term. I then model a listing's purpose label as a function of term frequencies. The exercise is 'supervised' machine learning in that I use labeled listings as a guide. The process involves choices of inputs (feature selection) and models. To arrive at a useful model, I use ten-fold cross-validation. For a given set of features or model specification, I iteratively partition the labeled listings into 'training' and 'test' groups (90% and 10% of labeled observations respectively), train the model on the training group, and assess model performance on the test group. I repeat this process ten times for each configuration and aggregate performance measures. Since test data are not used in the preparation of a given model, accurately predicting the labels of these data should give more confidence about performance with unlabeled data. Figures A.1 and A.2 in the Appendix describe the process graphically.

Below, I briefly describe the process of feature and model selection.

Feature Selection The inputs or 'features' for classification models may include single words or word combinations (e.g., "not bad"), word positions, parts of speech or linguistic structure, or combinations of thereof. The model presented below utilizes single words as well as pairs and triples of words (bigrams and trigrams). Table 1 summarizes the distributions of word counts before and after labeling.

Not all words are diagnostic of loan purpose. In the process of modeling purpose labels, I filtered the features to the set that provided the best available out-of-sample prediction accuracy. Opportunities for filtering include (a) the decision to use titles, descriptions, or both; (b) what sections of the description to use; and (c) more granular feature selection steps, including removing uninformative words ('stopwords'), stemming, and term frequency.

Text Fields. When preparing a description, borrowers encounter template which includes the prompts "Purpose of loan:", "My financial situation:", and several prompts for monthly budget figures, such as rent and debt expenses. The prompts were not compulsory and users could delete the template text, although reading hundreds of examples suggests descriptions generally follow this order. I found that the purpose and financial situation section both improved prediction performance, while monthly budget fields reduced it considerably (as much as an 11% loss in accuracy).

Titles and Descriptions. Both titles and descriptions may suggest the purpose of the loan.[4] My final model utilizes features from both titles and descriptions and performs better than models using only titles or only descriptions. Before cleaning, listing titles had 5.09 words on average (s.d. 2.5; min. 0, max. 17), with an average description word count of 182.16 (s.d. 127.18, min. 0, max. 763).[5]

Other Feature Selection Steps. To reduce overfitting, I constrained features to those that occurred at least 5 times both during the analysis period (untreated and treated periods) and the training period. In some models, I removed uninformative stopwords such as articles and prepositions. Stemming words (e.g., collapsing 'debt', 'debts', and 'debtor' to the stem, 'debt, or 'borrow', 'borrows', and 'borrowing' to the stem, 'borrow') did not aid performance. The final feature set consisted of 212,942 features—a notable decrease from the 1,582,826 features in the model without the 5-times-or-more constraint.

Model Selection. Maximum entropy classification (Manning and Schütze 1999) provided the best out-of-sample prediction performance of all machine learning approaches employed. Using the `maxent` package in R, I trained multinomial logistic regression models using a matrix representation of term frequencies and vector of known labels as inputs. To help ward against overfitting, I explored different regularization strategies. The model providing the best out-of-sample prediction performance used L2 regularization. After selecting the features above through cross-validation, I trained the final model using all labeled data.

Label Assignment and Coherence

For each listing, I predict the probability of membership in each of the seven purpose categories. I assign unlabeled listings to the category with the highest predicted probability. Labels are available for the latter four months of the data, but recall that I am interested in borrowers' positions in discourse space as defined by Prosper's

[4] I observe that some borrowers summarize why they need money in the title, and then elaborate in the description. Others use the title to convey more personally-identifying information, such as relationship status, geographic region, or social affiliation.

[5] It is possible for borrowers to leave titles or descriptions blank; after cleaning, more titles or descriptions may become blank. Maximum entropy classification assigns featureless listings as the most populous category observed in the training data.

purpose categories. Accordingly, I use the predicted purpose when referring to a listing's category membership.

The distribution of the predicted probabilities of membership is informative for the manner of claims a borrower is making. Consider two loan candidates: one with a high probability of membership in one category and low probabilities in the other six, and another candidate with moderate probabilities in three categories and low probabilities in the others. These candidates differ in how readily they are identified with one purpose claim, my theoretical construct of interest. To summarize a borrower's distribution of predicted category membership, I calculate the Herfindahl index of these probabilities:

$$coherence = \sum_{i=1}^{6} p_i^2$$

where p_i is the predicted probability that a listing is in category i.[6] Table 1 reports the distribution of this maximum probability and coherence across labeled and unlabeled listings. After describing the matching procedure below, I present example listings with varying levels of coherence (see Table 5).

Table 1 shows that, on average, listings in the labeled period receive higher probability scores and exhibit higher coherence (concentration in one purpose category) in the labeled period. This may reflect increased conformity of borrowers to categories, and it may also reflect differences in prediction accuracy across the two periods. I approach this issue in several ways. First, I reduce error in label prediction by ensuring a threshold of term frequency in the description corpus: all unigrams, bigrams, or trigrams must appear at least five times in the training and analysis periods. This helps ensure common support among features, and it helps remove many that are ultimately uninformative and could otherwise introduce overfitting.[7] Second, in robustness checks, I normalize the measure of coherence by weekly moving averages. This transformation incorporates how a given listing's level of coherence compares

[6]Since there are six categories, this measure is bound between 1/6 and 1, where the former would be achieved by a borrower with equal predicted probability in each of the six categories, and the latter is a borrower predicted to be in one category with certainty.

[7]For instance, suppose the bigram 'hello world' is not diagnostic of loan purpose, but that the term happens to appear only a few times in one category, Education, in the training data. This incidental concentration in Education may inappropriately pull listings test data towards Education, resulting in more uniformly distributed predicted probabilities.

Table 2: Top 10 Features for Label Classification

	Debt. Consolid.	Home Improv.	Business	Personal/Other	Education	Auto
1	debt	roof	inventory	vacation	college	car
2	debts	kitchen	capital	wedding	student	auto
3	free	garage	advertising	emergency	paying	buying
4	loans	addition	business	bills	school	buying car
5	credit	putting	equipment	personal	degree	vehicle[1]
6	cards	drive	payroll	surgery	tuition[1]	motorcycle
7	payoff	home	business[1]	furniture	education[1]	vehicle
8	chance	pool[1]	shop	boat	education	transportation
9	bills	room	design	repair	school[1]	auto loan[1]
10	consolidate[1]	adding	operating	daughter	classes[1]	car[1]

Note: Features with the highest weights for each purpose category, estimated through maximum entropy. The training data are all labeled listings created between December 8, 2007 and June 27, 2008. The superscript ([1]) indicates the feature is from Purpose or Financial Situation section. Otherwise, the feature is the from the listing title.

to that of contemporaneous listings, and should reduce bias attributable to absolute changes in coherence (and coherence measurement) across the labeled and unlabeled periods. Third, in further robustness checks, I match on listing attributes and on bins of coherence to ensure results do not merely reflect differences in the distribution of coherence. I find a consistent pattern of results.

Comparison of Purpose Categories

Table 2 reports for each purpose label the features that received the highest weights. These will be terms that were most highly associated with the purpose categories. While most top ten words seem to refer to what the borrower will do with the money, it is interesting that debt consolidation includes ‘free’ and ‘relief’, desired end states. Some overlap is evident, as seen in ‘paying bills’ and ‘pay bills pay’ in personal and other. In all, these terms provide face validity for the performance of the classification prediction.

Figure 3 reports the weekly distribution of predicted purpose labels over time. Debt consolidation is consistently the most popular category, followed by personal/other and business. The remaining four categories remain in the minority, averaging around 2% to 3% of the market.

There appears to be some change in the proportion of personal and business listings across market periods. As reported in Table 1, personal loans composed 34% of the unlabeled market and 27% of the labeled market, and business constituted

12% and 16% of each market period, respectively. These changes in the purpose distribution could reflect several scenarios and merit discussion. First, the introduction of labels could invite borrowers with a particular purpose to enter or exit the market. While Prosper did not appear to launch a campaign advertising the loan purpose categories, prospective borrowers browsing the website would be able to see labels among live listings. For instance, some prospective users who may not have considered Prosper as a source of business capital may choose to enter after seeing business listings funded. Second, the juxtaposition of labels could cause some borrowers to rethink their priorities of needs—e.g., some who would have requested a personal loan recognize that a business loan is the more pressing need. Third, in an extreme version of the second scenario, borrowers may infer a purpose preference among lenders for type of loan and may misrepresent their purpose needs accordingly. Fourth, the changes could reflect market shocks less amenable to control.

Recall that I am more interested in a property of features—coherence—than I am in specific categories themselves. Thus, I am concerned about these scenarios only to the extent that they introduce an alternate explanation for a relationship between coherence, funding, and labeling. First, if labels invite sorting into and out of the market, then I should be able to account for changes in observable attributes of listings both in my regression analyses and in my non-parametric matching strategy. To the second and third points: it is unclear that reprioritization or strategic manipulation would necessarily result in coherent or diffuse descriptions, and even if they did, it is less obvious how such borrowers would then leak some signal beyond observables that would be picked up by lenders.

Figure 4 presents a plot of average 100-quantiles of coherence over time. The plot shows between-category differences in levels of coherence. On average, Business listings present descriptions that have higher levels of coherence than other categories, followed by Debt Consolidation. The other remaining categories are generally below average on coherence in any given week. There does appear to be a slight increase in coherence of listings overall after labels are introduced, suggesting that the labels may have altered borrowers' general approach to self-description. While the increase is not sizeable, I address the potential issue of simultaneous effects of labeling by matching on coherence in robustness checks.

An alternate measure of (in)coherence is the second-highest predicted probability. For example, a listing that has predicted probabilities of 0.40 and 0.35 of being in Debt and Education, respectively, can be considered to be split between these purpose categories. This measure more directly corresponds to the concept of hybridity (Ruef and Patterson 2009). Second-highest probability is reported in Table 1. Consistent with the dynamics observed for coherence, there is a slight decrease in second-highest

Figure 3: Dynamic Plot of Purpose Frequency

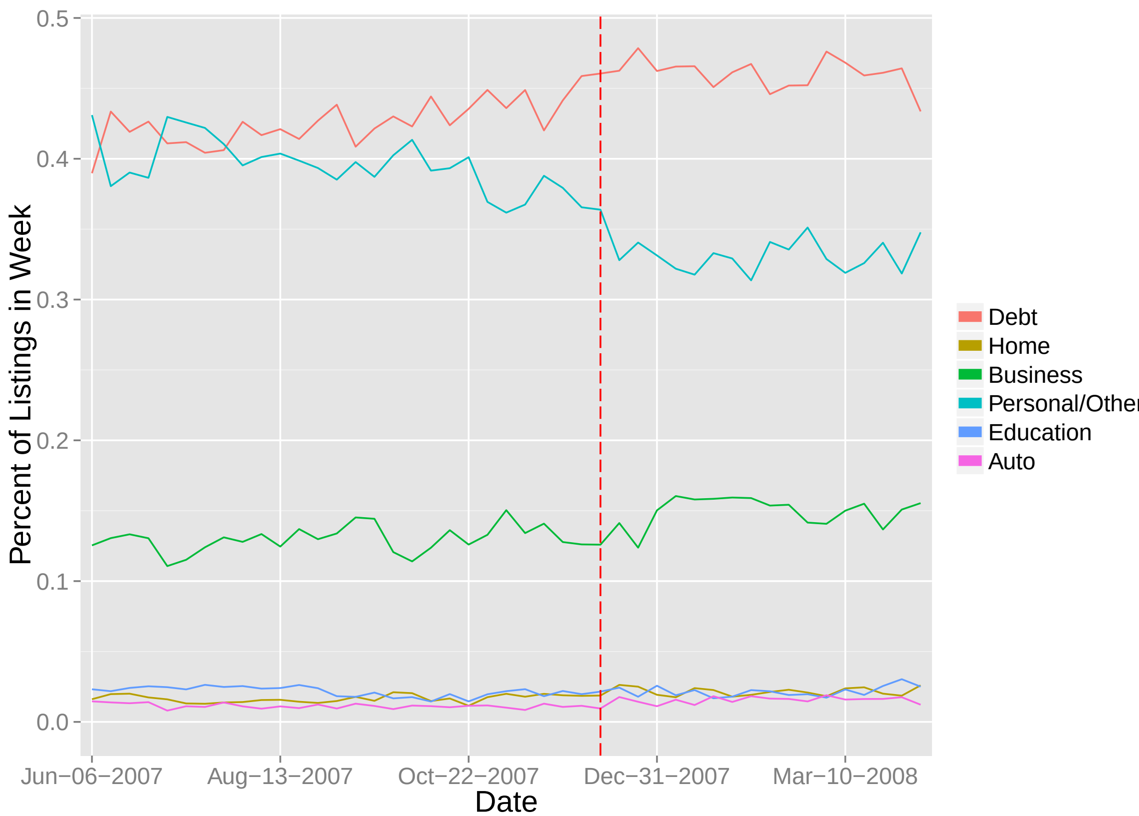

Note: Plot of weekly percent of listings orginated with a given predicted purpose category. Dates range from June 6, 2007 to April 14, 2008. The dashed red line represents when purpose labels were introduced (December 5, 2007).

probability between the unlabeled and labeled analysis periods. In ancillary analyses, I examine second-highest probability as an alternate measure.

In the following chapter, I consider another attribute of feature combinations: the extent to which features are typical of a candidate's typicality. I describe the measurement and effects of typicality in more depth in Chapter 3, but report the levels of typicality here for completeness.

Figure 4: Dynamic Plot of Coherence by Predicted Category

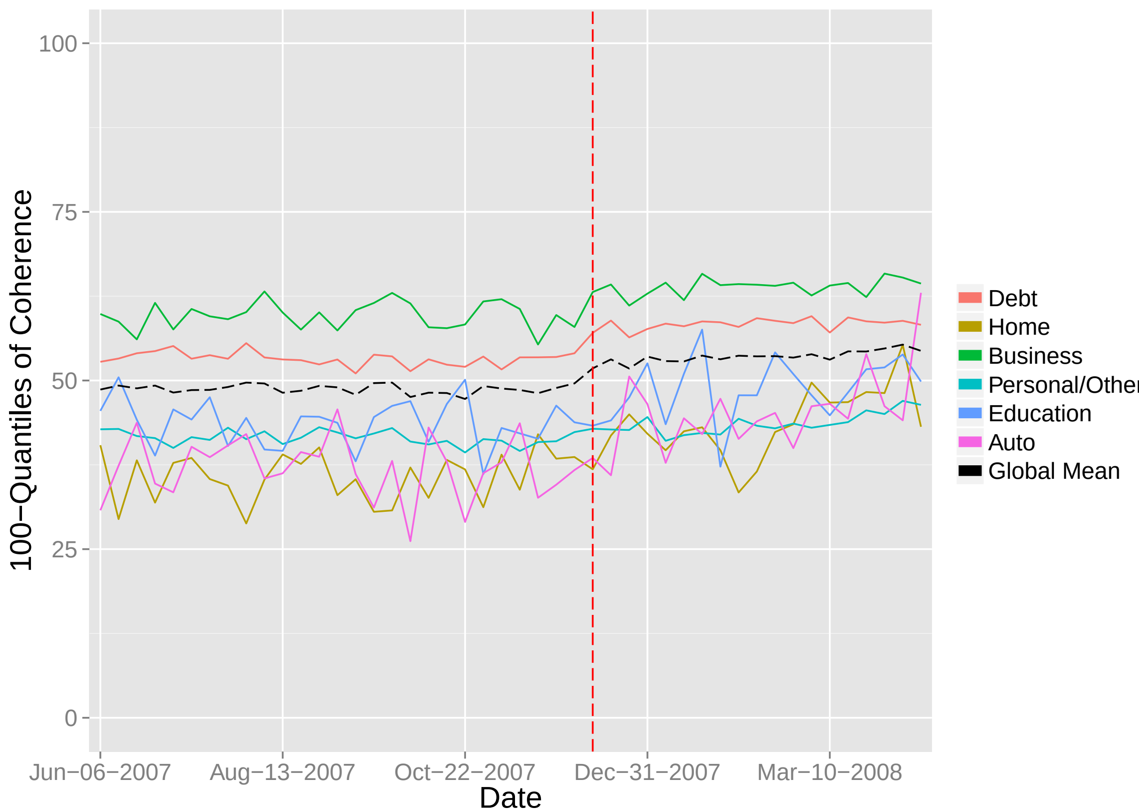

Note: Dynamic plot of 100-quantiles of listing coherence, by purpose category. Coherence is a measure of the distribution of a listing's prediction scores for each purpose label. Dates range from June 6, 2007 to April 14, 2008. The dashed red line represents when purpose labels were introduced (December 5, 2007). The dashed black line describes the global mean level of coherence.

Matching

To help ward against quality confounds in my analysis, I use coarsened exact matching (Iacus, King, and Porro 2012). The motivation for this method is to improve the comparability of observations in the 'treated' and 'untreated' groups. While including controls in regression analyses will help account for between-listing differences, if there are observations in either treated or control groups that would occur with zero probability in the other, then a linear specification will not 'condition away' the confound. Matching pares down a sample to observations that could have credibly occurred in either condition. This restricts the effective population from which inference is drawn, but reduces the bias contributed by observable characteristics.

In view of covariate imbalance suggested by Table 1, I construct a matched sample by performing exact or coarsened-exact matching on the following listing attributes: Prosper credit grade (a bin of Experian credit scores), predicted purpose category, funding option (whether the listing was closed immediately when funded, or remained open for the full stated duration, to collect competitive bids), amount requested, and number of title and description words. Table 3 summarizes descriptive statistics for the matched sample, which is comprised of 98,052 listings, down from 131,436 listings in the full sample. Covariate balance is substantially improved among variables of interest: among matching criteria variables described above, differences in means are no longer statistically significantly different at the 0.05 level. The sample size is sufficiently large that most other variables are statistically significantly different: notably, the ostensibly small difference in mean levels of coherence is significantly different ($t = 8.44$; $df = 97869.63$; $p < 0.001$). Ideally, a matched-sample design such as this will achieve perfect covariate balance, but in a setting such as Prosper, the high number of covariates incurs the 'curse of dimensionality'. If all covariates were balanced, I could trivially regress percent funded on an interaction of coherence and labels; since not all covariates are balanced, I retain covariates in regressions below. Also, in robustness checks described below, I vary the set of variables used to match. Particularly, I include coherence among the matching criteria. Results are robust to different choices of matching variables. Table 4 reports correlations for the matched sample.

The most notable change across Tables 1 and 3 is in the distribution of percent funded. Whereas Table 1 describes an increase in average percent funded across market periods (11% to 15%), Table 3 reports a decrease (17% to 15%). This is mostly likely attributable to the change in amount requested between the unlabeled and labeled markets. Whereas naïve comparisons of means in Table 1 would suggest the average effect of labeling on percent funding was positive, Table 3 suggests this effect was confounded with changes in amount requested. Matching helps avoid such sources of bias in estimation of the moderating effect of labeling.

Table 5 provides examples of listings drawn from the same matching stratum: listings that requested between \$2,500 and \$3,000, appeared in the 'high risk' (HR) credit grade, remained open for the fully stated duration (i.e., the listing would accept competitive bids after fully funded), used between 40 and 55 words in the title and description, and were predicted to be in the Debt Consolidation category. The first two examples and last two examples are drawn from the pre-label (Label

Table 3: Descriptive Statistics: Matched Listings

	Pre-Label Listings (N = 49,026)			Labeled Listings (N = 49,026)			All	
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.	Min.	Max.
Percent Funded	0.17	0.01	0.33	0.15	0.01	0.32	0.00	1.00
Nb. Words	187.38	147.00	128.13	179.40	142.00	123.60	1.00	770.00
Nb. Words (excluding budget)	147.07	114.00	113.41	147.30	113.00	114.12	1.00	770.00
Amount Requested	7324.35	5000.00	6494.10	7314.06	5000.00	6386.71	1000.00	25000.00
Max. Borrower Rate	0.18	0.17	0.07	0.19	0.18	0.09	0.00	0.36
Debt to Income Ratio[1]	0.54	0.26	1.32	0.43	0.26	0.98	0.00	10.01
Credit Grade: AA (760 and up)	0.04	0.00	0.19	0.04	0.00	0.19	0.00	1.00
Credit Grade: A (720-759)	0.05	0.00	0.22	0.05	0.00	0.22	0.00	1.00
Credit Grade: B (680-719)	0.08	0.00	0.27	0.08	0.00	0.27	0.00	1.00
Credit Grade: C (640-679)	0.14	0.00	0.34	0.14	0.00	0.34	0.00	1.00
Credit Grade: D (600-639)	0.18	0.00	0.38	0.18	0.00	0.38	0.00	1.00
Credit Grade: E (560-599)	0.18	0.00	0.38	0.18	0.00	0.38	0.00	1.00
Credit Grade: HR (520 to 559)	0.34	0.00	0.47	0.34	0.00	0.47	0.00	1.00
Duration	7.67	7.00	1.99	7.51	7.00	2.27	3.00	10.00
Homeowner	0.36	0.00	0.48	0.39	0.00	0.49	0.00	1.00
In Prosper Group	0.24	0.00	0.43	0.13	0.00	0.33	0.00	1.00
Automatic Funding	0.09	0.00	0.29	0.09	0.00	0.29	0.00	1.00
Nb. Attempt by Member	2.89	2.00	3.07	3.34	2.00	4.62	1.00	68.00
Number Live by Start	2621.32	2539.00	340.44	2372.96	2295.00	442.12	1684.00	3971.00
Percent Live in Same Category	0.36	0.45	0.16	0.32	0.27	0.16	0.01	0.52
Predicted Label: Debt	0.47	0.00	0.50	0.47	0.00	0.50	0.00	1.00
Predicted Label: Home	0.02	0.00	0.13	0.02	0.00	0.13	0.00	1.00
Predicted Label: Business	0.15	0.00	0.35	0.15	0.00	0.35	0.00	1.00
Predicted Label: Personal/Other	0.34	0.00	0.47	0.34	0.00	0.47	0.00	1.00
Predicted Label: Education	0.02	0.00	0.13	0.02	0.00	0.13	0.00	1.00
Predicted Label: Auto	0.01	0.00	0.10	0.01	0.00	0.10	0.00	1.00
Coherence	0.74	0.78	0.22	0.78	0.85	0.21	0.17	1.00
Second-Highest Probability	0.14	0.10	0.14	0.13	0.07	0.14	0.00	0.50
Typicality[2]	51.85	53.00	29.14	51.00	51.00	29.15	1.00	100.00

N Listings = 98,052. Note: Descriptive statistcs for listings that appeared on Prosper's website from June 6, 2007 through April 14, 2008. Labeled and unlabeled listings were matched exactly on Prosper credit grade, predicted category, and funding option, and coarsened-exact matched on amount requested and word count.
[1] Missing values for 6675 observations.
[2] Missing values for 7 observations. Typicality is the subject of Chapter 3.

= 'No') and post-label (Label = 'Yes') periods, respectively, and the examples vary in coherence. The first and third listings demonstrate lower ranges of the coherence distribution. Both have somewhat vague titles and incorporate personal details (e.g., family, medical status) in describing their debt consolidation plans.

Estimation Strategy

Listings vary in the attention they receive from lenders. One measure of this attention is the percentage of a borrower's requested amount that is funded by lenders (percent

Table 4: Correlations: Matched Listings

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1 Percent Funded														
2 Predicted Label: Debt	-0.02													
3 Predicted Label: Home	0.03	-0.12												
4 Predicted Label: Business	0.05	-0.39	-0.05											
5 Predicted Label: Personal/Other	-0.01	-0.67	-0.09	-0.30										
6 Predicted Label: Education	-0.02	-0.13	-0.02	-0.06	-0.10									
7 Predicted Label: Auto	-0.01	-0.10	-0.01	-0.04	-0.08	-0.01								
8 Coherence	0.04	0.14	-0.05	0.11	-0.19	-0.02	-0.05							
9 Typicality[1]	-0.10	0.18	0.02	-0.15	-0.13	0.07	0.13	0.02						
10 Second Probability	-0.04	-0.07	0.02	-0.12	0.14	0.01	0.03	-0.90	0.00					
11 Nb. Words	0.09	-0.02	-0.02	0.11	-0.05	0.01	-0.02	0.21	-0.43	-0.16				
12 Nb. Words (excluding budget)	0.09	-0.04	-0.02	0.13	-0.05	0.01	-0.03	0.22	-0.44	-0.16	0.95			
13 Amount Requested	0.02	-0.01	0.03	0.23	-0.15	-0.05	-0.04	0.08	-0.02	-0.06	0.02	0.05		
14 Max. Borrower Rate	0.12	0.00	0.00	0.02	-0.01	0.00	0.01	0.03	-0.07	-0.02	0.12	0.11	-0.08	
15 Debt to Income Ratio[2]	-0.07	-0.02	0.00	0.02	0.01	0.01	-0.01	-0.02	0.06	0.01	-0.03	-0.03	0.07	0.00
16 Credit Grade: AA (760 and up)	0.24	-0.07	0.03	0.10	0.00	-0.02	-0.02	0.01	-0.06	-0.02	-0.01	0.01	0.22	-0.16
17 Credit Grade: A (720-759)	0.18	-0.02	0.01	0.08	-0.03	-0.02	-0.02	0.04	-0.04	-0.04	0.01	0.03	0.23	-0.12
18 Credit Grade: B (680-719)	0.18	0.00	0.01	0.07	-0.04	-0.02	-0.02	0.01	-0.01	-0.02	-0.01	-0.01	0.21	-0.09
19 Credit Grade: C (640-679)	0.09	0.01	0.02	0.06	-0.05	-0.02	-0.01	0.02	0.00	-0.02	0.00	0.00	0.11	-0.03
20 Credit Grade: D (600-639)	-0.01	0.04	-0.01	-0.02	-0.02	-0.01	0.00	0.02	0.02	-0.02	0.02	0.01	-0.02	0.05
21 Credit Grade: E (560-599)	-0.10	0.04	-0.02	-0.06	0.01	-0.01	-0.02	-0.01	0.01	0.01	0.01	0.00	-0.13	0.10
22 Credit Grade: HR (520 to 559)	-0.26	-0.03	-0.02	-0.09	0.08	0.05	0.05	-0.05	0.03	0.05	-0.01	-0.02	-0.27	0.07
23 Duration	0.03	0.03	0.00	0.03	-0.05	0.00	-0.01	0.04	-0.05	-0.02	0.11	0.09	0.08	0.01
24 Homeowner	0.11	-0.01	0.09	0.07	-0.05	-0.05	-0.04	0.01	-0.02	-0.01	-0.01	0.00	0.22	-0.07
25 In Prosper Group	0.11	0.04	0.00	0.00	-0.04	-0.01	-0.01	0.03	-0.15	-0.03	0.28	0.23	-0.07	0.10
26 Automatic Funding	-0.02	0.00	-0.02	-0.03	0.03	0.01	0.00	0.00	-0.04	0.00	0.06	0.06	-0.10	0.00
27 Nb. Attempt by Member	-0.01	0.04	-0.01	-0.04	-0.01	-0.01	-0.01	0.01	-0.12	-0.02	0.20	0.18	-0.14	0.16
28 Number Live by Start	-0.03	0.01	-0.01	-0.01	0.00	0.00	0.00	-0.04	-0.01	0.02	0.01	0.00	-0.01	-0.04
29 Percent Live in Same Category	-0.04	0.59	-0.20	-0.44	-0.15	-0.20	-0.16	0.07	0.16	0.00	-0.03	-0.06	-0.05	-0.01

	15	16	17	18	19	20	21	22	23	24	25	26	27	28
15 Debt to Income Ratio[2]														
16 Credit Grade: AA (760 and up)	-0.01													
17 Credit Grade: A (720-759)	0.01	-0.05												
18 Credit Grade: B (680-719)	0.01	-0.06	-0.07											
19 Credit Grade: C (640-679)	0.01	-0.08	-0.09	-0.12										
20 Credit Grade: D (600-639)	0.01	-0.09	-0.11	-0.14	-0.18									
21 Credit Grade: E (560-599)	-0.01	-0.09	-0.11	-0.14	-0.18	-0.22								
22 Credit Grade: HR (520 to 559)	-0.01	-0.14	-0.16	-0.21	-0.28	-0.33	-0.33							
23 Duration	-0.01	0.00	0.01	0.02	0.00	0.01	-0.01	-0.02						
24 Homeowner	-0.01	0.16	0.10	0.11	0.13	-0.03	-0.04	-0.22	-0.03					
25 In Prosper Group	-0.01	-0.02	-0.02	-0.01	0.01	0.00	0.03	-0.01	0.10	0.00				
26 Automatic Funding	-0.02	-0.05	-0.05	-0.05	-0.03	0.00	0.04	0.06	-0.04	-0.02	0.03			
27 Nb. Attempt by Member	-0.02	-0.07	-0.06	-0.05	-0.03	0.00	0.08	0.04	0.06	-0.02	0.33	0.13		
28 Number Live by Start	0.01	-0.01	0.00	-0.01	0.00	0.00	0.01	0.00	0.00	0.01	0.02	0.02	-0.02	
29 Percent Live in Same Category	-0.01	-0.09	-0.05	-0.01	0.00	0.03	0.06	-0.01	0.03	-0.03	0.03	0.00	0.01	0.06

N Listings = 98,052. Note: Correlations for variables utilized in regressions. Data are from samples constructed through coarsened exact matching. All correlations greater than 0.01 in absolute magnitude are statistically significant at the 0.05 level.

[1] Missing values for 7 observations.

[2] Missing values for 6675 observations.

funded). I model percent funded as a function of listing attributes in generalized linear models of the following form:

Table 5: Sample Listings from Same Matching Stratum

Labeled	Coherence (100-Quantile)	Category (Probability)	Description
No	0.484 (13th)	Debt (0.556)	Title: In Great Need for HELP
		Home (0.010)	Purpose of loan: Bill Consolidation
		Business (0.003)	My financial situation: Single Mom of a 2 year
		Personal/Other (0.419)	old. I would love to consolidate my bills into one,
		Education (0.005)	and spend more of a peaceful lifestyle with my son.
		Auto (0.008)	Your help is greatly appreciated.
No	0.985 (84th)	Debt (0.993)	Title: Small CC debt
		Home (<0.001)	Purpose of loan: Need this loan to pay off my old
		Business (<0.001)	Credit card Debt
		Personal/Other (0.007)	My financial situation: I am working full time and
		Education (<0.001)	am in a comfortable position to pay off this loan.
		Auto (<0.001)	More over i do have enough savings though i just
			need them as any emergency cushion
Yes	0.508 (18th)	Debt (0.617)	Title: My Loan
		Home (0.005)	Purpose of loan: This loan will be used to pay
		Business (0.013)	down uninsured medical bills and credit cards used
		Personal/Other (0.357)	to pay off medical bills
		Education (0.003)	My financial situation: I am a good candidate for
		Auto (0.004)	this loan because I pay my bills on time.
Yes	0.996 (91st)	Debt (0.993)	Title: Paying off loans
		Home (<0.001)	Purpose of loan: This loan will be used to Consol-
		Business (<0.001)	idate Loans
		Personal/Other (0.007)	My financial situation: I am a good candidate
		Education (<0.001)	for this loan because I want to make sure that I
		Auto (<0.001)	pay everyone that deserves to be paid back what I
			bor[r]owed when I needed them.

Example listings are from the same matching stratum: requested between $2,500 and $3,000, appeared in the 'high risk' (HR) credit grade, would remain open for the full stated duration, used between 40 and 55 words in title and description, and were predicted to be in the Debt Consolidation category. 'No' in the 'Labeled' column indicates the listing appeared in the unlabeled period of the market; 'Yes' indicates the listing appeared in the labeled period of the market.

$$\begin{aligned} g(E[PercentFunded]) = \beta_0 &+ f_{1j}(CONTROLS_j) \\ &+ \beta_1 LABELED + f_{2k}(PURPOSE_k) \\ &+ f_3(COHERENCE) \\ &+ f_{2k}(PURPOSE_k) \times LABELED \\ &+ f_3(COHERENCE) \times LABELED \end{aligned} \tag{1}$$

where g is a canonical link function, the β_i are fixed parameters, and f_{1j}, f_{2k}, and f_3 are functions of corresponding covariates. Below, the f_{1j} are alternately untransformed variables, flexible indicator functions, or smooth regression splines of the following controls: amount requested, maximum borrower rate, debt-to-income ratio, credit grade, duration, homeownership status, Prosper Group membership, funding scheme (open for duration versus close when funded), and number of title and description words. The f_{2k} consist of five indicators of predicted loan purpose; debt consolidation, the most commonly predicted purpose, is the comparison group.

The motivation for using regression splines stems from (a) the fact that there is not a natural unit for computed variables such as coherence, and (b) the need to assess amplification or attenuation of the conditional effect of coherence on funding. Per the hypotheses described above, I expect a general negative relationship between coherence and percent funding, but hypotheses diverge with respect to the market conditions (labeled or unlabeled) wherein this negative relationship is most pronounced. Strong evidence for the normative account would be no relationship in the pre-label period, and a negative relationship afterwards. Weaker evidence would be negative relationships in both periods, but a more pronounced effect in the labeled period. The cognitive account would be most strongly supported by a clear negative relationship between coherence and labels in the unlabeled period, and a dampening of this relationship in the presence of labels. Ultimately, accentuation and attenuation of negative relationships are of immediate interest, and such changes in relationships are best assessed through graphs in which functional forms are not imposed on estimates. For instance, an attenuation of a negative effect could be complete dampening of an effect, or only dampening for observations not at the extremes of the distribution.

In initial models, I do proceed parametrically. For these models, f_3 is a pair of indicator functions corresponding to the highest and lowest third of the empirical distribution of the coherence measure: the comparison group is the middle 33%. Then, to better visualize the effect of coherence on percent funding in the unlabeled and labeled periods, I estimate a model in which f_3 consists of two suites of thin-plate cubic regression splines (Wood 2004), one suite for each of the unlabeled and labeled periods of the market. I penalize smooth terms under restricted maximum likelihood estimation, which pushes uninformative splines to zero and selects locations for the 'knots' based on the data.[8]

The models presented below assume a Gaussian Normal distribution, where g is the identity link. For fully parametric models (i.e., models not containing regression smooths), estimation may proceed through ordinary least squares. Since the dependent variable, percent funded, is bound between zero and one, I also estimated

[8]Thin-plate smoothing splines do not specify the location of knots, but rather begins with full-plate splines at every data point and then selects splines through penalization. See Wood (2004).

models assuming, alternately, binomial and beta distributions with logit link functions. Model diagnostics suggest models assuming a binomial distribution provide the best fit, but the various models provide highly similar results. For ease of interpretability and for the sake of parsimony, I present models assuming Gaussian Normal distributions below.

RESULTS

Effects of Labels on Funding Table 6 reports results of regression analyses using all listings that originated in the analysis period (the unmatched population). The dependent variable is the sum of dollars bid to a listing divided by the amount requested. Estimates were obtained through ordinary least squares. In these preliminary models, continuous control variables (e.g., amount requested, borrower maximum rate) are entered untransformed with linearity assumed. It is worth noting some of the linear effects recovered from these regressions. Across all models, the amount requested, starting borrower interest rate, credit grade, homeownership, and appearance in a Prosper group are listing attributes that are positively associated with funding. Debt-to-income ratio, the 'open funding' option (e.g., the listing seeks competitive bids before closing), and number of attempts are negatively associated with funding. Of particular interest is a positive association between number of words and funding. Controls for contemporary market conditions recover a positive effect for number of listings live at the time of a focal listing's posting, and, for some models, a negative effect for percent of those listings that are predicted to be in the same category as the focal listing. Models in future tables will enter nonparametric regression smooths for all continuous variables; this will permit consideration of nonlinearities in relationships among controls and funding.

Table 6: OLS Regression: Percent Funded - All Listings

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Intercept	0.003	0.002	-0.018^{**}	0.010	0.009	0.001
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Amount Requested	0.000^{***}	0.000^{***}	0.000^{***}	0.000^{***}	0.000^{***}	0.000^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Borrower Max. Rate	0.827^{***}	0.827^{***}	0.811^{***}	0.826^{***}	0.826^{***}	0.826^{***}
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
DIR	-0.013^{***}	-0.013^{***}	-0.012^{***}	-0.013^{***}	-0.013^{***}	-0.013^{***}
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Credit: AA	0.692^{***}	0.692^{***}	0.685^{***}	0.692^{***}	0.692^{***}	0.692^{***}
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Credit: A	0.531^{***}	0.531^{***}	0.525^{***}	0.531^{***}	0.531^{***}	0.531^{***}

Continued on next page

Table 6 – *Continued from previous page*

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Credit: B	0.429***	0.429***	0.423***	0.429***	0.429***	0.429***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Credit: C	0.259***	0.259***	0.255***	0.259***	0.259***	0.259***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Credit: D	0.135***	0.135***	0.132***	0.135***	0.135***	0.134***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Credit: E	0.040***	0.040***	0.039***	0.040***	0.040***	0.040***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Duration (Days): 3	0.002	0.002	−0.002	0.002	0.002	0.002
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Duration (Days): 5	0.007**	0.007*	0.005	0.007*	0.007*	0.007*
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Duration (Days): 10	0.012***	0.012***	0.011***	0.012***	0.012***	0.012***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Homeowner	0.003	0.003	0.002	0.003	0.003	0.003
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
In Prosper Group	0.050***	0.050***	0.055***	0.050***	0.050***	0.050***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Open Funding	−0.003	−0.003	−0.013***	−0.003	−0.003	−0.003
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Nb. of Attempt	−0.003***	−0.003***	−0.003***	−0.003***	−0.003***	−0.003***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Nb. Words	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Nb. Listings Live	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Perc. Same Category	−0.011	−0.009	−0.017***	−0.026***	−0.025***	−0.009
	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)	(0.006)
Purpose: Home	0.016**	0.021**				0.022**
	(0.006)	(0.007)				(0.008)
Purpose: Business	0.005	0.013***				0.012***
	(0.003)	(0.004)				(0.004)
Purpose: Personal/Other	0.001	0.000				0.001
	(0.002)	(0.002)				(0.002)
Purpose: Education	0.018***	0.012				0.013
	(0.006)	(0.007)				(0.007)
Purpose: Auto	0.019**	0.016				0.017
	(0.007)	(0.009)				(0.009)
Labeled	−0.031***	−0.030***		−0.032***	−0.029***	−0.029***
	(0.002)	(0.002)		(0.002)	(0.003)	(0.003)
Prp: Home × Label		−0.009				−0.010
		(0.011)				(0.011)
Prp: Business × Label		−0.019***				−0.018***
		(0.005)				(0.005)
Prp: Personal/Other × Label		0.004				0.004
		(0.003)				(0.003)
Prp: Education × Label		0.020				0.019
		(0.011)				(0.011)
Prp: Auto × Label		0.010				0.009
		(0.013)				(0.013)
Upper 33% Coherence			0.004*	0.005**	0.008***	0.007**
			(0.002)	(0.002)	(0.002)	(0.002)
Lower 33% Coherence			0.000	−0.002	−0.002	−0.002
			(0.002)	(0.002)	(0.002)	(0.002)

Continued on next page

Table 6 – *Continued from previous page*

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Upper 33% Coherence × Label					−0.007	−0.005
					(0.004)	(0.004)
Lower 33% Coherence × Label					0.001	0.001
					(0.004)	(0.004)
Log Likelihood	-3308.205	-3294.040	-3479.704	-3310.734	-3308.004	-3283.738
R^2	0.306	0.306	0.304	0.306	0.306	0.307
Num. obs.	122200	122200	122195	122195	122195	122195

$^{***}p < 0.001, ^{**}p < 0.01, ^{*}p < 0.05$

Model 1 of Table 6 recovers varying effects for purpose categories. Compared to the reference category, Debt Consolidation, listings predicted to be in the Home, Education, and Auto categories generally receive more funding. As Model 2 introduces interactions with the effect of labels, the effects for Education and Auto are attenuated. Also, a complex relationship for Business listings is revealed: Business listings generally fare better than Debt Consolidation listings, but this benefit is attenuated when labels are present. Home Improvement listings generally perform better than Debt Consolidation listings regardless of whether labels are present.

Models 3 through 6 attend to effects of more pressing interest for this chapter. Model 3 omits purpose category indicators and introduces indicators of whether listings are in the upper third or lower third of the coherence distribution. On average, highly coherent listings perform better than those in the middle of the distribution. This effect is robust to the inclusion of a main effect of labels in Model 4. Model 5 includes interactions between coherence levels and the label indicator. This model provides no evidence for an interaction between coherence and labeling, though the positive main effect for high coherence holds. Model 6 assesses whether the coherence effect is attributable to differences in purpose categories by reinstating the purpose category indicators and label interactions. The main effect for high coherence is robust.

While illustrative, the regressions in Table 6 are inadequate to assess the relationship among coherence, labels, and funding. As described above, the Prosper marketplace underwent other changes at the same time that labels were introduced: for example, the amount requested decreased. To address these potential confounds, I replicate the regression analyses of Table 6 using the matched sample described in Table 3. Table 7 reports these results.

Table 7: OLS Regression: Percent Funded - Matched Listings

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Intercept	−0.003	−0.005	−0.044***	0.005	0.004	−0.006
	(0.008)	(0.008)	(0.007)	(0.008)	(0.008)	(0.008)
Amount Requested	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Borrower Max. Rate	0.962***	0.962***	0.938***	0.961***	0.961***	0.961***
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
DIR	−0.017***	−0.017***	−0.016***	−0.017***	−0.017***	−0.017***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Credit: AA	0.715***	0.715***	0.714***	0.714***	0.714***	0.715***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Credit: A	0.555***	0.555***	0.553***	0.554***	0.554***	0.554***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Credit: B	0.448***	0.448***	0.446***	0.447***	0.447***	0.447***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Credit: C	0.269***	0.269***	0.268***	0.269***	0.269***	0.269***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Credit: D	0.142***	0.142***	0.141***	0.141***	0.141***	0.142***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Credit: E	0.049***	0.049***	0.049***	0.049***	0.049***	0.049***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Duration (Days): 3	0.006	0.006	0.001	0.006	0.006	0.006
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Duration (Days): 5	0.008*	0.008*	0.005	0.008*	0.008*	0.008*
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Duration (Days): 10	0.019***	0.019***	0.017***	0.019***	0.018***	0.019***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Homeowner	0.000	0.000	−0.001	0.001	0.001	0.000
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
In Prosper Group	0.061***	0.060***	0.069***	0.061***	0.060***	0.060***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Open Funding	−0.022***	−0.022***	−0.022***	−0.022***	−0.022***	−0.022***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Nb. of Attempt	−0.004***	−0.004***	−0.004***	−0.004***	−0.004***	−0.004***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Nb. Words	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Nb. Listings Live	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Perc. Same Category	−0.008	−0.004	−0.010	−0.022***	−0.021***	−0.005
	(0.008)	(0.008)	(0.006)	(0.006)	(0.006)	(0.008)
Purpose: Home	0.029***	0.063***				0.066***
	(0.008)	(0.010)				(0.010)
Purpose: Business	0.002	0.014**				0.012*
	(0.004)	(0.005)				(0.005)
Purpose: Personal/Other	0.002	0.001				0.004
	(0.002)	(0.003)				(0.003)
Purpose: Education	0.011	−0.005				−0.003
	(0.008)	(0.010)				(0.010)
Purpose: Auto	0.022*	0.020				0.024
	(0.009)	(0.013)				(0.013)
Labeled	−0.036***	−0.034***		−0.038***	−0.035***	−0.032***
	(0.002)	(0.003)		(0.002)	(0.003)	(0.004)
Prp: Home × Label		−0.064***				−0.068***

Continued on next page

Table 7 – *Continued from previous page*

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
		(0.014)				(0.014)
Prp: Business × Label		−0.021***				−0.019***
		(0.006)				(0.006)
Prp: Personal/Other × Label		0.003				0.000
		(0.004)				(0.004)
Prp: Education × Label		0.033*				0.031*
		(0.014)				(0.014)
Prp: Auto × Label		0.006				0.001
		(0.018)				(0.018)
Upper 33% Coherence			0.005*	0.007**	0.015***	0.015***
			(0.002)	(0.002)	(0.003)	(0.003)
Lower 33% Coherence			−0.002	−0.004	−0.007*	−0.007*
			(0.002)	(0.002)	(0.003)	(0.003)
Upper 33% Coherence × Label					−0.015***	−0.014**
					(0.004)	(0.004)
Lower 33% Coherence × Label					0.007	0.007
					(0.004)	(0.004)
Log Likelihood	-11008.681	-10986.274	-11190.668	-11007.494	-10994.761	-10962.238
R^2	0.302	0.303	0.300	0.302	0.303	0.303
Num. obs.	91377	91377	91371	91371	91371	91371

$^{***}p < 0.001, ^{**}p < 0.01, ^{*}p < 0.05$

The regressions in Table 7 generally replicate the main effects for controls reported earlier. The pattern of results for purpose categories is somewhat different: now the positive main effect for Business *and* Home Improvement (relative to Debt) is entirely offset by the interaction with labels. Model 3 through 6 recover positive main effects of high coherence as before, but this effect is offset by the appearance of labels. Also, with the inclusion of label interactions, there is a negative main effect for low coherence.

The relationship between coherence, labels, and funding is best understood through semiparametric regression. In another model, I replaced the indicators of coherence and label interactions with two suites of thin plate regression splines, with one suite corresponding to the main effect of coherence and the other to the interaction with labeling. I also entered smooths for all continuous controls. Spline terms were penalized under restricted maximum likelihood to prevent overfitting.

Figure 5: Smooth Estimates of the Effect of Coherence on Percent Funded

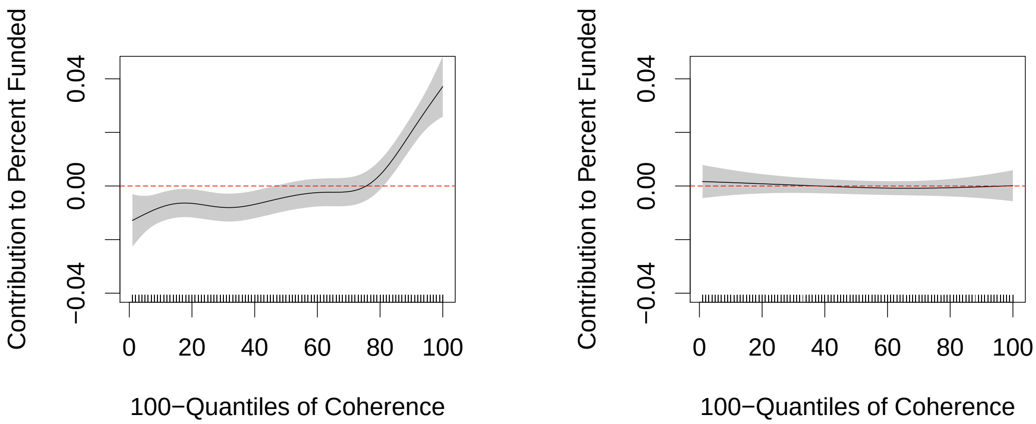

(a) Unlabeled (b) Labeled

Note: The solid black lines correspond to smooth term estimates stemming from iterative reweighted least squares regression.

Figure 5 visualizes the smooth estimates of the effect of coherence on percent funded in each market period. Panel (a) provides nuance to the parametric effects for coherence in Models 3 through 6 of Table 7 high levels of coherence receive more funding, and low levels of coherence receive less, and 95% confidence interval bands exclude zero at both extremes. In Panel (b), confidence bands generally include zero. This pattern of results suggests that labels give actors license to broaden claims in discourse.

Across Models 1 through 2 and 3 through 5, I recover a negative main effect of appearing in the labeled period of the market. Recall that identification of this effect is difficult since the empirical design does not afford calendar fixed effects to account for time-specific market shocks. My main interest is in the relationship between coherence and percent funded, similarity and percent funded, and the interaction with labels. Accordingly, omitted variables are problematic only insofar as they have bearing on the relationship between coherence and similarity and funding. For instance, the rise of competing crowdfunding platforms to Prosper is serially correlated with time, which could account for lower conditional funding rates in the labeled period. Yet we would not expect platform competition to alter the evaluation of coherence.

Mechanism: Risk Interaction

To this point, the results are consistent with a cognitive account of feature evaluation. Under a normative account, which suggests that labels are necessary to coordinate audience expectations of category membership, we would expect feature coherence to explain more variance in the performance of candidates *after* labels are imposed than before. Instead, results suggest that feature coherence separates candidates most in the period of the market before labels are present, and then explains no significant variance after labels are present. Having established the cognitive account, I conduct additional analysis below to illuminate mechanisms at work. Particularly, I exploit candidate variation in borrower risk.

Lenders face general risk in participating on Prosper: loans are unsecured, and minimal reparations are available in the event of a default. The quality of borrowers is uncertain. Prosper sought to stratify candidates by riskiness by collecting creditworthiness information from each prospective borrower, including their credit score, debt-to-income ratio, and home ownership status. Prosper's credit grades, which were constructed based on credit scores, provide an especially useful source of sample variation that may help identify the mechanism at work in feature valuation. Under a normative account, feature coherence is a signal of quality that audiences will attend to most under conditions of uncertainty. Borrowers that have very high credit and very low credit have more certain quality than those in the middle. Thus, under a normative account, we would expect the moderation effects reported in Figure 9 to vary with signals of borrower risk. Particularly, for those with extremes in creditworthiness, feature coherence will be least informative, and most informative for those with middle creditworthiness. Under a cognitive account, however, the moderation effects should not attend to lender uncertainty regarding quality.

I decompose my sample by Prosper credit grade, grouping "AA" through "B", "C" through "E", and "HR" (high risk)—roughly 17%, 49%, and 34% of the matched sample, respectively. I then replicate the regression analysis estimating the effect of coherence on percent funded for each subset. Figure 5 provides plots of regression smooths obtained from these regressions. I find that coherence provides no information to listings in the lowest credit grades, while it is most informative for those with middle and highest credit grades. See Figure 5.

Figure 6: Smooth Estimates of the Effect of Coherence by Credit Grade

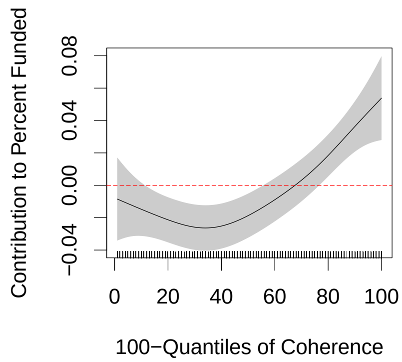

(a) Unlabeled: AA, A, B

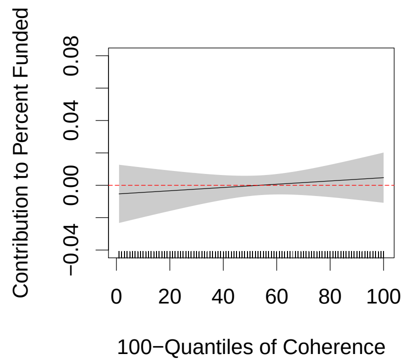

(b) Labeled: AA, A, B

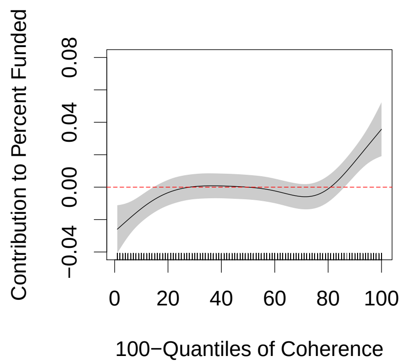

(c) Unlabeled: C, D, E

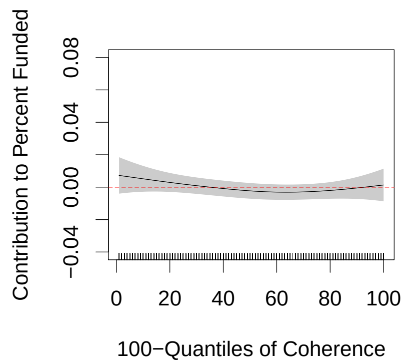

(d) Labeled: C, D, E

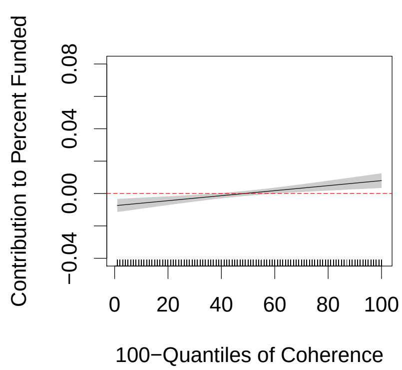

(e) Unlabeled: HR

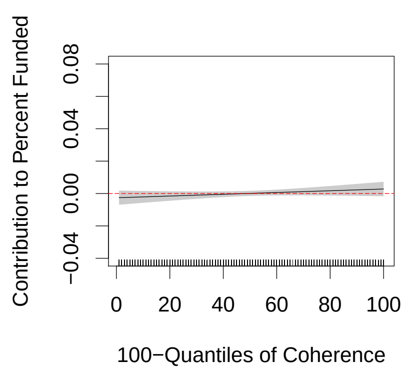

(f) Labeled: HR

The variance in percent funding explained by feature coherence in the pre-label period is proportionate to the variance in percent funding at each credit-grade group. For instance, the high risk (HR) borrowers are hardly ever funded. The decomposition by credit grade shows that the previous results are conservative in that coherence effects persist despite the inclusion of listings with little variation in the dependent variable. They also suggest that for this context, purpose categorization is not 'imperative' in the Zuckerman (1999) sense—that is, quality signals (credit grade) prescribe a general level of variance in evaluation, and purpose coherence contributes to variance within that bound.

This decomposition also lends evidence contrary to a normative account. The group with the highest credit grade (AA through B) has the highest variance in percent funding, and coherence explains the most variance in percent funding than elsewhere. Under a normative account, coherence would have been least important for the performance of borrowers with low risk *and* high risk, since their signal of high quality would not have invited close scrutiny. In all, my results are consistent with the cognitive account that labels give license to the claims actors can make in discourse.

The labels appear to provide a substitute for the information previously conveyed by feature coherence. A natural question is whether labels substitute for all feature information. For instance, individual features may communicate quality information, a dimension that is largely orthogonal to purpose information. For example, in the case of language features, misspellings or grammatical errors may be less egregious—or at the extreme, irrelevant—after labels are present. Yet Gao and Lin (2013), who study such linguisitic features in Prosper directly, find evidence that such attributes as readability and positivity do inform audience decisions. While these authors did not examine the consequences of labeling on these effects, the fact that these results were obtained from the labeled period of the Prosper marketplace suggests labels do not obviate all feature information. In the next chapter, I examine the effects of labeling on another dimension of feature information more directly related to the focus of category scholars: the degree to which features are typical of members of a given category.

Robustness Checks

As described above, listings that appear in the labeled period tend to exhibit higher levels of coherence. This difference may be due to the causal effect of labels on coherence itself, or it may reflect residual measurement error. Both are potential confounds to the labels-coherence moderation effects reported above. Measurement bias is reduced by training the model on listings that appear after the analysis data, but there may be outstanding concerns that the training data is more removed from listings in the unlabeled analysis period than from those in the labeled analysis period. I examine both concerns through the robustness checks described below.

In the analysis above, I calculate 100-quantiles of coherence using the entire distribution of coherence in the analysis period. I refer to this as the 'absolute coherence' approach. An alternate approach is to calculate quantiles within one-week cohorts: the 'relative coherence' approach. The latter reflects the notion that audiences make localized judgments of features, rather than using all previous experience. For example, a listing from week 10 of the analysis period (an unlabeled listing) and a listing from week 40 (a labeled listing) may exhibit the same level of absolute coherence, but may differ in their degree coherence relative to listings that appear contemporaneously. I re-estimated models using the relative coherence approach and find nearly identical results.

A remaining concern is that the matching process may assemble inappropriate comparisons. For example, unlabeled and labeled listing pairs may be matched on attributes such as amount requested and credit grade, but achieving covariate balance in these attributes may come at a cost of imbalance in coherence. To address this potential concern, I constructed another matched sample that included octiles of coherence as a matching dimension. The additional constraint resulted in a loss of about 12,000 observations. However, even with this reduced sample, the pattern of results is highly similar.

Inspection of example listings at the extreme values of coherence suggested that many highly coherent listings made claims to business use. While I include fixed effects for predicted purpose and interactions of purpose with labeling, I also replicate analyses omitting the business category altogether. Results were robust to this constraint.

As described above, an alternate measure of category spanning to consider is the second-highest predicted probability. Low levels of this probability would suggest claims are concentrated in one domain, while high levels would suggest straddling of two domains. I repeated analyses using reversed 100-quantiles of second highest probability. The pattern of interactions produces the expected reversed pattern (reverses, since second-highest probability is inversely correlated with coherence).

DISCUSSION

Before these results can be extended to other contexts, it is useful to consider several aspects of this setting.

First, the actors of interest were individuals. In his early formulations of signaling theory, Spence (1973) observed that signals would be most potent when the actors were in the market relatively infrequently. Organizations, however, seek permanent citizenship status in the market and develop idiosyncratic reputations. New organizations, however, for whom the longevity is still in question, may be similarly novel.

Second, the market communications I studied were supplied by actors themselves. In other settings, discourse is mediated. Previous work has shown that firms attend to rivals' portrayal in the news media over firm public statements Kennedy, 2005 4610. Perhaps my setting usefully simplifies the process, but the introduction of an intermediary invites new questions. While I show that labels liberate the scope of claims, might intermediaries, who may be more aware of categorical distinctions, curtail this range? This is fodder for future inquiry.

Third, my analysis considers discourse claims to particular domains—those that later receive a label by Prosper (debt, home improvement, business, personal, education, auto, and other). Furthermore, many of these labels are broad either by definition (the other category) or by dint of material differences. While I include fixed effects for predicted purpose in my models and show that my results are not attributable to one category, this does not make this set of labels representative of the underlying population of possible labels. I do consider these labels as descriptive of a nominal rather than ordinal position among candidates, which is a feature that can be assessed of other classification systems.

Fourth, the context of seeking to relay identity and creditworthiness may amplify the sensitivity to signals. Ancillary analyses revealed that coherence was most redemptive for borrowers with mid to lower creditworthiness, while for those in the highest credit tier (AA through B), coherence was uninformative in either period of the market. This suggests that other contexts in which there is low altercentric uncertainty will not yield the pattern of results reported here.

Finally, there are limitations to this analysis. As can be expected of an online business, Prosper made multiple and frequent changes to its website, introducing confounding to estimates of main effects of returns to labeling. Accordingly, categorization was not the only change to the site on December 5, 2007. Prosper also implemented a new blog, began providing 'daily differential' downloads of the market data, and linked the discussion forum to the main site. It is not obvious that a blog would have a systematic impact on loan candidate evaluation in the post-categorization period. The daily differential download would have made it easier for users accessing market data to obtain up-to-date information concerning loan payment histories, available listings, and bidding activity, but since APIs require considerable technical know-how, it does not seem that a majority of borrowers or lenders would have been able to exploit this development effectively.

CONCLUSION

Whereas much of previous research in market categorization has focused on the constraints of labels, I find that labels can be liberating in the domain of discourse. Absent of labels, audiences appraise discourse in a manner similar to the categorical imperative: focused identities are rewarded, and diffuse identities are discounted. When actors can provide audiences general cognitive moorings with labels, they may broaden their self-descriptions without penalty. Furthermore, I find preliminary evidence that while coherence criteria is supplanted by labels, lenders still attend to the typicality of listings given a purpose claim. This illustrates that categorization is not entirely dependent on the existence of labels.

Self-description and labels warrant more attention in the sociology of markets for several reasons. First, description constitutes a source of market information that is not immediately conducive to valuation. Description is often idiosyncratic. Individual resource holders need to appraise claims both based on personal preferences, but also on the expected valuations of others. The indeterminacy of text and the resulting reliance on others suggests this is a major source for social processes in price setting.

Second, economic sociology has revisited whether 'laws of the market' proceed from individual preferences, or are principles deliberately constructed by mediators. Two-sided platforms occupy a growing share of the economy. One of the challenges confronting managers of such markets is how to effectively structure the interface between buyers and third-party producers, especially when products are novel and highly heterogeneous. Research in two-sided networks has concentrated on pricing strategies and the magnitudes of network externalities (Hagiu 2009; Tucker 2008). The current paper speaks to the problem of managing the interface between platform sides. The focus on the valuation of public, self-description is especially relevant to platforms featuring creative projects and new ventures, such as crowdfunding markets. In many such markets, the producers or products defy easy categorization. Labels may provide a tool for market makers to present heterogeneous candidates as though they were homogenous, a prerequisite for liquidity (Carruthers and Stinchcombe 1999).

Third, while valuation researchers have previous investigated how markets assign value to heterogeneous candidates, they have rarely discussed how different kinds of candidate signals interact. Recent assessments of the state of valuation research call for explanations of when multiple systems of value are operative simultaneously (Lamont 2012). If labels and language undergo independent tracks of valuation, this could help explain when mixed rules exist.

3 Labels and the Returns to Typicality

ABSTRACT

In chapter 2, I examine the relationship between labels and how self-descriptions are evaluated. Particularly, I study *feature coherence*, or the extent to which a candidate's features are associated with a single or many categories, and how a market's adoption of labels changes the returns to this degree of focus.

The present chapter addresses a different question. Claiming membership in a particular category invites comparison with others currently or previously associated with that category. A candidate's features may be representative of those in the category, or they may diverge from what has appeared before. Being typical of a category can again aid audiences in making evaluations, but atypical actors may appeal to variety-seeking audiences. The dynamics of typicality have been studied recently in economic sociology, but again the general approach has been to use settings with well-established labels. This invites the question, can typicality even be assessed in the absence of labels? If so, how is typicality evaluated with and without labels present?

Extending expectations of the categorical imperative in market categories research, I hypothesize that typicality contributes positively to evaluation, but only in the presence of labels. Again leveraging the natural experiment on Prosper, I construct counterfactual levels of typicality for unlabeled loan listings. For example, suppose a given listing's description suggests it is in the Education category. I assess the typicality of the listing within the Education category. I model main effects of typicality, and I interact these effects with labels. Surprisingly, I find that typicality is negatively associated with funding before labels are present, but a curvilinear, u-shaped relationship is manifest once labels are adopted. These results, as well as the results from Chapter 2, persist as typicality and coherence are included in the same model. These results bring into question the long-standing notion in market

categories research that typicality is necessarily favorable.

INTRODUCTION

Recent extensions of market categories research has shifted attention from category spanning to candidates' within-category variation, such as partial category membership or typicality. Some of this interest has grown out of efforts to moderate the (negative) consequences of category spanning and to illustrate obstacles to 'perfect information' in markets. Kovcs and Hannan (2010) find that when a category is characterized by high variance in the grade of membership, having such a category among those multiply-spanned is less hazardous than spanning categories that are more 'crisp'. Bowers (2015), highlighting the importance of consideration sets in the interpretation of market information, shows that audiences rely on localized knowledge more when candidates are typical. Whether considered an attribute of a category or of a candidate, within-category heterogeneity has often been studied as a moderator of other information. This chapter seeks to understand the *direct* role of within-category membership information in evaluation.

THEORETICAL DEVELOPMENT

In considering the manner in which (within-) category information is assessed, it is useful to distinguish among *categories*, *labels* and *features*. While more formal definitions have been forwarded (see Hannan, Polos, and Carroll 2007), a category is a collection of objects identifiable with a common concept. Both labels and features can work to identify candidates with categories, though they differ in how and what kind of information they convey. The formality of labels (e.g., the process by which they are assigned) may vary between markets, but labels generally provide swift mappings from a signal (the label) to a conceptual category, and often do so independently. Perhaps because of this tight association, researcher usage sometimes implies categories and labels are synonyms.

Features are any source of information that are of a finer-grain than labels. Examples of features include attributes (the technical specifications of computer; a firm's balance sheet); actions (legal activity pursued by a law firm; dishes served at a restaurant), or words (the content of job talk; the words of a romance novel). Features may not have been affixed to candidates primarily or intentionally to convey category membership information, but audiences can utilize features for this purpose. For instance, the journals a scholar chooses to publish in wittingly (or unwittingly) characterizes their fit within and across disciplinary categories. Not all features are equally informative, and often must be considered in combination for category information to be assessed.

The distinction between categories, labels, and features has received more attention recently (Pontikes and Hannan 2014), though the focus has been largely reserved to category dynamics. This distinction can aid other, fundamental questions of categories research. Below, I outline two contributions the label-feature distinction offers market categories research: the different roles that labels and features play in categorization and differentiation, and the contingent nature of typicality.

The first contribution is illustrated in prior efforts to assess candidate typicality. Often candidates have been described as 'atypical' if they are affixed with multiple category labels. While it is certainly possible for a combination of labels to be more or less familiar to audiences, I argue that researchers tend to conflate two dimensions of candidate information: between-category membership, and within-category membership. When labels are considered in isolation, two candidates that are identically labeled are categorized the same: no within-category variation is evident. While a singly-labeled candidate may be inferred to be more typical of a given category than one that is also affixed with other labels, I contend that typicality is more often (and, more richly) inferred at the level of features. While acknowledging that the division between labels and features may not always be clean, I assert that between-category information is more readily inferred from labels than from the level of features, and within-category information from the level of features. As I demonstrated in Chapter 2, features can inform between-category judgments, but if labels are present, these labels generally supersede features.

The second contribution of a category-label-feature distinction is clarification of the construct of typicality. As what is typical of one category may not be of another, typicality needs to be considered in terms of a category. For purposes of analysis, it may be necessary to aggregate a candidate's typicality in each of its associated categories (as in Kovacs and Johnson, 2014, and below); but *conceptually*, typicality is a category-specific construct, a statement of how a candidate's features correspond to those of candidates currently or previously associated with that category. This conceptual clarification has more implications than for measurement alone: it raises the question of what typicality means in 'nascent' market settings in which categorization is unaided by labels. In such settings, audiences' categorization proceeds more independently than when labels are present. With labels, there is (at least the possibility of) a common reference for audiences regarding the categories to distinguish candidates between and within. Can typicality be assessed in the absence of labels, and if so, how is it evaluated? How do labels alter the evaluation of typicality?

HYPOTHESIS DEVELOPMENT

When individuals or firms compete for resources, they face a tradeoff in the manner in which they do or do not emulate others. On one hand, appearing like others can make it easier for resource-providers to recognize a candidate's offering and assign value. On the other hand, conformity to others' self-descriptions can invite competition.

Cognitive psychologists have identified two components of typicality: *representativeness*, or the extent that a member shares features of others in the category, and *certainty of membership*, which is more about the extent to which the candidate does or does not share features with candidates outside the category (Murphy and Ross 2005; Hannan, Pólos, and Carroll 2007). While both considerations ultimately figure into audiences' assessments of candidates, I focus on the representativeness component, as this more closely corresponds to a within-category comparison than an across-category comparison.[9]

Before developing specific hypotheses for labeled and unlabeled markets, I present arguments for why typicality in general should elicit positive evaluations.

Previous categories research suggests the general typicality-performance relationship is positive. Most of this follows from conceptual arguments about how audiences compare candidates to latent schemata, or the socially constructed expectations of what candidates appear in a category. Essentially, this argument suggests that feature information that does not disconfirm category membership facilitates ready evaluation. Typical candidates are recognizable.

Kovacs and Johnson (2014) provide one of the first empirical studies that directly examines relationship between typicality and evaluation. Retrieving ingredients from menus of San Francisco restaurants, these authors compute a typicality score for each ingredient and cuisine category pair. Some ingredients are more typical of some cuisines than others: for example, tomatoes are more typical of Italian cuisine than of Chinese cuisine. They find a positive effect of typicality on restaurant ratings, and further find that using typical ingredients is most helpful for restaurants that are of low and mid-quality. This finding suggests a cognate mechanism for the positive typicality-performance relationship: typicality is utilized to alleviate uncertainty concerning a candidate's quality. The two-stage model of categorization and differentiation described in the categorical imperative (Zuckerman 1999) reserves considerations of quality for the differentiation stage. That is, given that a candidate

[9] As I show below, I am able to separately measure feature coherence, which captures between-category membership, and typicality. Regressions that include coherence as a covariate alongside typicality recover effects that better correspond to the representativeness component of typicality.

is identifiable with a given category, the certainty of their quality of performance within that category is inferred from how well their features are representative of the category.

There are other reasons typicality could elicit positive evaluations. Evaluator's judgments are often not private, and in many cases evaluators are motivated to make the judgement that will be reached by others. For instance, in the stylized example of a beauty contest, a judge generally seeks to select the contestant that the maximum number of other observers will agree wins the contest. This task of rank ordering can be extended to the more general case of categories. For instance, returning to restaurants, when a large group of visitors asks an individual to recommend a local Mexican restaurant, the recommender will be more likely to suggest one with representative features of the 'Mexican restaurant' category, as this will be more likely to meet the expectations of the group.

At the same time, there are reasons typical candidates could perform worse than atypical candidates. Resource partitioning theory provides insight here. Under this theory, generalist actors seek to appeal to heterogeneous audiences while specialist actors pursue smaller, homogeneous segments (Carroll and Swaminathan 2000). In order to appeal to as broad a base as possible, generalists must adopt the most generic, typical position they can. Such typicality, while recognizable, also hazards substitutability or even denigration as 'bland' (e.g., the case of generalist beer manufacturers being labeled 'bland', 'industrial beer' (Carroll and Swaminathan 2000)). That is, high competition could negate the positive effects of being recognizable.

Typicality with and without Labels

The markets in which positive typicality effects have been theorized and found have generally been systems with labels present. Under these conditions, between-category membership is more readily inferred. This means that audiences can deploy more cognitive effort to assessing features' within-category implications, rather than both within- and between-category implications. Additionally, audiences are able to better coordinate their expectations of what categories to consider when adjudicating candidate features: those suggested by the labels. Accordingly, audiences will be better able to track what features have been associated with a given (labeled) category.

Taken together, this suggests that audiences are best able to discern typicality when labels are present. Another matter is how this typicality is then evaluated. Here, the nature of the evaluation task at hand is important—particularly, whether evaluators seek to render a judgment that accords with or departs from others.

In the Prosper marketplace, audiences seek to make judgments that others corroborate. The audience members (lenders) incur an opportunity cost if they to bid on listings that fail to obtain funding. Funds bid on listings that do not become loans are returned, but this is capital that could have been invested elsewhere and has been detained without a benefit. To reinvest this capital incurs additional search costs. As a result, lenders are incentivized to place bids on listings that will eventually become loans. Candidates with features that are typical of a labeled category are more recognized as representative of that category, and audiences can proceed with greater assurance that others will recognize this correspondence as well. Such candidates better approach the status of a commodity. Typical candidates are not differentiated by their features, but reap the benefit of anticipated recognizability. Together, this suggests the following:

Hypothesis 1: On average, typical candidates receive more funding than atypical candidates when labels are present.

The foregoing hypothesis corresponds to what market categories scholars have described and found for 'markets as usual', or those with formal classification systems. Can typicality be assessed in markets without labels? As I show in Chapter 2, categorization can and does happen even if labels are not present: in the unlabeled period of the Prosper marketplace, candidates with coherent features were more likely to be funded than those that were incoherent. Would this suggest that audiences can also conduct within-category judgments unaided by labels?

The categorical imperative (Zuckerman 1999) suggests that cognitive resources will be deployed primarily in the task of making between-category membership judgments. The cognitive cost of categorization in the absence of labels impedes within-category comparisons. Audiences may be able to discern focus and complexity, but consistent of mapping of features to categories is a more involved activity. Even if cognitive resources weren't prohibitively taxed, without labels to coordinate audiences' category bases, it would be difficult for audiences to share a reference point. Thus, while an all-seeing observer could perform the classification task I undertake with machine learning, and subsequently assess the typicality of candidates, these typicality judgments would not be shared by candidates with different consideration sets and different lay theories of the categorical structure of the market. In all, this suggests:

Hypothesis 2: On average, typical candidates receive no more funding than atypical candidates when labels are absent.

Together, Hypothesis 1 and 2 express a contingency view of typicality: labels must be present for typicality to be positively valued.

DATA AND METHODS

I utilize the Prosper marketplace data described in Chapter 2. The natural experiment of label introduction is most conducive to the study of the interplay of labels and feature typicality.

As described above, I posit that typicality is not systematically discernible in the absence of labels. Accordingly, what I measure for the pre-label period could be described as *counterfactual* typicality: the extent to which a candidate's features are representative of their predicted category, had they been explicitly labeled a member of that category at the time of appearing in the market. Below, I describe how I computed typicality and my estimation strategy.

Measuring Feature Typicality

My measure of typicality follows from Kovacs and Johnson (2014). In their study of restaurants listed in multiple food categories, these authors measure typicality by computing a Jaccard similarity weight for each menu item in each food category, then aggregating these typicality weights to give each restaurant an overall typicality score. I use a modified version of their approach that is more appropriate for a study of word features.

Let $freq_{ij}$ be the frequency of term i in category j, $freq_{i.}$ the total frequency of term i in all categories, and $freq_{.j}$ the total frequency of all terms in category j. Then the Jaccard similarity of the term i to category j is computed as follows:

$$Jaccard_{ij} = \frac{freq_{ij}}{freq_{i.} + freq_{.j} - freq_{ij}}$$

The more selectively a term occurs in one category, the higher the Jaccard similarity weight. My approach extends this method in a number of ways. First, recognizing that actors can vary in the number of features they include, I incorporate between-listing differences in description length by normalizing each listing's term frequency counts by the listing's total term occurrences. Next, to lessen the impact of terms that occur frequently in all documents, I used inverse document frequency (tf-idf) weighting instead of raw counts. The normalized term frequency (tf) weight is multiplied by the following:

$$idf_i = log_2 \frac{|D|}{|\{d|t_i \in d\}|}$$

where the numerator $|D|$ is the total number of documents and the denominator $|\{d|t_i \in d\}|$ is the number of documents containing the term t_i (Salton and Buckley 1988).

Lastly, my approach differs from Kovacs and Johnson's (2014) in the aggregation of computed term weights. Whereas these authors compute category-term weights only for categories for which restaurants explicitly have a label, and sum all of these weights to provide an overall typicality score. For example, a restaurant may use garlic, an ingredient that could suggest membership in various food categories, but category-specific weights are computed only for the categories with which a restaurant is explicitly labeled. Instead, I am initially agnostic as to the label listings are (predicted to be) identified with and compute all possible category-term weights for each term. I then calculate category totals and multiply these by the listing's predicted probability of being in each purpose category. This weighted average approach incorporates the uncertainty in the accuracy of purpose category prediction. Ultimately, this measure is highly correlated with the unweighted average (r = 0.978) and both produce qualitatively similar results, albeit the results using the unweighted average is noisier. The results reported below utilize the weighted average.

Figure 7 reports daily-average levels of 100-quantiles of typicality over time. Considering all categories together, there is not a discernible uptick in typicality after labels are present, providing confidence that estimation will not need to account for simultaneous effects of labeling on typicality and on the relationship between typicality and percent funding. Taking predicted categories separately, it appears that there is increased separation among the categories after labels are introduced. Also, debt consolidation listings change from exhibiting low typicality on average to high typicality on average after labels are introduced, while personal/other listings display the opposite trend. In supplemental analysis, I decompose regressions into predicted categories to examine the possibility of heterogeneous treatment effects. Ultimately, this change in the relative typicality of categories does not appear to drive results reported below.

Figure 7: Dynamic Plot of Typicality by Predicted Category

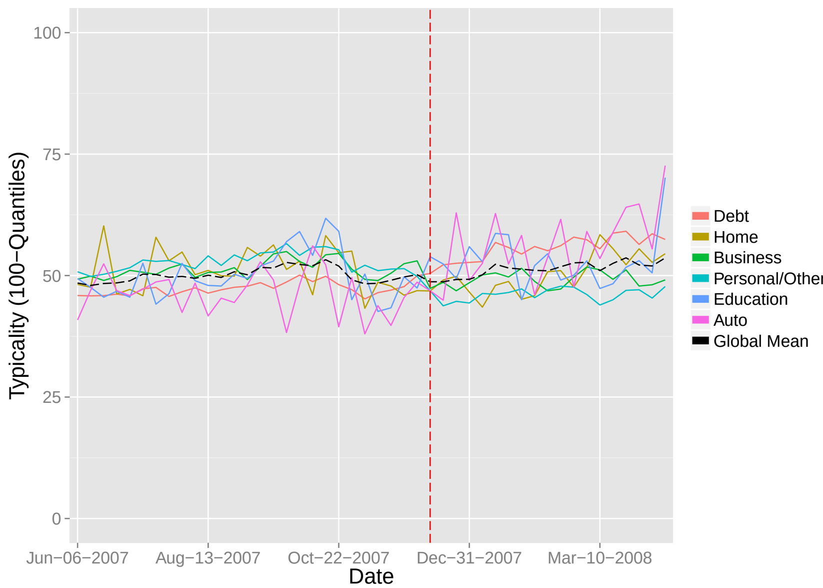

Note: Dynamic plot of 100-quantiles of typicality, by purpose category. Dates range from June 6, 2007 to April 14, 2008. The dashed red line in each figure represents when purpose labels were introduced (December 5, 2007). The dashed black line in (b) describes the global mean level of 100-quantiles of typicality.

For each purpose category, I identified the ten features yielding the highest typicality weights. Table 8 reports the top features. Recall that in Chapter 2 I produced a table of the highest model weights in the maximum entropy classification task. This table is presented here again for reference (see Table 9). Comparison of these tables is informative. As before, features may come from either the listing title or from the body of the description. Interestingly, the words reported in these two tables are similar, although the location of the words are different: features with the highest weights for classification tend to be drawn from titles, and the features with the highest typicality weights are from descriptions.[10]

Investigation of listings at different levels of the typicality measure reveals that highly typical listings—per the weighted measure described above—tend to use at

[10]In other analyses, I used the model weights from maximum entropy classification in a similar manner as the Jaccard typicality weights, but the descriptions manifesting the extremes of this alternate measure of typicality failed to provide similar face validity as the approach used below.

Table 8: Top 10 Features with Highest Typicality Weights

	Debt. Consolid.	Home Improv.	Business	Personal/Other	Education	Auto
1	cards[1]	home	business[1]	explain[1]	college	car
2	credit	home[1]	business	back[1]	school[1]	car[1]
3	debt[1]	improvement	explain[1]	paying[1]	college[1]	buying
4	credit[1]	house[1]	company[1]	bills[1]	tuition[1]	vehicle[1]
5	cards	roof[1]	equipment[1]	candidate[1]	school	purchase[1]
6	interest[1]	kitchen[1]	expand[1]	good[1]	degree[1]	buy[1]
7	paying	bathroom[1]	purchase[1]	bills	education[1]	reliable[1]
8	debt	roof	inventory[1]	pay[1]	semester[1]	auto
9	card[1]	kitchen	income[1]	job[1]	student[1]	work[1]
10	pay[1]	improvements	years[1]	time[1]	graduate[1]	work

Note: Features with the highest Jaccard typicality weights for each purpose category. The superscript ([1]) indicates the feature is from Purpose or Financial Situation section. Otherwise, the feature is the from the listing title.

Table 9: Top 10 Features for Label Classification

	Debt. Consolid.	Home Improv.	Business	Personal/Other	Education	Auto
1	debt	roof	inventory	vacation	college	car
2	debts	kitchen	capital	wedding	student	auto
3	free	garage	advertising	emergency	paying	buying
4	loans	addition	business	bills	school	buying car
5	credit	putting	equipment	personal	degree	vehicle[1]
6	cards	drive	payroll	surgery	tuition[1]	motorcycle
7	payoff	home	business[1]	furniture	education[1]	vehicle
8	chance	pool[1]	shop	boat	education	transportation
9	bills	room	design	repair	school[1]	auto loan[1]
10	consolidate[1]	adding	operating	daughter	classes[1]	car[1]

Note: Features with the highest weights for each purpose category, estimated through maximum entropy. The training data are all labeled listings created between December 8, 2007 and June 27, 2008. The superscript ([1]) indicates the feature is from Purpose or Financial Situation section. Otherwise, the feature is the from the listing title.

least some of the top ten words in the category. Also, a highly typical listing tends to use the same words in titles and descriptions. In doing so, highly typical listings forgo providing differentiating information. Table 10 provides examples of varying levels of typicality, both before and after labels are adopted. ‘No’ in the first column indicates that the listing appeared in the unlabeled period of the market; ‘Yes’ indicates the listing appeared in the labeled period.

Table 10: Sample Listings from Same Matching Stratum

Labeled	Typicality: 100-Quantile	Coherence: 100-Quantile	Description
No	21st	1st	Title: Duckbill Purpose of loan: For legal expenses My financial situation: Consultant with steady income
No	25th	35th	Investing Money in a Swim Club Purpose of loan: With the money from this loan I will invest it in a youth project to expand area Swim Team, for training, travel and other expenses.
No	87th	29th	starting my own business Purpose of loan: I am in the process of trying to start my own business. It's a residential/ commercial cleaning business. My financial situation: I am a hard worker and would just like someone to take a chance on me
No	96th	17th	support my business Purpose of loan:(explain what you will be using this loan for) My financial situation:(explain why you are a good candidate for paying back this loan)
Yes	3rd	6th	start chatrooms for prison inmates and excons Purpose of loanto start chatrooms for prison inmates and excons who want to chat with each other My financial situation: I am a good candidate for this loan because I am a conservative country man who is finacialy resposible.
Yes	8th	9th	Buying materials for California Highway Patrol Project Purpose of loan: This loan will be used to Materials for Project My financial situation: I am a good candidate for this loan because This project is for the state of California
Yes	82nd	27th	Buying inventory and equipment Purpose of loan: This loan will be used to buy stock My financial situation: I am a good candidate for this loan because I always pay my bills
Yes	94th	85th	Business Capital Purpose of loan: This loan will be used to business marketing/ capital for payroll My financial situation: I am a good candidate for this loan because my business is growing daily/ i believe in priority first which means my bills come before anything else

Example listings are from the same matching stratum: requested between $2,000 and $3,000, appeared in the 'high risk' (HR) credit grade, would remain open for the full stated duration, used between 1 and 46 words in title and description, and were predicted to be in the Business category. Values of 100-Quantiles of coherence and typicality are reported. 'No' in the 'Labeled' column indicates the listing appeared in the unlabeled period of the market; 'Yes' indicates the listing appeared in the labeled period of the market.

Note that original spelling and punctuation is preserved. Examples were drawn from the same matching stratum: listings requesting similar amounts, at the same level of risk, used a similar number of words, and predicted to be in the same purpose category (Business). A listing manifesting high typicality uses words that are to be expected from listings predicted to be in its respective category. Examples reported in Table 10 also report the level of coherence described in Chapter 2. These examples demonstrate that there is some correspondence between typicality and coherence: for instance, in the second example, a proposal to start a swim club is perhaps not expected of a business loan, and the coherence score correspondingly indicates this listing is predicted to be in multiple categories. However, the examples also show it is possible to be highly typical and yet less coherent. The analyses that follow will account for coherence and typicality in the same regressions.

Empirical Strategy

I use a methodological approach similar to what I described in chapter 2: I model a listing's percent funded as a function of its attributes and surrounding market conditions. As demonstrated in that chapter, it is important to account for potential confounds through exact and coarsened-exact matching. The regression models presented below use the same matched sample as described in Chapter 2. The matching strategy paired labeled and unlabeled listings based on coarsened bins of amount requested, borrower maximum rate, and number of words, as well as on indicators of predicted purpose category and loan funding option (open for the duration of the loan, or closed when funded). The descriptive statistics for this matched sample are reported in Table 11.

RESULTS

Table 12 reports results of ordinary least squares regression. Percent funded is regressed on listing attributes, market conditions, and discrete variables of interest.

As was the case for feature coherence in Chapter 2, there is not a natural unit of measurement for feature typicality. Accordingly, I construct indicators that the listing demonstrates the upper third or lower third of the distributions of typicality. Model 1 recovers a negative, statistically significant effect of high typicality, and a positive and significant effect for low typicality. To this point, the pattern of results is consistent with a *negative* main effect of typicality across all periods of the market. Model 2 introduces an indicator that the listing appeared in the labeled period of the market, and Model 3 includes interactions between typicality indicators and the label indicator. Neither main effect for high or low typicality is notably attenuated in either of these models. There is a positive and significant interaction between high typicality and the label indicator; this suggests that the introduction of labels has a different impact on listings with high typicality than on listings with low typicality.

Table 11: Descriptive Statistics: Matched Listings

	Pre-Label Listings (N = 49,026)			Labeled Listings (N = 49,026)			All	
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.	Min.	Max.
Percent Funded	0.17	0.01	0.33	0.15	0.01	0.32	0.00	1.00
Nb. Words	187.38	147.00	128.13	179.40	142.00	123.60	1.00	770.00
Nb. Words (excluding budget)	147.07	114.00	113.41	147.30	113.00	114.12	1.00	770.00
Amount Requested	7324.35	5000.00	6494.10	7314.06	5000.00	6386.71	1000.00	25000.00
Max. Borrower Rate	0.18	0.17	0.07	0.19	0.18	0.09	0.00	0.36
Debt to Income Ratio[1]	0.54	0.26	1.32	0.43	0.26	0.98	0.00	10.01
Credit Grade: AA (760 and up)	0.04	0.00	0.19	0.04	0.00	0.19	0.00	1.00
Credit Grade: A (720-759)	0.05	0.00	0.22	0.05	0.00	0.22	0.00	1.00
Credit Grade: B (680-719)	0.08	0.00	0.27	0.08	0.00	0.27	0.00	1.00
Credit Grade: C (640-679)	0.14	0.00	0.34	0.14	0.00	0.34	0.00	1.00
Credit Grade: D (600-639)	0.18	0.00	0.38	0.18	0.00	0.38	0.00	1.00
Credit Grade: E (560-599)	0.18	0.00	0.38	0.18	0.00	0.38	0.00	1.00
Credit Grade: HR (520 to 559)	0.34	0.00	0.47	0.34	0.00	0.47	0.00	1.00
Duration	7.67	7.00	1.99	7.51	7.00	2.27	3.00	10.00
Homeowner	0.36	0.00	0.48	0.39	0.00	0.49	0.00	1.00
In Prosper Group	0.24	0.00	0.43	0.13	0.00	0.33	0.00	1.00
Automatic Funding	0.09	0.00	0.29	0.09	0.00	0.29	0.00	1.00
Nb. Attempt by Member	2.89	2.00	3.07	3.34	2.00	4.62	1.00	68.00
Number Live by Start	2621.32	2539.00	340.44	2372.96	2295.00	442.12	1684.00	3971.00
Percent Live in Same Category	0.36	0.45	0.16	0.32	0.27	0.16	0.01	0.52
Predicted Label: Debt	0.47	0.00	0.50	0.47	0.00	0.50	0.00	1.00
Predicted Label: Home	0.02	0.00	0.13	0.02	0.00	0.13	0.00	1.00
Predicted Label: Business	0.15	0.00	0.35	0.15	0.00	0.35	0.00	1.00
Predicted Label: Personal/Other	0.34	0.00	0.47	0.34	0.00	0.47	0.00	1.00
Predicted Label: Education	0.02	0.00	0.13	0.02	0.00	0.13	0.00	1.00
Predicted Label: Auto	0.01	0.00	0.10	0.01	0.00	0.10	0.00	1.00
Coherence	0.74	0.78	0.22	0.78	0.85	0.21	0.17	1.00
Second-Highest Probability	0.14	0.10	0.14	0.13	0.07	0.14	0.00	0.50
Typicality[2]	51.85	53.00	29.14	51.00	51.00	29.15	1.00	100.00

N Listings = 98,052. Note: Descriptive statistcs for listings that appeared on Prosper's website from June 6, 2007 through April 14, 2008. Labeled and unlabeled listings were matched exactly on Prosper credit grade, predicted category, and funding option, and coarsened-exact matched on amount requested and word count.

[1] Missing values for 6675 observations.

[2] Missing values for 7 observations. Typicality is the subject of Chapter 3.

Visualization of this effect through plots of regression smooths will illustrate this effect more directly.

Model 4 includes indicators of predicted purpose label and interactions with the label indicator, and Model 5 adds the feature coherence indicators and interactions that were the primary interest in Chapter 2. The pattern of results for typicality are robust to the inclusion of these additional variables. Also, the models recover similar results for coherence as was found in Chapter 2.

Table 12: OLS Regression: Percent Funded - Matched Listings

	Model 1	Model 2	Model 3	Model 4	Model 5
Intercept	−0.040***	0.007	0.013	0.008	0.008
	(0.007)	(0.008)	(0.008)	(0.009)	(0.009)
Amount Requested	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Borrower Max. Rate	0.937***	0.960***	0.960***	0.960***	0.959***
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
DIR	−0.016***	−0.017***	−0.017***	−0.017***	−0.017***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Credit: AA	0.711***	0.712***	0.712***	0.712***	0.711***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Credit: A	0.552***	0.553***	0.553***	0.553***	0.552***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Credit: B	0.445***	0.447***	0.446***	0.447***	0.446***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Credit: C	0.268***	0.269***	0.269***	0.269***	0.268***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Credit: D	0.141***	0.142***	0.141***	0.142***	0.141***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Credit: E	0.049***	0.049***	0.049***	0.049***	0.049***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Duration (Days): 3	0.001	0.006	0.005	0.005	0.005
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Duration (Days): 5	0.004	0.007*	0.007	0.006	0.006
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Duration (Days): 10	0.016***	0.018***	0.018***	0.018***	0.018***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Homeowner	−0.001	0.000	0.001	0.000	0.000
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
In Prosper Group	0.068***	0.060***	0.059***	0.059***	0.058***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Open Funding	−0.022***	−0.022***	−0.022***	−0.022***	−0.022***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Nb. of Attempt	−0.004***	−0.004***	−0.004***	−0.004***	−0.004***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Nb. Words	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Nb. Listings Live	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Perc. Same Category	0.001	−0.011	−0.013*	0.002	0.002
	(0.006)	(0.006)	(0.006)	(0.008)	(0.008)
Upper 33% Typicality	−0.011***	−0.011***	−0.026***	−0.026***	−0.027***
	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
Lower 33% Typicality	0.008***	0.008**	0.009**	0.010**	0.012***
	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
Labeled		−0.037***	−0.047***	−0.046***	−0.044***
		(0.002)	(0.003)	(0.004)	(0.004)
Upper 33% Typicality × Label			0.030***	0.030***	0.028***
			(0.004)	(0.004)	(0.004)
Lower 33% Typicality × Label			−0.003	−0.003	−0.004
			(0.004)	(0.005)	(0.005)
Purpose: Home				0.064***	0.068***
				(0.010)	(0.010)
Purpose: Business				0.008	0.005

Continued on next page

Table 12 – *Continued from previous page*

	Model 1	Model 2	Model 3	Model 4	Model 5
				(0.005)	(0.005)
Purpose: Personal/Other				−0.004	−0.001
				(0.003)	(0.003)
Purpose: Education				0.002	0.005
				(0.010)	(0.010)
Purpose: Auto				0.027*	0.032*
				(0.013)	(0.013)
Prp: Home × Label				−0.065***	−0.069***
				(0.014)	(0.014)
Prp: Business × Label				−0.014*	−0.012*
				(0.006)	(0.006)
Prp: Personal/Other × Label				0.007	0.005
				(0.004)	(0.004)
Prp: Education × Label				0.028*	0.026
				(0.014)	(0.014)
Prp: Auto × Label				0.000	−0.004
				(0.018)	(0.018)
Upper 33% Coherence					0.016***
					(0.003)
Lower 33% Coherence					−0.009**
					(0.003)
Upper 33% Coherence × Label					−0.013**
					(0.004)
Lower 33% Coherence × Label					0.009
					(0.004)
Log Likelihood	-11169.423	-10994.281	-10958.700	-10924.501	-10894.813
R^2	0.300	0.303	0.303	0.304	0.304
Num. obs.	91370	91370	91370	91370	91370

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

To visualize the interactions described in the parametric model, I conducted iterative reweighted least squares regression, substituting indicators of levels of typicality with regression smooths. All covariates reported in Model 5 have been retained, with regression smooths for continuous variables substituted for raw variables. Such semiparametric modeling seeks to account for nonlinearities in covariates. Figure 8 is then the semiparametric counterpart to the parametric effects obtained in the OLS regressions. As seen in the figure, the effects of typicality on percent funded have been disaggregated by whether the listings appeared in the labeled or unlabeled period of the market. Consistent with the negative main effect suggested by earlier models, I recover a negative, linear effect for typicality in the pre-label period. Subsequently, however, the effect for typicality is strikingly different: listings in the highest end of the distribution of typicality receive significantly more funding than listings in the middle of the distribution, and the positive effect for atypical listings is generally attenuated. In all, the effect of typicality on percent funded reflects a nonlinear, u-shaped curve: those at either extreme of typicality perform better than those that are moderately typical.

Figure 8: Smooth Estimates of the Effect of Typicality on Percent Funded

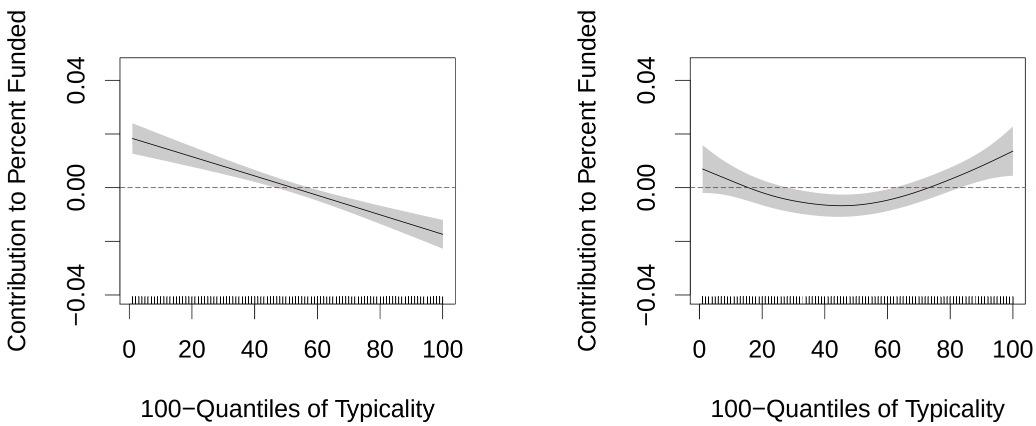

(a) Unlabeled (b) Labeled

Note: The solid black lines correspond to smooth term estimates stemming from iterative reweighted least squares regression.

The negative effect for typicality suggests that there is discernment of within-category membership even if labels are absent, and that typicality in such circumstances is *discounted* relative to typicality under labels. Inspection of many listing descriptions manifesting the extremes of typicality both before and after labels are adopted failed to find systematic changes in the content of descriptions. Recall that care was taken to ensure that typicality is comparable across market periods: the 100-quantiles of typicality were calculated using the entire analysis sample, meaning, for instance, that an unlabeled listing in the 23rd percentile of typicality would register in the 23rd percentile in the labeled period as well.

Returning to the listing descriptions, it is clear that to take a typical position in a purpose category is to forgo differentiation: such listings seldom communicate details about themselves or their intended use of the loan. This differentiating information increases the likelihood of being noticed in general. To the extent audiences seek to select the same candidates as others, a natural inference is that highly differentiated candidates are more likely to be selected than those that are not differentiated, or that have taken a generic position. However, in a labeled world, such conformity is valuable since it invites confidence that others will be able to apply category-specific schemata successfully.

It is still surprising that there does not seem to be a diminishing effect of atypicality. It would seem that at the extreme, atypical positions may be noticed but not funded. Replication of the results without controlling for coherence shows some evidence in support of the diminishing returns expectation.

To further examine mechanisms at work, I disaggregate the effect of typicality by estimating separate regressions for bins of Prosper credit grade: high credit grade (AA, A and B); medium credit grade (C, D, and E); and low credit grade (HR, or high risk). Recall that these groups represent 17%, 49%, and 34% of the matched sample, respectively. Figure 9 reports the regression smooths for typicality estimated in each of these regressions. In the pre-label period, the effect of typicality on percent funded is consistently negative, although slopes are declining in credit risk. Inspection of descriptive statistics for the high, medium and low credit-grade groups suggest that variance in percent funding declines with credit grade. Thus, the pattern reflected in the unlabeled period in the respective credit groups is less about creditworthiness positively typicality across credit grades, but differences in the total available variance to be explained. Thus, the (negative) contribution of typicality to percent funded in the pre-label period appears to be proportional to the total variance in funding in credit-grade groups.

The disaggregation reveals a different story for the labeled period. For the high credit-grade group, the presence of labels has rendered the negative, linear effect of typicality a positive, nearly linear effect. It is among medium credit-grade group that the curvilinear, u-shaped effect of typicality is most apparent. Lastly, among low-credit listings, the slight, negative effect of typicality is entirely attenuated. The latter pattern is not terribly surprising, given what is known about the funding rate of high-risk credit listings, but the change in effects for high and medium credit groups is striking.

Figure 9: Smooth Estimates of the Effect of Typicality by Credit Grade

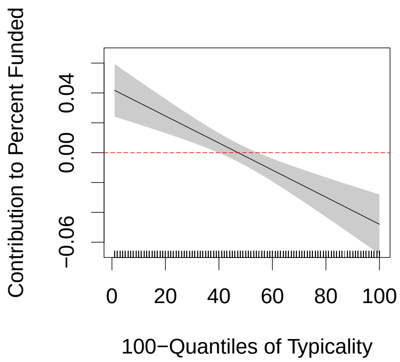

(a) Unlabeled: AA, A, B

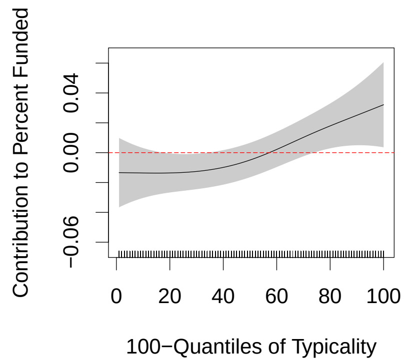

(b) Labeled: AA, A, B

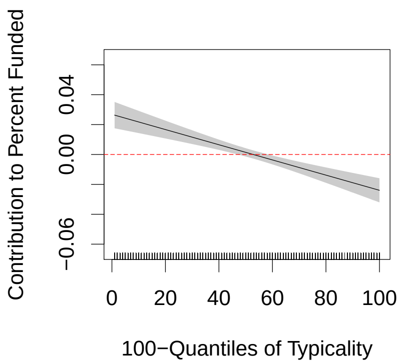

(c) Unlabeled: C, D, E

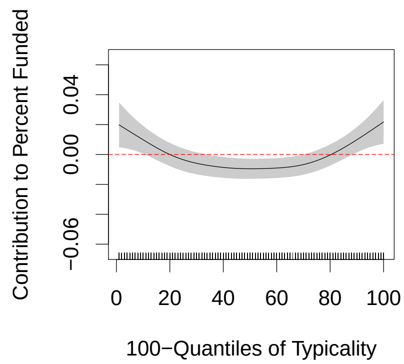

(d) Labeled: C, D, E

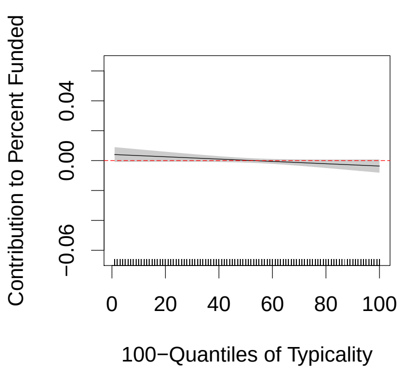

(e) Unlabeled: HR

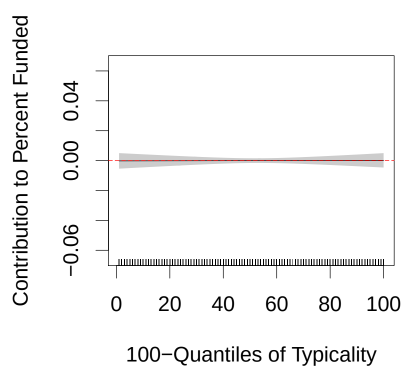

(f) Labeled: HR

It is notable that the pattern displayed for high credit grades reflects the expectation of the literature for labeled markets (Hypothesis 1). The differences between high and middle credit grade groups may reflect differences in audience-side diffusion of market conventions. As listings with the highest credit grades undergo the most volume of funding, it may that the value of typicality—reliability in valuation—is most readily recognized for these listings. If this is so, it means that Prosper lenders' assessments of typicality are credit-grade dependent once labels are introduced, because otherwise, the pattern for the middle credit grades would be linear and positive as well. The middle credit grade may exhibit a transitional state: high typicality has become valued, as it has for high-credit grade listings, and the most atypical are still defying the move to commodification.

Other Analysis

In all models, I control for the degree of crowding in the focal listing's category (percentage live in category). This is to account for an important time-variant difference between purpose categories. Crowding has itself been the focus of other research. For example, Barroso et al. (2014) assess the effect of crowding within television series themes on series longevity, finding that crowding is actually positively associated with series longevity. In my study, the effect of crowding on funding is rather noisy in the pre-label period, and then linear (positive) in the post-label period. This generally replicates the results of Barroso and colleagues, though the effect is not as pronounced as for typicality. This is likely due to the relative coarseness of the categories I considered: rather than rely on labels (e.g., TV genres), Barroso and colleagues use unsupervised machine learning to identify distinct conceptual themes in television series.[11] Figure 10 reports regression smooths for crowding obtained through semiparametric modeling.

[11]Preliminary application of unsupervised machine learning methods to Prosper descriptions produced yielded clusters with poor face validity. Future research may return to this task to extend Barroso and colleagues' work more directly.

Figure 10: Smooth Estimates of the Effect of Crowding on Percent Funded

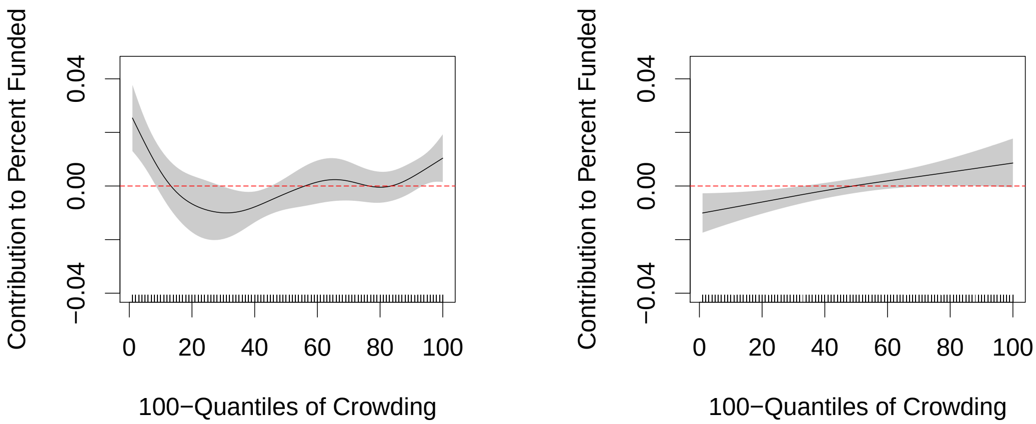

(a) Unlabeled (b) Labeled

Note: The solid black lines correspond to smooth term estimates stemming from iterative reweighted least squares regression.

Figure 11: Smooth Estimates of the Effect of Coherence on Percent Funded

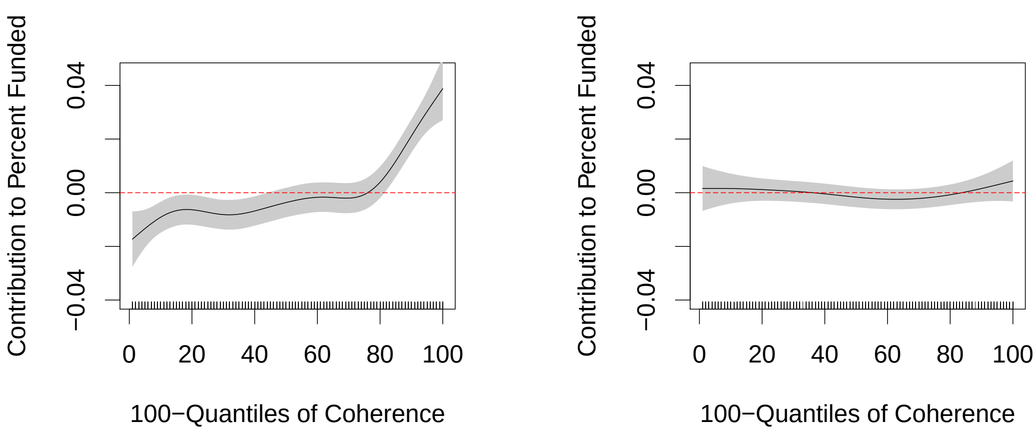

(a) Unlabeled (b) Labeled

Note: The solid black lines correspond to smooth term estimates stemming from iterative reweighted least squares regression.

Figure 11 reports results of semiparametric modeling for coherence. This is a robustness check. The same pattern of results as demonstrated in chapter 2 is found here as well. The separate effects of crowding, typicality, and coherence on percent funded are robust to the inclusion of the other variables in the model.

DISCUSSION AND CONCLUSION

This chapter extended the discussion of the difference between categories, features, and labels. In focusing on audiences' use of features to make within-category judgments of candidates, this chapter complements the program of Chapter 2, which focused on audiences' between-category judgments.

Taken together with Chapter 2, the results of this chapter suggest that the introduction of labels does not obviate the role of features entirely. Rather, where within-category considerations are concerned, audiences pay attention to how a candidate's features conform to what does or has previously appeared in a category. This affirms the pairing of labels with between-category judgments and features with within-category judgments.

The nature of the relationship between typicality and evaluation depends critically on whether labels are present. In market conditions that include labels, I am able to replicate the positive relationship suggested by prior research, but only for candidates with low altercentric uncertainty (Podolny 2001). For a sizeable subset of my data, those evoking some levels of uncertainty in quality, candidates atypical of their predicted category receive higher valuations as well. And, quite contrary to the expectations of the literature, I find that typicality is *discounted* generally when labels are absent. In all, I find partial support for Hypothesis 1, and I find evidence to reject Hypothesis 2.

Perhaps the results ultimately raise more questions than provide answers. Qualitative analysis of listing descriptions at different extremes of the typicality distribution suggested that the content of descdesriptions may not have changed much, but it seems the same information may be viewed differently depending on whether labels are present. Highly typical listings more closely resemble commodities. Indeed, it appears that Prosper eventually sought to make the market less about the unique story of the borrower and more about borrowers' opportunities to gain access to capital.

In summary, the results of Chapter 2 and this chapter suggest that in the absence of labels, audiences assess between-category membership—belonging to one or multiple categories—but that fidelity to a particular category, a within-category comparison, is not valued unless labels are present. Absent of labels, candidates are in general better off being distinct from others, but the premium for distinctiveness (for some) may be attenuated in the presence of labels.

There are some limitations to the approach taken in this chapter. Essentially, the calculation of typicality either assumes (a) some omniscience on the part of audiences: that they view all possible candidates they make comparisons among all words (or actors); or (b) that observers' judgments will, in the aggregate, approach what I model with my measure. Bowers (2015) suggested similar limitations in categories research with respect to consideration sets. Were data on Prosper borrowers' precise consideration sets available, I could calculate within-borrower measures of typicality based on the candidates that they observe. In the next chapter, I do something approaching this in studying evaluations of unconventional category label combinations.

The present study was intentionally bound to a relatively short period of time so as to identify the effects of labels on typicality with minimal external market confounds, and also to keep the participants and kind of content presented relatively stable. The matter of dynamics in typicality would be a natural extension. Before, I hinted at the possibility of satiation effects in typicality. Barroso et al. (2014) studied satiation in prime-time television, though the attention in that study was on using features to define niches, wherein the members of a niche are assumed to be the same. Typicality may afford candidates the benefits of early acceptance into a category, but also hazard obsolescence as audience preferences change.

Future research into the dynamics of typicality could examine the viability of strategic re-entry into markets. Possible 'moves' could be to assume a more typical position than before, a less typical position, or to remain the same. Empirically, typicality could be assessed by collapsing all market data, as is done here, or it could be calculated dynamically based on changes in what candidates have arrived.

4 Variety Is the Spice of Life: Audience Engagement and the Preference for Unconventional Category Combinations

(Coauthored with Ming D. Leung)

ABSTRACT

Extant work in market categorization has documented the hazard of category spanning. While recent research has shifted attention from candidate-side to audience-side explanations, we argue this literature should focus more on audience motivations to explain when uncommon category combinations are devalued or ignored. We propose that in markets where resource holders are motivated to seek diversity, unconventional candidates are advantaged because by differing from others, such candidates represent unique opportunities. Furthermore, we suggest that more engaged audience members are better at recognizing this unconventionality and are therefore more likely to choose these candidates. We find support for our contentions with data from an online peer-to-peer lending market, Prosper.com. Contrary to expectations of the current categories literature, borrowers who utilize less commonly paired category labels receive more funding. Lenders who participate more in the market are more likely to loan money to those same unconventional borrowers.

INTRODUCTION

Abundant findings in economic sociology suggest that social actors who combine disparate elements from prevailing social categories are disadvantaged (Hannan, Pólos, and Carroll 2007; Hsu, Koçak, and Hannan 2009; Zuckerman 1999). Markets consist of actors in two roles: audience members, who control resources; and candidates, who desire those resources. Audiences search and evaluate candidates relative to expectations of a recognized categorical niche (Pontikes 2012; Rao, Monin, and Durand 2005; Zuckerman et al. 2003). When a given candidate defies expectations, they are generally passed over or devalued. Therefore, firms that combine several industry groups are ignored by finance analysts who are focused on a single one (Zuckerman 1999), chefs who span Nouvelle and Classical cuisines are initially punished by the Guide Michelin (Rao, Monin, and Durand 2005), wines which are produced by less focused winemakers receive lower ratings (Negro and Leung 2013), movies that combine more film genres are less well reviewed (Hsu 2006), and eBay sellers who attempt to sell across disparate product categories are less successful (Hsu, Koçak, and Hannan 2009). In short, audiences tend to select focused candidates over ones that span elements from multiple categories, or what Ruef and Patterson (2009) call "hybrids."

We question the scope of these theories because other scholars supply instances in which unconventional actors can be advantaged (Kleinbaum 2012; Leung 2014; Padgett and Ansell 1993; Pontikes 2012; Smith 2011; Zuckerman et al. 2003). We augment these findings by identifying a circumstance in which audiences may instead prefer candidates who span unconventional combinations of social categories. We leverage theories of diversification in markets and audience engagement to demonstrate that an understanding of the motivations of the audience help us to identify the conditions under which items that associate with unconventional combinations of categorical elements can be successful.

Previous studies provide grounding for our assertion of audience preference for spanning, though we depart from this work in several noteworthy ways. Ruef and Patterson (2009) demonstrated that hybrid organizations were penalized less when an audience's category understanding was not well-defined. However, their findings could not demonstrate how rare hybrids and those organizations that spanned particularly difficult boundaries may be rewarded even in the presence of an established categorical system. Smith (2011) showed that rewards for non-conformity to categorical conventions were forthcoming to those social actors who were successful despite their less well-understood position. Smith's account relies on redemptive signals of competence; we present the possibility that non-conformity itself may be valued. Pontikes (2012) theorized that multiple roles existed in the marketplace, and that those roles which had the ability to construct the market by manipulating categories and the social actors within them were more likely to prefer firms that maintained an ambiguous identity because this flexibility was attractive. Our theory does not rely upon different roles and instead proposes a single audience can be composed of heterogeneous members who have different preferences for novelty depending on their experience.

We examine an online market for peer-to-peer lending, Prosper.com. This market connects individuals who wish to borrow money with other individuals that wish to lend it. We focus on lenders' motivations in making lending decisions. In doing so, we question the notion that an audience may be confused by hybrid social actors who attempt to combine multiple social categories (Hannan, Pólos, and Carroll 2007; Zuckerman 1999). Instead, we highlight the fact that the motivations of an audience should be examined more closely to reveal whether category spanning is indeed disadvantaged because acceptance of such behavior is ultimately a function of audience preferences (Merluzzi and Phillips 2014).

In financial markets such as this one, prevailing theories of categorization would suggest that in order for an individual investor to do well they should become steeped in understanding one particular type of loan - thereby being able to adequately assess an offering's risk and return. In short, we should expect investors to become specialists and prefer loans that fit nicely into a 'categorical imperative.' This is because, similar to the finance analysts in Zuckerman's (1999) investigation, having in-depth knowledge is helpful when making comparisons. On the other hand, there is a salient belief among participants in financial markets that diversification is a successful investment strategy (Markowitz 1952). In this case, we should instead see the individual investors here preferring to invest in a variety of loans, and not hewing to loans which are clearly categorized.

We find that the lenders on Prosper diversify their holdings. We contend that one way to accomplish diversification would be to lend to individuals who are least like one another.—that is, less similar to the average investment. If this is the case, we should expect to see borrowers who belong to unconventional combinations of social categories to benefit. This is because these non-familiar combinations present opportunities for lenders to increase the differentiation among their investments.

This chapter also tackles the issue of audience heterogeneity. We point to a finding in Hsu (2006) which demonstrated that movie critics were more likely to review a film which spanned a larger number of genres, while those precise films were punished by the mass audience. Film critics and the mass audience differ in how they view and interact with the market. We contend that audiences can vary in terms of how much knowledge they have of the market (Hsu, Koçak, and Hannan 2009). We hypothesize that for audience members who are more involved with the marketplace, those that are more 'engaged', that they would be more likely to recognize unconventional borrowers. This is because more engaged lenders would have developed a better understanding for the nuances of social categories and borrowers. This understanding assists them in recognizing better opportunities to diversify.

THEORETICAL DEVELOPMENT

Categorization in Markets

Per Zuckerman (1999), a market is an interface between participants in two distinct roles: candidates seeking resources, and audiences that have these resources. Transactions occur when audience members search and find candidates who fit their requirements. Categories are socially agreed upon groupings of like-objects (Hannan, Pólos, and Carroll 2007; Hsu 2006; Hsu and Hannan 2005) that bound the limits of search and thus facilitate these transactions (Rao, Monin, and Durand 2005; Zuckerman et al. 2003; Zuckerman 1999). Industry groups, represented by SIC codes, are one example because they cluster similar firms together (Zuckerman 1999). Examples of a less formal categorization system could be genres of films, such as comedy or drama, which serve to partition different types of movies (Hsu 2006). Categories represent a bundle of characteristics, which differentiate members from non-members and act as a cognitive device to both aggregate and separate objects (Zerubavel 1997). For example, grocery items found in the "vegetable" aisle will share similarities with one another, but be different from those items found in the "baking" aisle. Categories assist in search as they reduce the cognitive effort required to understand each object and helps us narrow our choices quickly. Grouping and labeling similar candidates eases the identification and comparison process for audiences. Market transactions are thereby facilitated.

In instances where audiences need to fulfill a well-defined requirement, candidates who attempt to combine elements from disparate categories have been found to be disadvantaged (Hsu 2006). These candidates fit poorly with an audience's expectations of a category in which they are searching and are therefore ignored. Zuckerman (1999) demonstrated that finance analysts who were tasked with the responsibility of evaluating firms within a particular industry were more likely to ignore firms that combined elements from multiple, disparate industries. These were more difficult to compare and evaluate than their industry focused peers. Instead of attempting to understand and value them, financial analysts merely ignored them.

Note that audience members are assumed to be limited to searching for candidates to fulfill a particular categorical niche. Referring to a scope assumption of his findings, Zuckerman (1999) writes, ". . . the market must possess certain structural features . . . an influential class of critics who specialize by category" (1405).

The audience he describes is therefore limited to considering only those candidates belonging to a categorical niche. However, recent advances in research on categorization processes demonstrate conditions in which non-conformity is advantaged. Smith's (2011) investigation of the hedge fund industry found that non-conforming hedge funds were more likely to get investor funding (than conforming ones) following increases in short-term performance. He suggested that investors were "drawn to the new and different if and when the new and different demonstrates competence" (68). The non-conforming, or novel, hedge funds stood out from a crowd, and therefore were subsequently "excessively" rewarded when successful. Pontikes's (2012) asserted that venture capitalists see firms that are categorical misfits as most promising because "ambiguous labels are less constraining and give organizations room to develop industry-changing products." Venture capitalists have the luxury of scanning for firms across established categorical boundaries. What these studies demonstrate is that audience members do not necessarily limit themselves to considering categorically narrow firms only—they can instead either take notice of non-conforming candidates (Smith 2011) or are incentivized to prefer ambiguous ones (Pontikes 2012).

Prosper Groups and Group Category Labels

Prosper provides a fitting setting to study the issue of conformity. As previously described, there are two important and distinct roles on this website—individuals that wish to borrow money and those that wish to lend it—paralleling the candidate and audience conception, respectively. Candidates, or users of the website who wish to borrow money, can post an unsecured loan request or listing for up to $25,000 to be paid back over three years. These prospective borrowers provide a short description of the purpose of the loan and submit financial information (e.g., credit rating, income, debts) verified by a third party. In addition to extensive financial and loan-related information, a listing shows other facets of a prospective borrower's profile, including the borrower's occupation, their state of residence, and (the focus of our investigation) group affiliation.

Figure 12: Sample Prosper Listings

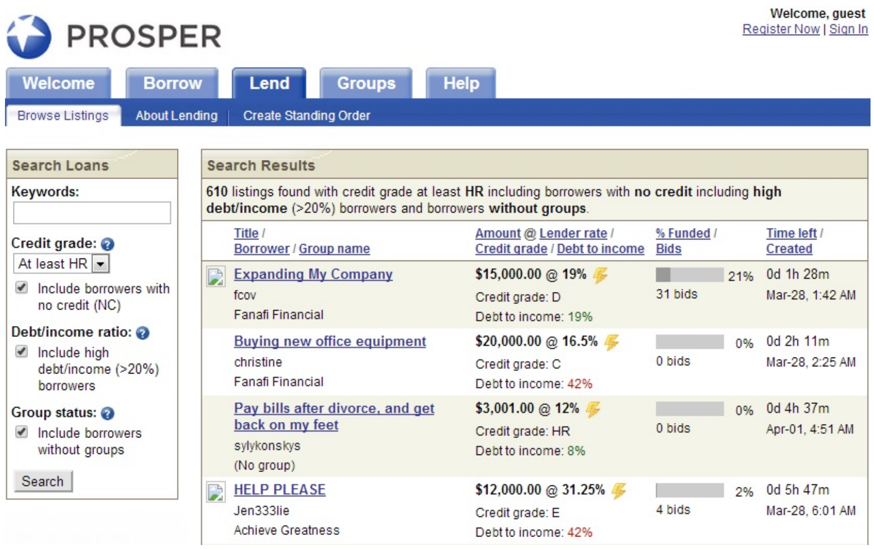

Note: A screenshot of Prosper - April 11, 2006.

Seeking to create a sense of community among borrowers, Prosper enabled participants to establish self-organized groups. In the words of the website, "Prosper Groups are a way for tightly affiliated communities to help their members through peer-to-peer lending." In Figure 12 and Figure 13 (examples that also appeared in Chapter 2), we can see examples of how Prosper Groups were visible to potential lenders: Group names were visible on browsing pages and on specific listing pages. Clicking on the group link on the group member's listing took lenders to a group page with a more detailed description of the group, along with the labels for the categories with which the group was affiliated. See Figure 14 below for an example of a group page.

Group membership is relevant to lenders' lending decision for at least three reasons. First, group membership should increase a sense of community among borrowers and therefore encourage them to be more likely to pay back their loans. Lenders should therefore be sensitive to the groups borrowers belong to. Second, groups screen their members, so group membership should suggest a minimal level of quality. Lastly, group membership communicates additional information regarding the borrower, thereby affecting how attractive they appear to lenders, who presumably wish to gather as much information as possible regarding their investments. In sum, group membership should factor into a lender's consideration as to whether to fund a loan or not.

Figure 13: Sample Listing: "Expanding My Company"

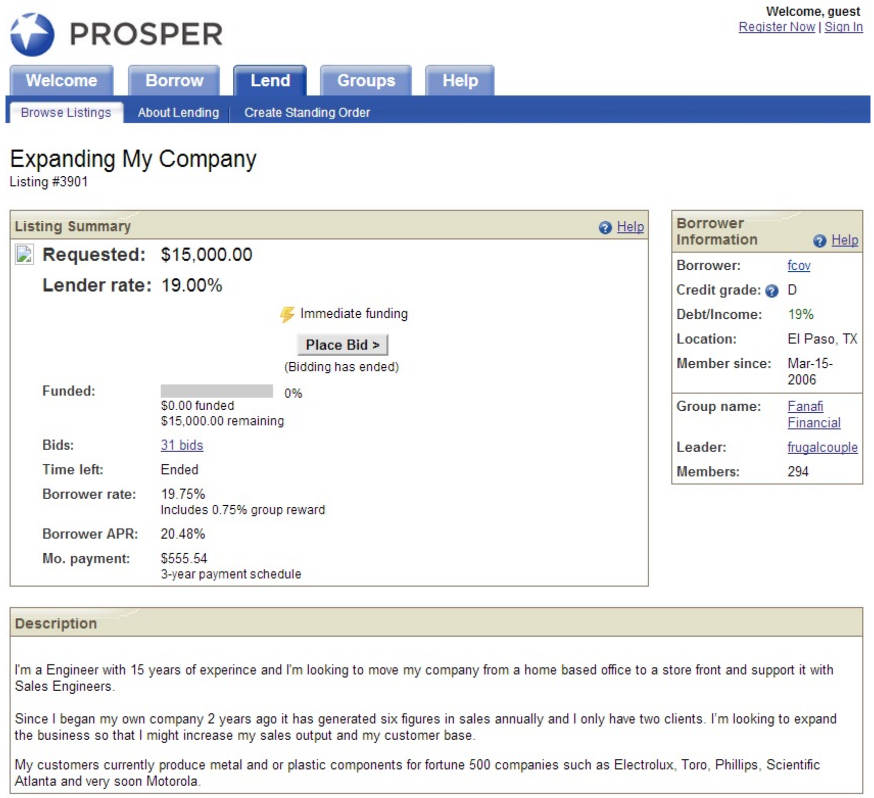

Note: A screenshot of a Prosper Listing.

Each group was required to select at least one and up to nine of 1,552 available category labels. Groups chose how to portray themselves through the use of category labels, which were chosen by self-appointed group leaders. Non-group members cannot affiliate with categories themselves. Categories could signify occupations (i.e. nurses, IT professionals, consultants); geographic locations (i.e. Bay Area, Midwest, East Coast); alumni affiliations (Stanford University MBAs, Penn State Nittany Lions, Colorado); as well as other socially or professionally distinctive dimensions. The variety of category labels presents the opportunity for hybrid candidates. Following Albert and Whetten (Albert and Whetten 1985), we define a *hybrid* as a social actor "whose identity is composed of two or more types that would not normally be expected to go together." Combinations of categories with which groups identify can range from very pedestrian, such as "Financial planning" and "Investments," to more unconventional combinations such as "Startups" and "Pharmaceutical", or "Vermont"

Figure 14: Screen Shot of Group Page with Category Labels

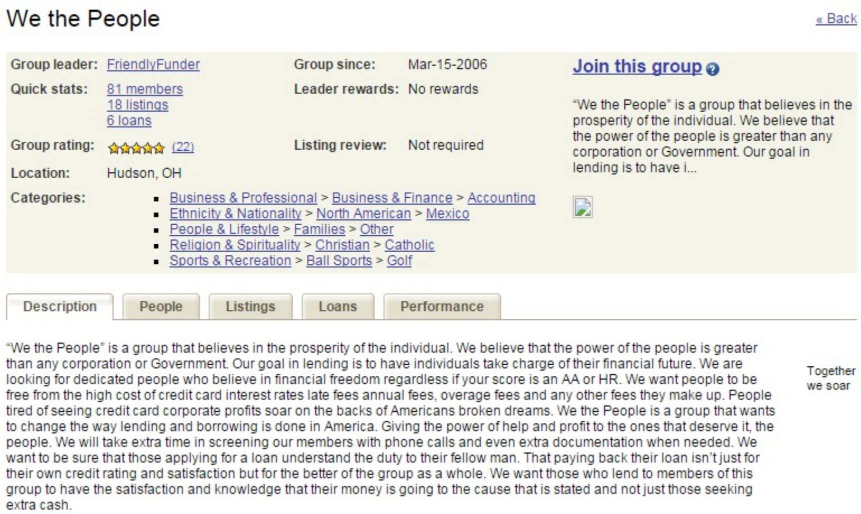

Note: A screenshot of a Prosper Listing.

and "Neighborhood Organizations". We define unconventional groups as those that affiliate with categories that are not normally seen together.

Groups used these categories to communicate their commonalities, both to attract new group members as well as to entice lenders to pay attention to them. Users on the site searched for groups to join through these category labels. If they did join a group, they were limited to membership in one group only, thereby tying their public identity to the set of labels attached to that group. When viewing a borrower's loan request, prospective lenders can see details of the group to which the member belongs. Prominent among this information is the list of category labels with which their group is affiliated. Group categories are an extension of the borrower's identity. Just as names have been shown to affect audience perceptions of a firm's identity or legitimacy or success in the marketplace (Lee 2001), we posit that categories provide additional information to a lender in their funding decisions (Zhao and Zhou 2011).

HYPOTHESES

Diversification and Unconventional Borrowers

The crux of the 'categorical imperative' observed by Zuckerman (1999) hinges on the fact that the objective of the resource holders in the marketplace are focused on selecting objects based on familiarity and similarity in order to compare and value them. This article established the primacy of the audience as the arbiter of whether a candidate succeeds or not in getting their offerings accepted. Though theoretical developments have enhanced our understanding of audiences (Hannan, Pólos, and Carroll 2007), empirical investigation as to how audiences may process category spanning differently has only recently taken hold. We seek to elucidate the audience motivations for candidate selection.

Investors who learn to spot undervalued investments are advantaged. Learning to recognize these opportunities may require narrowing the range of investments they consider in order to focus their attention on understanding few of them well. Similar to the finance analysts in Zuckerman's (1999) seminal study, the best way to truly valuate investments is to narrowly circumscribe the types one decides to investigate. By doing so, an investor can become specialized in particular areas and therefore develops an ability to distinguish a 'good' investment or not. For example, when choosing of job candidates, it becomes much more difficult to compare applicants with a social psychological background, sociology backgrounds, and economics training partly because they are very different, but partly because we are products of our specialized training and therefore we would likely only be comfortable focusing on interviewing, evaluating, and hiring from one (familiar) discipline.

On the other hand, audiences (in our case: individual investors) may be motivated by very different reasons than the finance analysts in Zuckerman's (1999) study. For example, venture capitalists are constantly looking to redefine the market and therefore are attracted to 'ambiguous' firms as investments (Pontikes 2012). In this setting, we believe that the objective of the audience (the lenders) leads to a very different conclusion. In particular, audiences may instead be rewarded for the breadth of options they consider. In these settings, we should expect there to be an incentive to pay attention and consider less prevalent candidates. We expand on this idea below.

Prosper is a financial marketplace for unsecured consumer debt. The resource holders are those with money to lend. Their objective, as with most investment arenas, is to maximize their return while simultaneously reducing their risk. One touted strategy of successful investing in financial markets is to diversify one's holdings (Markowitz 1952). As Cervantes suggests in Don Quixote, "Tis the part of a wise man to keep himself today for tomorrow, and not venture all his eggs in one basket," (trans. Motteux 1719; emphasis added) the risk of an investment portfolio is reduced if investments are spread across several different securities. The theory behind this belief is that holding a portfolio of securities which move in an uncorrelated fashion with one another will have less risk because securities that lose value will be offset by uncorrelated securities which may gain value under the same macroeconomic circumstances. A fully diversified portfolio is then no longer subject to unsystematic risk.

In lay terms, diversification can mean investing in securities that are not similar to one another along some recognizable dimension. Active conversations on Prosper-related message boards concur and even provide specific characteristics in a borrower that may be diversifiable:

> "To truly diversify, it's important to distribute those eggs across multiple baskets. Spread loans across multiple states to help avoid local downturns and across different credit grades, incomes, and jobs to target different steps on the socio-economic ladder. By spreading out your lending, you can reduce correlation between loans and reap the benefits of true diversification."
>
> (Mike, Prosperous Land Blog, 2008)

As this lender suggests, Prosper users should aim to invest in borrowers who are different from one another on any observable dimensions. We believe that the audience in this market setting will actively seek out borrowers who have little in common with one another. In fact, they are continually motivated to look for sources of information to identify potential differences between borrowers, as evidenced by how this lender laments the lack of tools Prosper provides to perform such analyses:

> "A while back, I had written to encourage Prosper lenders to truly diversify their loans. It's not enough to spread your loans across different credit grades if they're all in California, for example. ...There are no readily available tools, Prosper or otherwise, to succinctly report performance by state or a lender's loan distribution over states. Beyond this, there are self-reported criteria from borrowers like occupation and the borrower's reason for getting the loan."

Mike (Prosperous Land Blog, 2008).

In contrast to the market characterized by Zuckerman (Zuckerman 1999), markets where the audience (or lender) is incented to seek out and identify a diverse set of candidates should be reflected by a preference for less conventional, versus prevalent, hybrids items. In fact, Prosper suggests that lenders look beyond the obvious dimensions of diversification to more nuanced ones:

> ". . . lenders should consider diversifying by investing in different types of Prosper Notes. And don't just limit yourself to a wide range of Prosper Ratings! You can also consider:
>
> - The term of the loan (one-year, three-year, and five-year)
> - The purpose of the loan (debt consolidation, home improvement, etc.)
> - The state where the borrower lives
> - The borrower's occupation"
>
> (Prosper Help, 2011).

We believe that this urge to diversify will motivate lenders on Prosper to seek out and fund listings from borrowers that have a more unique set of characteristics. Unconventional group combinations represent one such an opportunity.

Candidates that exhibit familiar identities may avoid the threat of illegitimacy, but they are also more exchangeable with other candidates (Merluzzi and Phillips 2014). Lenders should be able to valuate borrowers with familiar category pairings more readily because of their own experience with borrowers with similar profiles. They are also better able to validate their valuations by examining the valuations of others. As lenders reach a consensus in valuation, lenders are less likely to see a borrower with a familiar identity as a 'diamond in the rough' from which they can extract more value. In contrast, as uncommon category combinations defy consensus in valuation, lenders are more likely to see these candidates as untapped opportunities.

Borrowers with conventional identities are also less attractive from a diversification standpoint. For example, a borrower that is a represented by common categories, such as "Financial planner" and "Investments," will be more prevalent, and therefore more likely to be held already as an investment. Borrowers that identify with a less-common combination of categorical identities as a function of their group membership represent unique opportunities for lenders to invest in loans which are less likely to be correlated with any other ones they may already be holding. This is because unconventional borrowers represent a rare intersection of identities. For example, the categories of "Environmental" and "Entrepreneurs" do not occur together very often. To the extent that these affiliations represent a borrower's background and situation, and that these elements bear some relationship with their ability to pay back their loans, then they represent a rare instance of a borrower. The intersection of these categorical identities represents a distinctive investment in the sense that the characteristics it brings are not likely related to an investor's other holdings.

The preceding suggests that to the extent a setting consists of audience members who are motivated to diversity their resource distributions, we should expect:

Hypothesis 1: The more unconventional the candidate, the more funding they will receive.

Engagement

Participating in the Prosper market requires learning to understand and recognize how the borrowers are organized. In particular, we contend that lenders can vary in their ability to recognize these unconventional borrowers. Specifically, we suggest that those lenders who have more experience searching for and comparing borrower listings - that they are more engaged in this marketplace - will be more successful at findings these unconventional investments. Because the more often a lender encounters categorized objects, the better they becomes at recognizing them. This increased facility with categorical conventions, or "fluency," as Koçak, Hannan, and Hsu Kocak, Hannan, and Hsu (2014) (Kocak, Hannan, and Hsu 2014) suggested, follows from ". . . engaged members communicat[ing] with others about the category and develop greater consensus about the meaning of its label, these audience members can discuss producers and offerings in more specific detail." This stems from more engaged audience members having a more complex category understanding to draw upon.

By engaged, we mean more experience with the market domain. For example, increased engagement can be reflected by an audience member's ability to demonstrate use of categorical conventions, such as esoteric abbreviations in the case of eBay seller (Hsu, Koçak, and Hannan 2009). With regard to hybridity, Rao and his co-authors (Rao, Monin, and Durand 2005) have demonstrated in the case of French cuisine that even professional critics have to gain experience with a novel cuisine type before being able to adequately evaluate it. The critics with the Guide Michelin were dismissive of Nouvelle Cuisine in the beginning because they found it difficult to understand: Guide Michelin critics had to 'learn' how best to recognize and value these novel combinations of cuisines. As audiences observe more instances of hybridization, they learned to recognize and evaluate them.

A thought exercise may elucidate our assertion for our setting. Take two lenders: one who uses a simple categorization system (Lender 1) and one who relies on a more complex one (Lender 2). Lender 1's categorization system may only be able to recognize three types of listings, say categories A, B, and C. As Lender 1 examines listings, they only recognize the differences that partition the listings into these three categories. Therefore, as Lender 1 is searching listings among which to diversify their holdings, they may only look to invest in these three types of listings (A, B, C), which comprises the universe of their recognition. On the other hand, hypothetical Lender 2 is more engaged, and therefore has a more detailed categorical understanding of the listing universe. Lender 2 recognizes up to 5 categories of listings, A, B, C, D, and E. In Lender 2's case, they would be willing to invest in each of these 5 listing categories, thereby covering a more diverse range. Given this, if a listing is unconventional, lender 1 will be less likely to recognize it because they would be more likely to bundle it into a more inclusive category - thereby watering down its uniqueness. Lender 2 will be more likely to have created more cognitive niches in their understanding of the market for listings, and therefore be better able to identify a less-common listing. This does not necessarily imply Lender 2 will hold a greater number of listings, as Lender 1 could hold the same number, but within fewer recognized categories. Therefore, we expect that:

Hypothesis 2: The more engaged a lender is, the more likely they will invest in unconventional candidates.

METHODS

Data and Sample

Prosper freely provides data of lending activity on their website. We retrieved data on listings posted prior to September 12, 2007, the date when the group category labels were no longer visible to potential lenders. We eliminated listings that were cancelled by Prosper and that were withdrawn by borrowers within a day of posting. This data cleaning yielded a population of 127,700 valid listings. We accessed information on the borrowers and groups associated with these listings. Of the listings we examined, 59,930 (46.93%) were affiliated with a Prosper group. We utilized this listing-level data, including the personal and financial information associated with each listing, to test our expectations regarding the attention accrued to unconventional candidates.

Our second hypothesis suggests that engaged experts bid differently than less engaged ones, and therefore is examined at the bid-level. We retrieved all recorded bids the 127,700 valid listings received, including the bid creation time, amount of the bid, and the outcome (i.e., whether the bid won or not). In all, 28,379 lenders made 2,019,830 bids on group- and non-group-affiliated listings, with an average of 71.07 bids per lender. We are interested in propensity of engaged audiences to support unconventional candidates; thus, our population of interest is all possible lender-listing combinations. Since the total number of possible combinations is enormous (on the order of 3.624×10), we draw a random 1% sample (k = 288) from the lenders that ever bid on a group-affiliated listing and populated the possible lender-listing pairings for these users (288 lenders × 127,700 valid listings = 36,777,600 potential observations). Sampling at the lender level is appropriate since we need the comprehensive investment histories of these lenders. In all, our 288 lenders made 14,393 bids on our valid set of listings.

We eliminated pairings in which the listing originated before the lender was a registered member of Prosper, and we dropped lender-listing pairings that occurred after the lenders' last recorded investment. This truncation was necessary since we wanted to examine lender's bidding activity during times when they could reasonably be at risk to bid. Based on these parameters each lender was at risk of bidding on 47,370 listings on average (minimum: 1,672; maximum: 127,623; standard deviation: 35,370). Lenders may bid multiple times on the same listing for a variety of reasons, including when they are 'outbid' by other lenders, or when a lender wants to provide additional funding to the borrower. Each lender bids once or twice on a listing on average (mean 1.23, std. dev. 0.72). For the purpose of creating time-varying lender variables (e.g., engagement, described below), we aggregated the total amount a lender successfully invested in a listing. Our analysis dataset for Hypothesis 2 consisted of 13,656,612 lender-listing observations.

Analyses for both hypotheses depended on Prosper group information. We accessed category labels, creation dates, and loan histories for all groups whose members created listings before September 12, 2007. Our listing population consisted of loan requests from 66,746 distinct borrowers, of whom 25,123 (37.64%) were affiliated with a Prosper group. Members from 1,061 different groups created loan requests during the period studied.

Note that two classes of listings deserve attention: listings with only one group category, and listings with no group affiliation (and thus no group category). Together, these listing types comprise 56.81% of the listings on Prosper. We are primarily interested in how actors who span dissimilar categories are evaluated; such spanning obviously requires claims of membership in multiple categories. We follow prior research which has reasonably assumed that single-category candidates do not present the opportunity for hybridity, and thus are more conventional (Hannan, Pólos, and Carroll 2007; Hsu 2006; Kovcs and Hannan 2013; Rao, Monin, and Durand 2005). Similarly, we argue that listings without any labels also provide no jarring contrast, and thus will appear more conventional relative to listings that span multiple categories. Thus, single-label- and no-label- listings may be evaluated similarly by audiences.

Our hypotheses above rely on the fact that borrower group membership is recognized by the market. In particular, we assume that market participants consider this as a factor in their decision to fund a loan. Because borrowers on Prosper comprise of both group and non-group members, we can test the assertion as to whether the market is sensitive to, or otherwise take into account, whether a borrower is part of a group or not.

Dependent Variables

Percent Funded. For the first hypothesis, we analyzed the effect of the unconventionality of a group's category combination on the percent funding a listing received. This was calculated by dividing how much money borrowers asked for by how much total money the listing was able to raise from potential lenders. This captures how much interest or appeal a listing generated. This variable ranged from 0% to 100%. Listings became loans only if they receive 100% of the amount requested by the prospective borrower. Of the 127,700 valid listings that where posted to the website before September 12, 2007, 14,619 (11.4%) were funded and became loans, including 9,323 listings from group members. Of the 59,930 group-affiliated loan requests we examined, approximately 12.9% become loans.

Bid. For the second hypothesis, we examined whether or not a lender decided to bid on a listing. We coded the dependent variable bid based on whether the lender actually bid on the listing (1) or not (0).

Independent Variables

Unconventionality. Our main independent variable of interest is the unconventionality of listing in terms of the categories associated with its Prosper group. We operationalize unconventionality as the average categorical distance among a listing's pairs of labels. In our analysis, 'distance' is a function of how often a group's categories have co-occurred in the marketplace. We computed a listing's unconventionality in three steps. First, we obtained cumulative counts of category- and category-pair-occurrences for each listing, and used these counts to calculate the Jaccard index (1901) of similarity for each category pair. For each pair of labels associated with each listing, the similarity index is calculated as follows:

$$similarity_{xij} = \frac{|i \cap j|}{|i \cup j|}$$

which is simply the ratio of the number of times category i and j have occurred together (including the current co-occurrence) to the number of times they have occurred together and separately. By this definition, similarity is greater than 0 and less than or equal to 1, with higher values corresponding to higher similarity.

Second, we transformed the similarity indices to obtain measures of distance between category pairs. Research in psychology has posited a negative exponential relationship between generalization and distance (Shepard 1987); we calculated distance accordingly ($distance_{ij} = -ln(sim_{ij})$).

Third, following Kovács and Hannan Kovcs and Hannan (2013) , we calculated the listing's average distance ($\bar{d}_x$) among category pairs:

$$\bar{d}_x = \frac{\sum_{i,j \in J_x} distance_{ij}}{n(n-1)}, \quad n > 1$$

where n is the number of categories associated with the group to which a listing belongs, and J_x is the set of all category pairs associated with that listing. We defined unconventionality as follows:

$$unconventionality_x = \begin{cases} \bar{d}_x & \text{if } n > 1 \\ 0 & \text{otherwise} \end{cases}$$

By design (Kovcs and Hannan 2013), this measure of unconventionality (1) has a minimal value of zero (when the listing has only one category), and (2) increases as the total distance among categories increases:

$$\frac{\partial \bar{d}_x}{\partial \sum_{i,j \in J_x} distance_{ij}} = \frac{1}{n(n-1)} > 0, \quad n > 1.$$

In the simple case that a listing's group was identified with either only one category or no categories, the listing spanned no categories and thus had zero distance among pairs. Such listings received an unconventionality score of 0. The categories we consider are the labels associated with a listing's Prosper group. Since only group-affiliated listings have labels, our final models utilize only the 59,930 listings with these labels.

Engagement. The second hypothesis pertained to audience engagement in the market. Engagement was operationalized as the moving average of dollars invested per day of a member's history in Prosper. For each lender-listing pair, we calculated the cumulative dollars the lender had invested in Prosper by the time the focal listing originated. We then divided this cumulative sum by the number of days the lender had been registered on Prosper by the time the focal listing went live. This is described in the following equation:

$$engagement_{k,x} = \frac{\sum_{n \in N_x} dollars\, invested_n}{days_x}$$

where N_x refers to the set of loans the lender (k) had ever funded (not simply bid on) by the time of the listing's (x) origination, $dollars\, invested_n$ refers to the total amount that the lender funded a previous loan (n), and $days_x$ is the number of days the lender had been registered with Prosper by the time the focal listing (x) originated.

Control Variables

Many (unobservable) factors contribute to whether a listing is funded or not. Our data do not give us access to lenders' specific preferences or when exactly they were browsing the available listings online. However, we are not trying to model a lender's bidding behavior exhaustively, but are instead attempting to isolate the effect of unconventionality on audience attention from other effects. Our approach benefits from the fact that we have access to virtually every detail market audiences would have had in evaluating these listings.

To help identify the effect of a listing's unconventionality on its degree of funding—and on a lender's propensity to bid on a listing—we control for several borrower-level attributes collected by Prosper. These included the borrower's amount requested, the maximum borrowing rate they are willing to endure, debt-to-income ratio (logged), homeowner status, credit rating, and the interest rate of the loan. We created a dummy variable image indicating whether the borrower included an image in the listing (1) or not (0). Prosper borrowers can choose to have their listing close once it has received enough funding from lenders, or to keep their listing open as additional bids drive down the borrowing rate. We included a dummy variable funding option, indicating whether the borrower chose the close-when-funded (0) or open-for-duration (1) option. Each of these variables was included to dismiss the possibility that the benefit of unconventionality is attributable to some systematic difference other than spanning distant categories.

We control for possible group-level differences (other than unconventionality) by including cumulative counts (logged) of the number of listings and the loans generated by each group. In testing our first hypothesis, we leverage the panel nature of our data and include group random effects. These controls should address concerns that effects for group unconventionality reflect other underlying differences between groups.

Finally, we also included the total number of categories affiliated with each listing's group, as past research has demonstrated how the greater number, the more confusing they will be perceived (Hannan, Pólos, and Carroll 2007; Hsu 2006; Leung and Sharkey 2014). Table 13 summarizes our variables and reports their correlations.

Models

For the first hypothesis, we model the percent funding that a listing eventually received as a function of the listing's unconventionality using ordinary least squares (OLS) regression. Alternate models not presented here—including tobit regression accounting for the boundedness of percent funding between 0 and 1, and logistic regression with percent funding recoded as fully funded (1) or not (0)—were also estimated, and provided results consistent with those of OLS. We begin with the full population of Prosper listings—both group and non-group—to assess the general benefit of group membership and unconventionality, respectively. Next, we restrict our analysis to group-affiliated listings to test our hypothesis more stringently.

For the second hypothesis, we modeled the likelihood that a lender would bid on a listing given the listing's unconventionality and the lender's degree of engagement. Because the dependent variable is dichotomous, logistic regression would normally be appropriate. One important consideration in testing the second hypothesis is unobserved lender heterogeneity, which, if not addressed, could threaten the identification of our effect of interest. Because we have repeated observations on lenders' bidding behavior, we can control for such lender differences using within-lender fixed effects. Due to the size of the analysis dataset, however, pre-programmed commands for fixed-effects logistic regression (e.g., xtlogit in Stata) cannot readily provide a result. As a solution, we constructed a within-estimator (Wooldridge 2010) by first de-meaning the dependent and independent variables and then estimating OLS regression. We corrected the standard errors using the usual procedure (multiply OLS standard errors by $\sqrt{K/(K-1)}$, where K is the number of lenders; Wooldridge, 2010). While the coefficients resulting from this procedure are not interpretable as odds ratios, as in the case of logistic regression, the tests of significance on these coefficients provide reasonable inference, since the demeaned-dependent variable becomes essentially a continuous measure varying between 0 and 1. Similar to the analysis for hypothesis 1, we first estimate models using all valid data (group- and non-group-affiliated listing-lender pairs), then constrain our analysis to group-affiliated data only. We estimated a control model, then added unconventionality, engagement, and an unconventionality-by-engagement interaction in a stepwise manner.

Table 13: Descriptive Statistics

Variable	Mean	s.d.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1 Percent Funded	0.132	0.297	1															
2 Amount Requested	7316.641	6251.426	-0.054	1														
3 Borrower Max Rate	0.175	0.071	0.17	-0.057	1													
4 Credit AA	0.023	0.15	0.206	0.082	-0.166	1												
5 Credit A	0.024	0.153	0.172	0.115	-0.129	-0.024	1											
6 Credit B	0.039	0.194	0.165	0.163	-0.1	-0.031	-0.032	1										
7 Credit C	0.075	0.263	0.136	0.163	-0.068	-0.044	-0.045	-0.057	1									
8 Credit D	0.116	0.32	0.075	0.102	-0.022	-0.056	-0.057	-0.073	-0.103	1								
9 Credit E	0.185	0.389	-0.035	-0.022	0.066	-0.073	-0.075	-0.096	-0.135	-0.173	1							
10 Credit NC	0.004	0.066	0.008	-0.029	-0.001[a]	-0.01	-0.01	-0.013	-0.019	-0.024	-0.032	1						
11 Image	0.479	0.5	0.15	0.018	0.109	0.027	0.026	0.036	0.025	0.022	-0.011	-0.002[a]	1					
12 Homeowner	0.286	0.452	0.103	0.19	-0.071	0.154	0.09	0.104	0.141	0.036	0.019	-0.041	-0.008	1				
13 Debt-to-Income (logged)	0.607	1.616	-0.101	0.295	0.019	-0.054	0.018	0.058	0.08	0.098	0.058	-0.067	-0.006[a]	0.06	1			
14 Funding: Open	-1.426	1.126	0.113	0.096	-0.01	0.092	0.083	0.089	0.083	0.066	-0.018	-0.016	0.081	0.053	0.053	1		
15 Group Label Count	0.588	0.492	0.141	-0.093	0.136	-0.02	0.004[a]	0.017	0.021	0.013	0.01	0.012	0.168	-0.006[a]	-0.065	0.082	1	
16 Unconventionality	2.025	2.312	0.187	-0.04	0.126	0.004[a]	0.018	0.029	0.036	0.029	0.003[a]	-0.001[a]	0.144	0.01	-0.015	0.1	0.654	1

Note: [a]Not statistically significant. All other correlations are significant at $p < .01$

RESULTS

Figure 15 compares the densities of unconventionality among listings that became loans and those that did not. Note that the listings are monotonically decreasing in prevalence as they get more unconventional, as expected. More important is the fact listings that eventually become loans tend to be less conventional than listings that are not funded. The graph demonstrates that the unconditional chance of a listing being fully funded is highest for those listings that are not the most common. The likelihood of the rarest listings being funded is difficult to discern, as there are very few observations. These graphical results are merely suggestive of our contention and need to be validated with statistical tests, which we perform below.

Table 14 contains the results of OLS and random-effects regression used to examine hypothesis 1. The coefficients corresponding to control variables are in the expected directions. The greater the maximum borrowing rate the borrower was willing to endure, the greater percentage funding a listing received. Better credit scores were also associated with higher levels of funding, as was homeownership and including an image with the listing. Borrowers with high debt-to-income ratios (logged) received less funding.

We test our hypothesis on unconventionality in Model 2 on the complete population of borrowers (group and non-group) and in Model 5 on the more restricted sample of just group affiliated borrowers. As expected, the coefficients associated with unconventionality are positive and statistically significant. The greater the unconventionality of a listing's group affiliations, the greater percentage funding the listing received (Model 2: 0.054, $p < 0.01$; Model 5: 0.063, $p < 0.01$). These results are not sensitive to the choice of sample (i.e., all listings vs. group-affiliated listings only). This provides broad support for hypothesis 1 and suggests that unconventional listings not only fare better relative to unaffiliated listings, they also garner more funding compared to more conventionally categorized listings. Results were not sensitive to alternative model specifications, such as the logistic and Tobit models. For brevity, these models are not reported.

The results for category count also deserve discussion. The coefficient is negative both in the presence of controls for group affiliation (Model 1) and our measure of unconventionality (Models 2, 5). At the base level, this is consistent with the literature on the negative effects of category spanning. More interestingly, our positive result for unconventionality suggests a more nuanced view. Spanning disparate categories can be beneficial in certain markets. Additionally, the effect for group is positive (Model 2: b = 0.097, $p < 0.01$). This suggests that membership in a group is observed by lenders and can have a positive impact on the likelihood of funding.

Figure 15: Density of Unconventionality by Whether a Listing Became a Loan

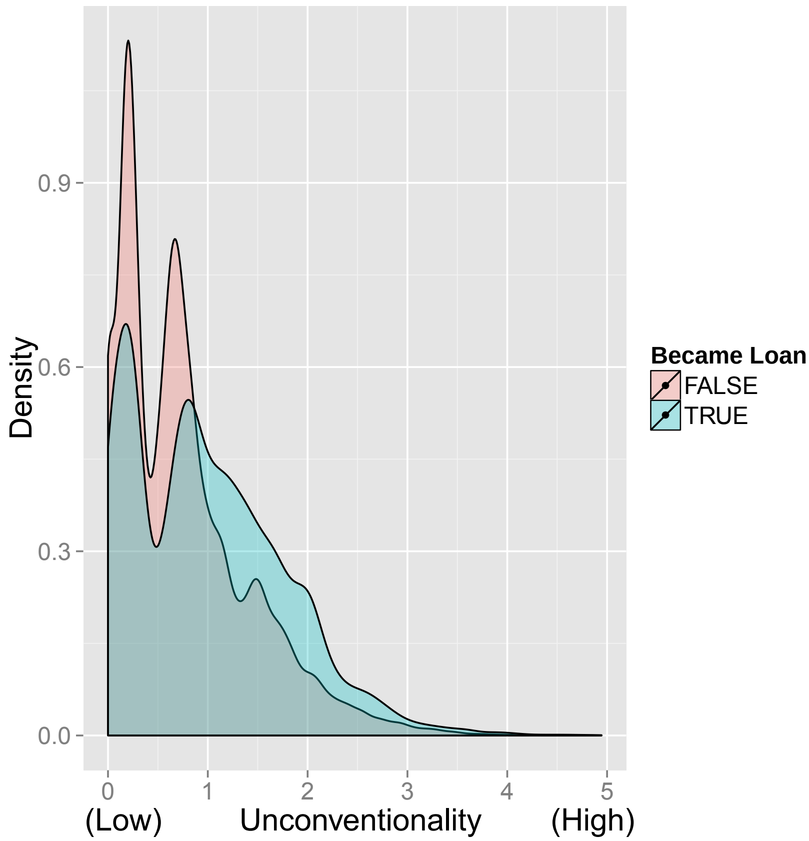

Note: Density plot of unconventionality decomposed by listing's final status: whether or not it became a loan. Data are the 59,189 listings that have group label information available.

Table 14: OLS Regression: Percent Funded

	Model 1	Model 2	Model 3	Model 4	Model 5
Amount Requested	-0.000**	-0.000**	-0.000**	-0.000**	-0.000**
	(-1.25E-07)	(-1.24E-07)	(-2.20E-07)	(-2.20E-07)	(-2.20E-07)
Borrower Max Rate	1.152**	1.137**	1.415**	1.418**	1.414
	(-0.01)	(-0.01)	-0.018	(-0.018)	(-0.018)
Credit AA	0.592**	0.585**	0.653**	0.652**	0.653
	(-0.005)	(-0.005)	(-0.009)	(-0.009)	(-0.009)
Credit A	0.522**	0.516**	0.567**	0.566**	0.567
	(-0.005)	(-0.005)	(-0.008)	(-0.008)	(-0.008)
Credit B	0.422**	0.417**	0.468**	0.468**	0.469

Continued on next page

Table 14 – *Continued from previous page*

	Model 1	Model 2	Model 3	Model 4	Model 5
	-(0.004)	(-0.004)	(-0.006)	(-0.006)	(-0.006)
Credit C	0.292**	0.287**	0.343**	0.344**	0.345
	(-0.003)	(-0.003)	(-0.005)	(-0.005)	(-0.005)
Credit D	0.185**	0.181**	0.221**	0.221**	0.222
	(-0.002)	(-0.002)	(-0.004)	(-0.004)	(-0.004)
Credit E	0.066**	0.066**	0.079**	0.079**	0.079
	(-0.002)	(-0.002)	(-0.003)	(-0.003)	(-0.003)
Credit NC	0.086**	0.088**	0.110**	0.107**	0.107
	(-0.011)	(-0.011)	(-0.016)	(-0.016)	(-0.016)
Has Image	0.043**	0.044**	0.039**	0.040**	0.04
	(-0.001)	(-0.001)	(-0.002)	(-0.002)	(-0.002)
Is Homeowner	0.013**	0.012**	(0.018**	0.019**	0.018
	(-0.002)	(-0.002)	(-0.003)	(-0.003)	(-0.003)
Debt to Income (logged)	-0.027**	-0.028**	-0.035**	-0.034**	-0.034
	(-0.001)	(-0.001)	(-0.001)	(-0.001)	(-0.001)
Funding: Open	0.010**	0.009**	-0.001	-0.001	-0.002
	(-0.001)	(-0.001)	(-0.002)	(-0.002)	(-0.002)
Category Count	-0.012**	-0.002**	-0.002	0.001	-0.022**
	(-0.001)	(-4.04E-04)	(-0.004)	(-0.004)	(-0.004)
In Group	0.097**				
	(-0.002)				
Unconventionality		0.054**			0.063**
		(-0.003)			(-0.005)
Group: Cum Listings				-0.017**	-0.006*
				(-0.003)	(-0.003)
Group: Cum Loans				0.009**	0.006*
				(-0.003)	(-0.002)
Constant	-0.201**	-0.192**	-0.154**	-0.143**	-0.171**
	(-0.02	(-0.011)	(-0.016)	(-0.016)	(-0.016)
N	127,700	127,700	59,930	59,930	59,930
R2	0.29	0.3	0.29	0.3	0.3
F(15,127684)	3534.99**	3595.45**	-	-	-
Wald chi2	-	-	18752.20**	18897.06**	19109.97**
Df	-	-	14	16	17
Group RE	No	No	Yes	Yes	Yes
Number of Groups	-	-	1060	1060	1060

$^{**}p < 0.01, ^{*}p < 0.05$

Results of tests of the second hypothesis are presented in Table 15. Model 1 and 4 estimate control variables only. Borrowers who were asking for a larger loan amount, offering to pay a higher interest rate, having a better credit rating, including an image, and owning a home are all positively correlated with the event of being bid on by a lender. Having a larger debt-to-income ratio was negatively associated with receiving a bid by any lender. At the group level, group members from groups which had more cumulative listings were less likely to receive a bid, but the more cumulative successful listings, the more likely another group member would receive a bid. Of note, similar to our finding above and consistent with previous literature, the greater the number of categories with which a listing is affiliated, the less likely it will be bid on by any lender (Leung and Sharkey 2014).

Table 15: OLS Regression: Bid - Within-Lender Estimators

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Amount Requested	6.99e-08**	6.82e-08**	1.16e-07**	1.16e-07**	1.14e-07**	1.15e-07**
	(-2.12E-09)	(-2.12E-09)	(-4.28E-09)	(-4.28E-09)	(-4.28E-09)	(-4.28E-09)
Borrower Max Rate	0.021**	0.021**	0.032**	0.031**	0.031**	0.031**
	(-1.86E-04)	(-1.86E-04)	(-3.55E-04)	(-3.57E-04)	(-3.58E-04)	(-3.58E-04)
Credit AA	0.012**	0.012**	0.016**	0.015**	0.015**	0.015**
	(-8.42E-05)	(-8.43E-05)	(-1.70E-04)	(-1.71E-04)	(-1.72E-04)	(-1.72E-04)
Credit A	0.012**	0.011**	0.016**	0.015**	0.015**	0.015**
	(-8.06E-05)	(-8.06E-05)	(-1.52E-04)	(-1.53E-04)	(-1.53E-04)	(-1.53E-04)
Credit B	0.009**	0.009**	0.012**	0.012**	0.012**	0.012**
	(-6.47E-05)	(-6.47E-05)	(-1.19E-04)	(-1.20E-04)	(-1.20E-04)	(-1.20E-04)
Credit C	0.005**	0.005**	0.008**	0.007**	0.007**	0.007**
	(-4.88E-05)	(-4.88E-05)	(-9.05E-05)	(-9.15E-05)	(-9.16E-05)	(-9.15E-05)
Credit D	0.003**	0.003**	0.005**	0.005**	0.004**	0.004**
	(-3.97E-05)	(-3.97E-05)	(-7.47E-05)	(-7.54E-05)	(-7.54E-05)	(-7.54E-05)
Credit E	0.001**	0.001**	0.001**	0.001**	0.001**	0.001**
	(-3.37E-05)	(-3.37E-05)	(-6.26E-05)	(-6.27E-05)	(-6.28E-05)	(-6.28E-05)
Credit NC	0.001**	0.001**	0.001**	0.001*	0.001*	0.001*
	(-2.27E-04)	(-2.27E-05)	(-3.75E-04)	(-3.75E-04)	(-3.75E-04)	(-3.75E-04)
Has Image	0.001**	0.001**	0.001**	0.001**	0.001**	0.001**
	(-2.53E-05)	(-2.52E-05)	(-4.72E-05)	(-4.73E-05)	(-4.73E-05)	(-4.73E-05)
Is Homeowner	2.48e-04**	2.47e-04**	2.31e-04**	2.66e-04**	2.77e-04**	2.76e-04**
	(-2.82E-05)	(-2.82E-05)	(-5.28E-05)	(-5.28E-05)	(-5.29E-05)	(-5.29E-05)
Debt to Income (logged)	-4.56e-04**	-4.64e-04**	-0.001**	-7.24e-05**	-7.23e-04**	-7.22e-04**
	(-1.14E-05)	(-1.14E-05)	(-2.19E-05)	(-2.19E-05)	(-2.19E-05)	(-2.19E-05)
Funding: Open	0.001**	0.001**	0.001**	0.001**	0.001**	0.001**
	(-2.60E-05)	(-2.60E-05)	(-5.02E-05)	(-5.04E-05)	(-5.05E-05)	(-5.05E-05)
Category Count	-1.07e-04**	9.32e-05**	-1.66e-04**	-1.35e-04**	-8.08E-06	-1.03E-06
	(-1.54E-05)	(-7.58E-06)	(-1.87E-05)	(-1.93E-05)	(-2.08E-05)	(-2.08E-05)
Lender Engagement	2.47e-04**	2.36e-04**	1.91E-05	3.15e-04**	3.33e-04**	1.97E-05
	(-1.07E-05)	(-1.07E-05)	(-2.33E-05)	(-1.95E-05)	(-1.95E-05)	(-2.34E-05)
In Group	0.002**					
	(-7.17E-05)					
Unconventionality		0.001**	1.36e-04**		-0.001**	-0.001**
		(-2.70E-05)	(-3.71E-05)		(-3.15E-05)	-3.15E-05
Group: Cum Listings (logged)				-0.001**	-0.001**	-0.001**
				(-2.79E-05)	(-3.07E-05)	(-3.07E-05)
Group: Cum Loans (logged)				0.001**	0.001**	0.001**
				(-3.03E-05)	(-4.83E-05)	(-5.07E-05)
Unconventionality × Engagement			3.55e-04**			0.0004**
			(-1.56E-05)			(-1.56E-05)
Constant	0	0	0.001**	0.001**	0.001**	4.96e-05**
	(1.23e-5)	(1.44e-5)	(4.96e-05)	(2.42e-05)	(2.42e-05)	(2.42e-05)
N	13,656,612	13,656,612	6,007,072	6,007,072	6,007,072	6,007,072
R2	0.01	0.01	0.01	0.01	0.01	0.01
F-test	5445.77**	5142.84**	5160.30**	2549.49**	2595.14**	2466.81**
df	(15, 13656612)	(16, 13656595)	(16, 13656595)	(17, 6007054)	(17, 6007054)	(18, 6007053)

Continued on next page

Table 15 – *Continued from previous page*

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
$^{**}p < 0.01$,$^{*}p < 0.05$						

Model 3 of Table 3 contains the main effects for listing unconventionality and lender engagement. Results support our contention that the greater the engagement a lender has with this market, the more likely they will be to bid on a listing from a borrower that is affiliated with rare category combinations (Model 2: b = 0.0004, $p < 0.01$). Our coefficient of interest is the odds ratio associated with the engagement × unconventionality interaction effect. The positive and significant estimate on the interaction of a lenders engagement with a listing's measure of unconventionality suggests that as a lender becomes more engaged in the marketplace, they will be increasingly likely to bid on these less conventional listings. Our hypothesis holds in model 6 as well, which is estimated only on the population of listings from borrowers which belonged to a group (Model 6: b = 0.0004, $p < 0.01$).

Robustness Checks

There are two potential alternative explanations that we investigated. First, we considered whether these listings are actually "better" listings to invest in. While we are controlling for a host of observable measures which affect an individual listing's quality (e.g. credit rating and debt-to-income ratio), one could still suggest that perhaps they are less likely to default on their loans. The market, learning this, would be able to identify these better investments, and therefore, were more likely to fund them. To address this possible complaint, we estimated the likelihood that a loan will be fully paid off as a function of the borrower's unconventionality of category combinations. In results not reported for brevity (but available from the authors), we find no evidence that unconventionality of a listing is correlated with being a 'better' borrower.

We also considered whether the positive effect for unconventionality reflects possible market-level difference in the timing of when these listings were offered. Certainly, the fortunes of financial market participants can change with the timing of when they enter. We therefore need to assess whether unconventional listings happened to be more likely to be listed during a particularly fecund time for any loan to get funded. For example, perhaps these listings occurred later in Prosper's history, thereby benefitting from the increased legitimacy of this market. We guard against this potential confound in two ways. First, we included a time trend control, which is a linear measure of the days since Prosper began accepting listings. This linearly increasing variable was not significant and did not affect the significance of our measure of unconventionality. Second, we include dummies for each week over our observation window. These dummy variables, in unreported analyses, also did not affect the significance of our effect.

DISCUSSION

This chapter addresses the imbalance to date which suggests that social actors who span disparate categorical distinctions are disadvantaged (Hannan, Pólos, and Carroll 2007; Hsu 2006; Zuckerman 1999). We identified a market context where the resource holders are motivated to diversity their allocations across candidates. In doing so, candidates who are less conventional were rewarded. In addition, we also identify how audiences can vary in their level of engagement in a market, with those that were more engaged preferring these unconventional candidates.

These findings are consistent with the more recent literature on categories and markets which examine how non-conformity may not always be detrimental (Ruef and Patterson 2009; Smith 2011). However, this chapter builds on this work by advancing the notion that the objective of market participants may vary (Pontikes 2012), and therefore, assumptions regarding their motivations should be carefully considered. Second, this chapter explores the fact that perceptions of audience members may vary depending on their past experiences. What is particularly striking is that we find that unconventional borrowers are rewarded for their less recognizable combinations of category labels. However, a reasonable person could suggest that there is a limit to just how unconventional a borrower can be before it becomes a detriment. In short, that there should be an inverse-U shaped relationship between being unconventional and success in this market. In analyses not reported for brevity, we did find such an effect which points to the fact that even in markets where resource holders prefer less conventional participants, there still is a limit to their ability to appreciate increasingly "strange" participants.

The idea that perceptual difficulties underlie the decisions that market participants make is certainly familiar to behavioral economists (Mullainathan and Tahler 2001) or decision making scholars (Bazerman and Moore 2013). However, for macro organizational scholars, the challenge has been to test micro-level theories of behavior with macro-level data. This is because it is often difficult to observe at a granular level with archival data to examine micro-level (individual) processes. On the other hand, lab studies, which excel at identifying mechanisms at the individual level, have difficulty resonating with research of macro-level phenomena. Because of the bid-level of data we had access to on Prosper (individual bids as well as listing performance) we were able to posit hypotheses at both the listing-level and lender-level. This study therefore represents an attempt to connect the micro mechanisms at the individual level with macro-level phenomena.

Another notable contribution of this chapter is evidence that the relationship between categories matters. Categories have been represented as separate and distinct entities, with little regard for how they may be related. Previous studies have assumed that organizations which span categories are treated similarly. In this study we consider how common a paring of categories are, and theorize as to how that may affect how they are perceived. Certainly, this is a more realistic view of categorization processes—for example who would be confused if a film spanned the categories of Romance and Comedy, while there would certainly be more confusion is a film attempted to incorporate elements from Documentaries and Romance.

In addition to extending the categorization literature, this chapter also elucidates an online peer-to-peer lending market, an increasingly popular, alternative form of financing. Recent turmoil in the more traditional markets for consumer debt should encourage market participants to seek additional understanding as to what factors affect decision making in financial markets. Beyond this, technological innovation has allowed these non-traditional markets to supplant what were traditionally bank or savings and loan-based businesses. However, in doing so, the online market has shifted how decisions are made on whether individual loans get approved from professional decision makers (a mediator) to a collection of individuals without experience to evaluate such loans (Iyer et al. 2009). This has implications for both candidates and audiences, which we discuss below.

First, much research in the finance has demonstrated the cognitive biases that investors suffer in more traditional markets, such as overreaction to earning reports (De Bondt and Thaler 1985) or a reluctance to realize losses (Odean 1998). While the focus of this chapter was not to understand an investor's success or failure in the marketplace, it does highlight a potential cognitive mechanism by which audiences, in this case investors, choose to make their investments. Particularly, since this market is comprised of non-professionals, our findings suggest that individual investors should be aware of potential shifts in their preferences for certain types of investment may be due to their past experiences and behaviors. Second, from the perspective of a candidate who is seeking attention from resource holders, this chapter would imply that if the resource holders are motivated to diversify their investments, then it behooves a candidate to attempt to portray themselves in as unique a way as possible. These findings then clearly have implications for other settings. Take labor markets. If an individual's identity can be seen as more or less novel as a function of their past experiences (Leung 2014), what is it that drives job candidates to choose to apply for particular jobs? Future research could extend on these findings by examining what factors lead social actors to compile a particular combination of categorical experiences.

A skeptical reader could suggest that these unconventional borrowers could be trying harder in some unobservable way, which leads to their being more likely to be funded. Perhaps because they are aware of their non-conforming position, they attempt to compensate in other, unobservable by us, ways which are reflected in their greater success. There are at least two reasons we do not believe this to be the case. First, as this is an online market, we have data on almost all the observables that affect ones likelihood of receiving funding. Therefore, we have been able to control for most of the individual listing level observables which the borrower can alter - for example, the amount of the loan or the interest rate they are willing to pay. Second, we have evidence which demonstrated that even within each individual lender there was heterogeneity in how their preferences changed. This suggests that, at least part, of our results are contingent on variation in the lenders and not the borrowers. However, it remains an open question as to whether these unconventional borrowers differed in a dramatically unobservable way from more conventional ones. Future research could examine this issue more directly.

We are making assumptions regarding the recognition or attraction that a certain combination of categories may elicit. For example, the combination the categories “Entrepreneur” and “Bay Area” and “Stanford” may not necessarily occur often, but is very well-understood. Also, certain combinations may also elicit more affinity from an audience regardless of how commonly they occur, “Veteran” and “Disabled” for example. Given the sheer number of possible combinations of merely two labels (1,552 choose 2 = 1,203,576) it would have been very difficult for us to code each combination for content. However, these category combinations can be controlled for in a lab environment, which certainly represents a way forward for this line of research and is something which the authors are exploring in future work.

Another particularly fruitful way forward from these findings could be to continue to investigate how heterogeneity in audiences affects other aspects of their evaluations. For example, to the extent that expertise in a domain causes the perceiver to hold more complex and nuanced views of the marketplace, then perhaps these same experienced audiences would evaluate candidates based on more detailed knowledge than less experienced audiences. Taking the example of employers in a labor market, this could imply that less experienced market participants could be more drawn to external signals of validity, such as symbols of status or reputation. On the other hand, more experienced audiences may be more comfortable evaluating the actual details of a candidate's past accomplishments and may not be swayed by status or recognition, which may seem superfluous to them.

5 Conclusion

The general project of this dissertation has been to investigate the role of labels in the processes whereby audiences assign value to candidates. On the whole, this work corroborates the continued assertion of economic sociology that markets are social systems consisting of actors subject to normative and cognitive constraints. At the same time, I have forwarded arguments and supplied evidence that break from previous studies of market categorization. Having presented my main empirical chapters, I now consider these essays together and outline paths for future research.

While the central focus is on labels and what their presence, absence, and combination does to candidate evaluation, this dissertation ultimately serves to broaden categories research to consider features and diversity among audiences. The first two empirical chapters of this dissertation (Chapters 2 and 3) address audience perceptions of candidates' between-category membership and candidates' within-category membership, respectively. I show that features inform audience judgments of both dimensions. The last empirical chapter (4) considers how unconventional label combinations are actually favorable to some audience segments. Each chapter utilizes the rich Prosper marketplace data to examine these questions.

In Chapter 2, I find that in the absence of labels, features can stand in to inform audiences whether candidates profess affiliation with one or multiple categories; consistent with the findings for multiply-labeled candidates, actors whose features span categories are discounted. Yet when labels are present, candidates are able to utilize category-spanning features without penalty. This presents a way that labels relax the cognitive constraints candidates and audiences otherwise confront.

In Chapter 3, I examine how a within-category dimension—candidates' typicality—contributes to evaluation. The received wisdom from previous market categories research suggests that it is generally good to be typical. Yet I find that this positive evaluation of typicality is contingent on the presence of labels. Strikingly, loan requests that are highly typical of their purpose category fare *poorly* when labels are absent, and *superbly* when labels are present. I posit typicality conveys different information depending on the context in which it is communicated. What was dull and terse in one condition (no labels) became a liquid commodity in another (labels).

Chapter 4 is most closely related to Chapter 2 in that it also considers category spanning: how conventional and unconventional label combinations—and candidates' consequent categorization—contribute to evaluation. This chapter goes further, however, in its attention to audiences. On average, unconventional label combinations are more likely to receive funding, and the more experience a lender has had with the Prosper marketplace, the more likely they are to bid on a listing with unconventional label combinations. In all, this pattern of results shows that multiple-category discounts are context- and audience-contingent.

Comparing the results of these chapters together provides further nuance and suggests avenues for future research. In Chapter 2, I pit normative and cognitive accounts of feature evaluation against one another. Finding that incoherence is no more scrutinized under conditions of uncertainty than conditions of certainty, I conclude that in nascent markets, assessments of category spanning are more a matter of a cognition rather than normative enforcement. That is, candidates with category-spanning features perform poorly because audiences cannot evaluate them readily, not so much because they actively devalue them (Zuckerman 1999). Yet in view of the results of Chapter 3, I see it would be erroneous to dismiss normative mechanisms from nascent markets entirely. In making a similar decomposition by borrower riskiness/quality, I find that borrowers with known high quality (high credit grade) exhibit a different relationship between typicality and funding than do borrowers with uncertain quality (those with neither high nor very low credit grade). This suggests that typicality is scrutinized more in the presence of uncertainty, which suggests a conscious, normative mechanism. Borrowers that have managed to differentiate themselves from others in their category, or that conform heavily, are alleviated from the uncertainty discount. This suggests that normative and cognitive mechanisms can be operative simultaneous, attending to different aspects of the same signals—here, coherence and typicality, respectively.

Furthermore, my study was (intentionally) restricted to a limited window of time; previous studies of norm enforcement have spanned many years of time (e.g., Ruef and Patterson 2009). Were macroeconomic conditions to permit more of a longitudinal design with Prosper, I could examine whether feature coherence figures into evaluations as lenders construct shared expectations of what feature combinations are permissible. Yet the finding that feature coherence separates candidates in unlabeled markets does revisit claims that formal structure is needed for categorization processes to have systematic consequences.

The u-shaped relationship between typicality and evaluation in the presence of uncertainty merits further research. The strategy of studying audience heterogeneity

 Chapter 5 Conclusion

in Chapter 4 could be a useful path forward: less-experienced audience members may be more dependent on typicality than more experienced audience members. Such would corroborate findings for a normative mechanism.

Another fruitful combination of the approaches of Chapters 3 and 4 concerns a different sense of 'typicality': the extent to which candidates affiliated with a superordinate group use similar or disparate words. This is a question similar to how the average grade of membership of a category explains the hazard of category spanning (Kovcs and Hannan 2010), but differs in that the 'category' of interest is a social entity. Comparing social groups that have similar longevity, membership, and activity, but with different degrees of conformity in features, could recover the effects of group cohesion on turnover, network vulnerability, and individual member evaluation. A similar consideration—for settings in which candidates enter evaluative contexts multiple times—is a candidate's consistency in features. In their study of tweet propagation, Tan, Lee, and Pang (2014) find that conformity to the conventions of Twitter, as well as to one's previous pattern of activity, is positively related with message propagation. The Twitter authors in this study were necessarily established members of the Twitter 'market' (e.g., had at least 5,000 followers), but this invites the question of generalizability to more emergent actors. My findings in Chapter 3 suggest that when markets are nascent, conformity can be especially hazardous, but the findings of Tan, Lee, and Pang (2014) suggest conformity to a community may make conformity tolerable.

The thesis of chapter 3 particularly could be extended by the recent insights of Bowers (2015), who argued that categories research should account for the limited and varied consideration sets of members of an audience. Bowers finds that stock analysts interpret a given firm's performance information differently because of differing consideration sets, but that this reliance on relative comparison depends on the typicality of the firm within the category: when firms are atypical, relative comparison is less potent in evaluations.

Lastly, while my research did not consider the transference of performance (quality) signals into evaluation, the focus on localized comparison is relevant. Presently, measures of coherence and typicality/unconventionality are calculated using all available data simultaneously. This approach assumes either that every audience member views all candidates and incorporates this information, or that it is a good approximation of individual audience members' localized assessments. Indeed, there may be both within- and between-audience-member variation in assessments of coherence and typicality, depending on experience in viewing candidates in that category, or through more substantive relations or interactions with candidates in that category. The preceding analysis presents an initial foray into the matters of coherence and

Chapter 5 Conclusion

typicality. Subsequent efforts could examine more localized variation in how between- and within-category comparisons are made.

REFERENCES

Agrawal, Ajay K., Christian Catalini, and Avi Goldfarb. 2011. "The Geography of Crowdfunding."

Albert, Stuart, and David A. Whetten. 1985. "Organizational identity." *Research in Organizational Behavior* 7:263–295.

Ashforth, B. E., and R. H. Humphrey. 1997. "The ubiquity and potency of labeling in organizations." *Organization Science* 8 (1): 43–58.

Barroso, Alicia, Marco S. Giarratana, Samira Reis, and Olav Sorenson. 2014. "Crowding, satiation, and saturation: The days of television series' lives." *Strategic Management Journal:*n/a–n/a.

Bazerman, Max H., and Don A. Moore. 2013. *Judgment in managerial decision making.* 8th. vii, 277 p. New York: Wiley.

Berger, Peter L., and Thomas Luckmann. 1967. "Society as objective reality." Chap. 2 in *The Social Construction of Reality: A Treatise in the Sociology of Knowledge.* Garden City, NY: Doubleday / Anchor Books.

Bowers, A. 2015. "Relative Comparison and Category Membership: The Case of Equity Analysts." *Organization Science* 26 (2): 571–583.

Burtch, G., A. Ghose, and S. Wattal. 2013. "An Empirical Examination of the Antecedents and Consequences of Contribution Patterns in Crowd-Funded Markets." *Information Systems Research* 24 (3): 499–519.

Carroll, G. R., and A. Swaminathan. 2000. "Why the microbrewery movement? Organizational dynamics of resource partitioning in the US brewing industry." *American Journal of Sociology* 106 (3): 715–762.

Carruthers, Bruce G., and Arthur L. Stinchcombe. 1999. "The social structure of liquidity: Flexibility, markets, and states." *Theory and Society* 28 (3): 353–382.

De Bondt, Werner F. M., and Richard Thaler. 1985. "Does the Stock Market Overreact?" *The Journal of Finance* 40 (3): 793–805.

DiMaggio, Paul. 1987. "Classification in Art." *American Sociological Review* 52 (4): 440–455.

Gao, Qiang, and Mingfeng Lin. 2013. "Linguistic features and peer-to-peer loan quality: A machine learning approach."

Hagiu, Andrei. 2009. "Two-sided platforms: Product variety and pricing structures." *Journal of Economics and Management Strategy* 18 (4): 1011–1043.

Hannan, Michael T., László Pólos, and Glenn Carroll. 2007. *Logics of organization theory : audiences, codes, and ecologies.* xiii, 364 p. Princeton, N.J.: Princeton University Press.

Herzenstein, Michal, Scott Sonenshein, and Utpal M. Dholakia. 2011. "Tell Me a Good Story and I May Lend You Money: The Role of Narratives in Peer-to-Peer Lending Decisions." *Journal of Marketing Research* 48:S138–S149.

Hsu, Greta. 2006. "Jacks of all trades and masters of none: Audiences' reactions to spanning genres in feature film marketing." *Administrative Science Quarterly* 51:420–450.

Hsu, Greta, and Michael T. Hannan. 2005. "Identities, genres, and organizational forms." *Organization Science* 16 (5): 474–490.

Hsu, Greta, Michael T. Hannan, and László Pólos. 2011. "Typecasting, Legitimation, and Form Emergence: A Formal Theory." *Sociological Theory* 29 (2): 97–123.

Hsu, Greta, Özgecan Koçak, and Michael T. Hannan. 2009. "Multiple category memberships in markets: An integrative theory and two empirical tests." *American Sociological Review* 74 (1): 150–169.

Iacus, Stefano M., Gary King, and Giuseppe Porro. 2012. "Causal Inference without Balance Checking: Coarsened Exact Matching." *Political Analysis* 20 (1): 1–24.

Iyer, R., A. I. Khwaja, F. P. Luttmer, and K. Shue. 2009. "Screening in new credit markets: Can individual lenders infer borrower creditworthiness in peer-to-peer lending?"

Jensen, Michael, and Bo Kyung Kim. 2014. "Great, Madama Butterfly Again! How Robust Market Identity Shapes Opera Repertoires." *Organization Science* 25 (1): 109–126.

Kleinbaum, Adam M. 2012. "Organizational misfits and the origins of brokerage in intrafirm networks." *Administrative Science Quarterly* 57 (3): 407–452.

Kocak, Ozgecan, Michael T. Hannan, and Greta Hsu. 2014. "Emergence of Market Orders: Audience Interaction and Vanguard Influence." *Organization Studies* 35 (5): 765–790.

Kovacs, Balazs, and Rebeka Johnson. 2014. “Contrasting alternative explanations for the consequences of category spanning: A study of restaurant reviews and menus in San Francisco.” *Strategic Organization* 12 (1): 7–37.

Kovcs, B., and M. T. Hannan. 2010. “The consequences of category spanning depend on contrast.” In *Categories in Markets: Origins and Evolution. Research in the Sociology of Organizations,* ed. by G. Hsu, G. Negro, and O. Kocak, 31:175–202. New York: Emerald Group Publishing.

———. 2013. “Category spanning, interpretability, and appeal.”

Lamont, M. 2012. “Toward a Comparative Sociology of Valuation and Evaluation.” *Annual Review of Sociology* 38:201–221.

Lee, Peggy M. 2001. “What’s in a name. com?: The effects of . comname changes on stock prices and trading activity.” *Strategic Management Journal* 22 (8): 793–804.

Leung, M. D., and A. J. Sharkey. 2014. “Out of Sight, Out of Mind? Evidence of Perceptual Factors in the Multiple-Category Discount.” *Organization Science* 25 (1): 171–184.

Leung, Ming D. 2014. “Dilettante or Renaissance Person? How the Order of Job Experiences Affects Hiring in an External Labor Market.” *American Sociological Review* 79 (1): 136–158.

Manning, Christopher D, and Hinrich Schütze. 1999. *Foundations of statistical natural language processing.* Vol. 999. MIT Press.

Markowitz, Harry. 1952. “Portfolio selection.” *The Journal of Finance* 7 (1): 77–91.

Merluzzi, Jennifer, and Damon Phillips. 2014. “The Specialist Discount: Negative returns for MBAs with focused profiles in investment banking.”

Mollick, Ethan R. 2014. “The Dynamics of Crowdfunding: An Exploratory Study.” *Journal of Business Venturing* 29 (1): 1–16.

Mullainathan, S., and R. H. Tahler. 2001. “Behavioral Economics.” In *International Encyclopedia of the Social and Behavioral Sciences,* ed. by Neil J. Smelser and Paul B. Baltes, 1094–1100. Pergamon, Oxford.

Murphy, Gregory L, and Brian H Ross. 2005. “The two faces of typicality in category-based induction.” *Cognition* 95 (2): 175–200.

Negro, Giacomo, and Ming D. Leung. 2013. “”Actual” and Perceptual Effects of Category Spanning.” *Organization Science* 24 (3): 684–696.

Odean, T. 1998. “Are investors reluctant to realize their losses?” *Journal of Finance* 53 (5): 1775–1798.

Ody-Brasier, Amandine, and Freek Vermeulen. 2014. “The Price You Pay: Price-setting as a Response to Norm Violations in the Market for Champagne Grapes.” *Administrative Science Quarterly* 59 (1): 109–144.

Padgett, J. F., and C. K. Ansell. 1993. “Robust action and the rise of the Medici, 1400-1434.” *American Journal of Sociology* 98 (6): 1259–1319.

Phillips, Damon J., Catherine J. Turco, and Ezra W. Zuckerman. 2013. “Betrayal as Market Barrier: Identity-Based Limits to Diversification among High-Status Corporate Law Firms.” *American Journal of Sociology* 118 (4): 1023–1054.

Podolny, JoelM. 2001. “Networks as the Pipes and Prisms of the Market.” *The American Journal of Sociology* 107 (1): 33–60.

Pontikes, Elizabeth G. 2012. “Two sides of the same coin: How ambiguous classification affects multiple audiences’ evaluations.” *Administrative Science Quarterly* forthcoming.

Pontikes, Elizabeth, and Michael T. Hannan. 2014. “An Ecology of Social Categories.” *Sociological Science* 1 (August): 311–343.

Prosper.com. 2012.

Rao, H., P. Monin, and R. Durand. 2003. “Institutional change in Toque Ville: Nouvelle cuisine as an identity movement in French gastronomy.” *American Journal of Sociology* 108 (4): 795–843.

———. 2005. “Border crossing: Bricolage and the erosion of categorical boundaries in French gastronomy.” *American Sociological Review* 70 (6): 968–991.

Rosa, J. A., J. F. Porac, J. Runser-Spanjol, and M. S. Saxon. 1999. “Sociocognitive dynamics in a product market.” *Journal of Marketing* 63:64–77.

Ruef, Martin, and Kelly Patterson. 2009. “Credit and classfication: The impact of industry boundaries in 19th century America.” *Administrative Science Quarterly* 54:486–520.

Salton, Gerard, and Christopher Buckley. 1988. “Term-weighting approaches in automatic text retrieval.” *Inf. Process. Manage.* 24 (5): 513–523.

SEC. 2008. *In the matter of Prosper Marketplace, Inc.*

Smith, E. B. 2011. "Identities as Lenses: How Organizational Identity Affects Audiences' Evaluation of Organizational Performance." *Administrative Science Quarterly* 56 (1): 61–94.

Spence, Michael. 1973. "Job market signaling." *Quarterly Journal of Economics* 87 (3): 355–374.

Tan, Chenhao, Lillian Lee, and Bo Pang. 2014. "The effect of wording on message propagation: Topic-and author-controlled natural experiments on Twitter." *arXiv preprint arXiv:1405.1438.*

Tucker, Catherine. 2008. "Identifying Formal and Informal Influence in Technology Adoption with Network Externalities." *Management Science* 54 (12): 2024–2038.

Waguespack, D. M., and O. Sorenson. 2011. "The Ratings Game: Asymmetry in Classification." *Organization Science* 22 (3): 541–553.

White, H. C. 1981. "Where do markets come from?" *American Journal of Sociology* 87 (3): 517–547.

Wooldridge, Jeffrey M. 2010. *Econometric analysis of cross section and panel data.* 2nd. xxvii, 1064 p. Cambridge, Mass.: MIT Press.

Zerubavel, Eviatar. 1997. *Social Mindscapes: An Invitation to Cognitive Sociology.* Cambridge, Massachusetts: Harvard University Press.

Zhao, Wei, and Xueguang Zhou. 2011. "Status Inconsistency and Product Valuation in the California Wine Market." *Organization Science* 22 (6): 1435–1448.

Zuckerman, Ezra W. 1999. "The categorical imperative: Securities analysts and the illegitimacy discount." *American Journal of Sociology* 104 (5): 1398–1438.

———. 2000. "Focusing the corporate product: Securities analysts and de-diversification." *Administrative Science Quarterly* 45:591–619.

Zuckerman, Ezra W., Tai-Young Kim, Kalinda Ukanwa, and James von Rittmann. 2003. "Robust identities or nonentities? Typecasting in the feature-film labor market." *American Journal of Sociology* 108 (5): 1018–1074.

Table A.1: Changes to Prosper Website

Date	Change
April 2, 2007	New web design and real-time API
May 2, 2007	Lenders can create and share a list of watched listings.
June 6, 2007	Borrowers enabled to preview listings, save drafts of listings,
(Begin analysis period)	and use templates
Sept. 12, 2007	Group leader rewards and group categories discontinued
October 30, 2007	Maximum borrower rate changes from 30% to 36% in some states
	Lenders can choose from up to four automated portfolio plans
December 5, 2007	**Borrowers select a category describing purpose of their loan**
(Categories adopted)	Debt consolidation, home improvement, business, personal,
	education, auto, and other
	Lenders can search and bid by category
	New Prosper blog
	Prosper discussion forums now tied to the main Prosper web site
January 4, 2008	Fees: Borrower origination fees change to 1% for AA, 2% for A-B,
	and 3% for C-HR.
February 23, 2008	Lenders can now create, share, and copy portfolio plans
	API users can now place bids via the API
April 15, 2008	Maximum borrower rate changes to 36% in almost every state
(End analysis period)	Fees: Lender annual servicing fees change to 1% for AA,
	matching that for all grades
	Listing duration changed to 7 days
May 19, 2008	Institutional lenders welcomed on Prosper

Note: Description of major changes to the Prosper website before, during, and immediately after the analysis window. Information was obtained from Prosper blog postings.

Figure A.1: Illustration of Purpose Prediction

	Listing	Purpose	due	due balance	due tuition	dues	duplex	…
	1	**Auto**	1	0	0	0	0	…
Unlabeled (N = 81,477	2	**Education**	1	0	1	0	0	…
	3	**Personal**	0	0	0	0	0	…
	…	…	…	…	…	…	…	…
	81478	Education	1	0	1	1	0	…
Labeled (N = 85,505)	81479	Auto	1	1	0	0	0	…
	81480	Home	0	0	0	0	1	…
	…	…	…	…	…	…	…	…

Note: Purpose classifiers were modeled on observed term distributions in the labeled period. These classifiers were used to predict purpose labels of pre-label listings.

Figure A.2: Summary of Cross-validation Process

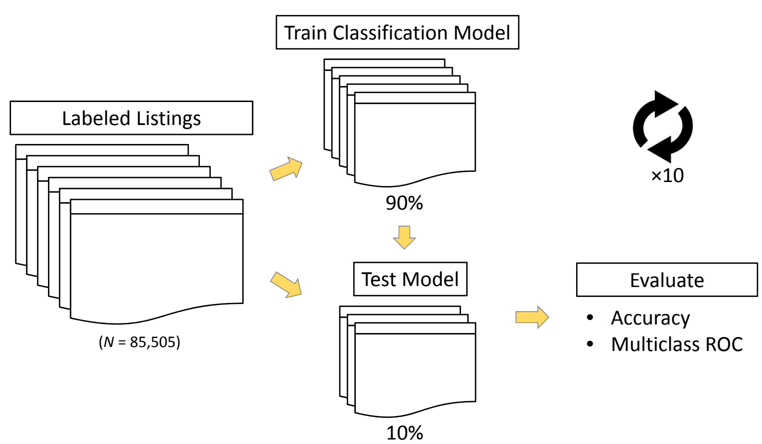

Note: Classification models are evaluated through ten-fold cross-validation. In each step, a random 10% of the labeled data is reserved for testing, and a model is trained using the remaining 90%. Since the purpose labels of the testing data are known, it is possible to assess the performance of a given model. Models that yield high average performance with few terms are preferred.