UC Berkeley
Working Paper Series

Title

An idiosyncrasy credit or a generalist discount? Conditional advantages to working broadly in a virtual labor market

Permalink

https://escholarship.org/uc/item/2m57x2tm

Authors

Leung, Ming D.
Ng, Weiyi

Publication Date

2014-03-21

An idiosyncrasy credit or a generalist discount? Conditional advantages to working broadly in a virtual labor market

Ming D. Leung and Weiyi Ng

irle.berkeley.edu/workingpapers

An idiosyncrasy credit or a generalist discount?
Conditional advantages to working broadly in a virtual labor market*

Ming D. Leung
Haas School of Business
University of California, Berkeley
2220 Piedmont Avenue
Berkeley, CA 94720
1-510-642-4952
mingdleung@haas.berkeley.edu

Weiyi Ng
Haas School of Business
University of California, Berkeley
2220 Piedmont Avenue
Berkeley, CA 94720
weiyi_ng@haas.berkeley.edu

March 21, 2014

ABSTRACT

This article challenges the widespread notion that spanning recognized categories, or generalism, is disadvantageous. In markets, buyers face the fundamental challenge of gaining insight into the underlying ability of sellers. Economic sociologists rely on theories of categorization to explain how participants navigate this assessment. Social actors who span multiple categories have been found to be disadvantaged. We do not believe this is always the case and seek to identify conditions under which generalism by job seekers is valued. We start with the contention that spanning is an ambiguous indicator of quality, and suggest how it is perceived is dependent on other, salient indicators. We investigate two – relevance and reputation – which are often confounded with generalism. We propose that the same spanning behavior can be advantageous or disadvantageous depending on these other indicators. We hypothesize that those who evince the necessary attributes an employer is looking for by being relevant are advantaged when they also span disparate areas because employers value broader experiences. Spanning should also be advantageous for highly reputable applicants because it is more consistent with the belief they possess superior skills. In short, we identify *idiosyncrasy credits*. We test and find support for these hypotheses in the context of a virtual labor market for freelancers – Elance.com.

KEYWORDS

Categories, Online Labor Markets, Freelancing

In markets, buyers face the fundamental challenge of gaining insight into the underlying ability of sellers. Flourishing work by economic sociologists has relied on theories of categorization to explain how market participants navigate this assessment (Hannan et al 2007, Hsu 2006). They begin with the observation that continuously varying items are often organized into discrete and identifiable sets (Rosch 1978); thereby easing the task of comprehension and comparison (Zerbabuval 1997, Hsu and Hannan 2005). In doing so, those social actors who are more apt to display features or characteristics of a single socially defined category, and not others, appear more categorically coherent (Zuckerman et al 2003). They are therefore advantaged in the marketplace.

Theorists who proceed along this line of investigation speculate that these categorical assessments are useful because in the absence of more quantifiable measures of ability or quality, demonstrations of past categorically distinct work serve as a proxy of this difficult to detect characteristic (Zuckerman et al 2003). Because buyers tend to search for sellers within distinct categorical confines (Rao et al 2003, Smith 2010, Pontikes 2012), products or producers who span, or attempt to combine elements from, several established circumscribed boundaries are disadvantaged because they may present identities which are hard to compare (Zuckerman 1999), are assumed to be of lower quality (Negro and Leung 2013), or transgress socially sanctioned boundaries (Phillips and Zuckerman 2001; Rao et al 2003). Evidence for this empirical regularity has been identified across a broad range of markets such as feature films and eBay sellers (Hsu et al 2009), Italian Barolo and Barbaresco wines (Negro et al 2011), consumer debt (Leung and Sharkey 2013), restaurants (Kovacs and Hannan 2013), and feature film actors (Zuckerman et al 2003). For a review see Hannan (2010).

Yet this belief has been challenged by the more recent literature which reveals that categorical spanning effects continue to operate in the presence of visible quality assessments (Negro and Leung 2013) or that a more nuanced and interdependent relationship persists between the status of an organization and categorical divisions such as industry groups (Sharkey 2013). This belief also contradicts reputational theorists who propose that visible indicators of better or worse performance have much more diffuse effects, such as your ability to survive (Rao 1994) or long term financial performance (Roberts and Dowling 2002). Precisely because reputation leaks across different domains, generalism should be advantageous to some. Indeed more recent theorizing has called for a tighter link between reputation and categories (Jensen et al 2012). For a recent review see Barnett and Pollack (2012).

The stance that generalism is universally disadvantaged limits the usefulness of categorization theory when it is applied to labor markets which are increasingly characterized by shorter term relationships between employers and employees making spanning more common (Arthur and Rousseau 1996). Movement between varying jobs, whether to gain additional skills (O'Mahony and Bechky 2006, Bidwell and Briscoe 2010) or because of the temporary nature of your position (Kunda and Barley 1997, Osnowitz 2010) increasingly requires employers to ascertain suitability of an employee with a set of diverse experiences which may span disparate domains. The rise of external labor markets (Polivka et al. 2000) and non-standard work arrangements (Kalleberg 2000), where workers are no longer expected to develop a career within a single organization only highlights these issues and leads us to ask: Under what conditions does working broadly across socially defined categorical boundaries advantageous to job seekers?

We extend on burgeoning work that attempts to identify the conditions under which spanning, or generalism, may be less penalized or even advantaged (Kleinbaum 2012, Merluzzi and Phillips 2014). We suggest one way extant work can benefit is by disambiguating previous conceptions of generalism into two additional independent concepts – that of relevance, reputation. Because spanning disparate areas, or generalism, can either evince superior skills or a lack of any for a job seeker (Leung 2014, Ferguson and Hasan 2013, Zuckerman et al 2003), it is a fuzzy signal for underlying ability. We therefore expect to identify contextual factors that affect whether spanning, or generalism, is perceived as more or less advantageous. We propose two.

First, in the presence of more quantifiable measures of overall quality, spanning for those who are also able to demonstrate a better reputation incurs additional advantages. This is because employers will code spanning disparate areas as congruent with evincing broader skill or ability only in the presence of corroborating signals such as a better reputation. Second, indicators of relevance of past experiences are most important in seeking an employee. However, beyond this 1st order expectation, evidence of a broader portfolio of past experiences should make a candidate more attractive over one who is narrowly experienced. This is because employers often hire an employee not merely the particular task at hand, in which case broader experiences may be valuable later. Also, the abilities necessary to perform a particular task may not always be easily specified, in which case evidence of broader ability may be more useful.

Our setting is a virtual market for freelance contract labor, www.elance.com, where employers of contract labor seek and hire temporary employees willing to work on a contract basis – all completed online. This setting presents a particularly suitable site for investigating how employers view and choose employees. First, as a result of the large volume of business, all

jobs on this website are listed in a well-defined category. Furthermore, because all the transactions are recorded by the site, a freelancer's past jobs on the website are visible for an employer to view as well as past measures of the freelancer's reputation, as reflected in the feedback scores they receive on past jobs. Empirical advantages aside, the virtual nature of this setting is also notable as an increasing proportion of the labor market may continue to move towards these virtually mediated labor markets.

THEORY

Categories in Labor Markets

Categories are socially recognized groupings of like-objects, and serve to circumscribe similar items and exclude dissimilar ones (Rosch 1978; Hannan et al. 2007). This parallels our universal inclination to partition an assortment of complex items or objects into manageable and publicly understood classificatory clusters (Fiske and Taylor 1991; Zerubavel 1997). Categories are useful because they invoke assumptions about what being a member of a category entails. This reduces the cognitive effort required to sort, search, and compare these items. For example, feature films are usually identified by recognized and socially accepted genres. Moviegoers know that an action movie will usually include plot elements such as car chases and stunts. So when a film is billed as an action movie, it will elicit expectations by filmgoers of these plot elements and not others. Categories therefore facilitate transaction in markets.

Given this, buyers look within a particular socially recognized category when searching for suitable transaction partners. So purchasers of software with a requirement for an enterprise software system will search for software companies that produce such a product (Pontikes 2012). In doing so, software makers who produce several different types of software, perhaps the disparate categories of enterprise software and customer relationship management software are

disadvantaged. This is because buyers rely on these categorical distinctions when assessing the quality of a seller (Rao et al 2005). Products, which combine elements of disparate categories present incoherent identities. They are harder to understand, evaluate, and value (Zuckerman 1999) and therefore ignored (for review, see Hannan 2010). This is why buyers prefer sellers who are able to present a single, distinct identity (Zuckerman et al 2003).

In labor markets, categories ease comprehensibility for employers because jobs labeled in one category are understood to be similar to each other and dissimilar to those in another. Those seeking work as well as those seeking workers hew to these recognized category groupings. Categories in this sense evoke a set of associated skills and training particularized to those domains. For example, the market for graduates of PhD programs in business are distinguished into recognized categories such as marketing, operations, organizational behavior, or strategy – representing the departments from where they studied. Junior faculty searches proceed within these boundaries – with schools seeking students who are educated within these domains presumably because they are educated with skills specific to those areas and not others.

The seminal example of how categories operate in labor markets is the phenomenon of typecasting (Faulkner 1983). *A priori* assessment of a film actor's ability is difficult. Therefore, because acting in a dramatic film is generally seen as different from acting in a comedic one, previous demonstrations of success in a dramatic role act as *prima facie* evidence of an actor's skill in that genre and suggests an inability to act in comedic ones. So actors who focus their experiences in a single movie genre are more likely to be subsequently hired for an identical role (Zuckerman et al. 2003), but are quickly screened out from consideration for different roles. This is why specialization in labor markets, as conceptualized by concentration in a single category, is

theorized to signal greater skill earlier in ones career than generalists who compile multiple experiences (Ferguson and Hasan 2013).

However, our theoretical apparatus is unable to explain how past experiences may vary in relevance to the role under consideration (particularly in cases where experiences are not perfectly matched to needs) and instead merely leaves us with the unsatisfying belief that not having performed in a specific role unequivocally disqualifies you. How then can we understand the chances a person will be able to switch jobs? Some nascent work by Kovacs and Hannan (2013) has begun to explore this particular conundrum by suggesting that more prevalent combinations of categories are less likely to elicit censure when combined. They contend that repeated observations of a particular paring of categories necessarily leads to less cognitive difficulty in understanding that combination. Yet their theory falls short of helping us develop an understanding as to how distant a singular category instance is from another.

Another shortcoming of this extant theory is the fact that it cannot distinguish between those employees who present a myriad of different past experiences because they are broadly skilled or because they are unskilled at any one. Instead, past work has focused on situations where these category spanners are *assumed* to be of lower ability. In particular, these categorical pressures are hypothesized to operate in the absence of more easily observable measures of quality (Zuckerman et al 2003). In these situations, category spanning is more likely to be disadvantageous because it is difficult to distinguish between multi-talented and untalented employees. However, more recent work has demonstrated that externally valid indicators of quality operate independently of demonstrations of spanning (Negro and Leung 2013), in a narrow band of ability (Ferguson and Hasan 2013), or may be an advantage when the modal candidate is highly specialized (Merluzzi and Phillips 2014) and leads us to question under what

conditions a broad work background can be advantageous. We turn now to contract labor and an online market for freelancing.

A Virtual Market for Contract Labor

Variously described as freelancers, contract labor, or independent contractors (Barley and Kunda 2004; Osnowitz 2010), an increasing proportion of the labor market has been leaving permanent employment for self-employment representing an extreme instantiation of a "boundaryless" career (Arthur and Rousseau 1996). Falling under the umbrella term of "non-standard employment," some estimates suggest that contingent work may comprise 29% of all jobs (Kalleberg 2000). While some examinations have highlighted the fact that these are mostly "bad jobs" and represent poor employment opportunities in terms of pay, benefits, or advancement (Kalleberg 2011; Moss et al. 2000), non-standard workers actually encompass a surprisingly heterogeneous group (Kalleberg 2000).

We focus on the recent trend for the skilled and technical labor force to increasingly embrace a contract-based arrangement, providing skilled labor on a project basis and often commanding a pay premium over their similarly tasked full-time counterparts (Polivka et al. 2000). The domains of writing and editing or programming and engineering fields, for example, have well-institutionalized contractor and client bases (Osnowitz 2010). Structural as well as individual changes have contributed to this shift leading people to choose to enter the contract labor force, including the speed of technological advancements, corporate restructuring, and worker preferences (for review see Ganz et al. 2000). By some reports, this skilled portion of the temporary workforce has grown dramatically in the past couple decades to comprise approximately a quarter of all non-standard contingent workers (Barley and Kunda 2004) – with

proportions estimated to be greater in particular domains, such as programming or multimedia design.

Following a similar trend in tangible product markets, labor market hiring has begun to move online. Witness the recent proliferation of online "crowd-sourced" labor markets that mediate employers and employees, such as Monster or Career Builder; or ones that specialize in temporary contract labor, such as oDesk or Elance. Elance.com, the site under study here, is the oldest firm in this arena and acts as a virtual marketplace where buyers of a broad range of business services find and hire independent professionals on a contract basis to work remotely. Self-styled "the future of hiring," Elance provides a venue for would-be employers of freelancers to solicit the skills and services of an online freelancing workforce. This is achieved by posting a "job listing." Potential employers (henceforth *buyers*) provide a job description, detail requirements and set a timeframe for both the job award and deliverables. See Figure 1 for a sample job listing. Freelancers (henceforth *sellers*) proceed to bid on these jobs. After reviewing the past job history and performance of bidding sellers, the buyer then establishes a winner; the job is thus awarded. All deliverables are accomplished online. There are currently over 100,000 jobs posted each month and over 2 million freelancers located worldwide on the website. Since founding in late 1999, cumulative transactions worth over $900 Million have been completed on the website.

[insert Figure 1 about here]

As a necessity, given the volume of transactions, Elance.com job listings are organized into job categories that represent conventionally recognized divisions of works. Each job is therefore listed in one, and only one category. Examples include "Website Programming", "Administrative Assistance", "Translation Services", and "Logo Design." (See Appendix for

complete list of the 168 job categories.) Categories cover a range of services from graphic design to web programming. Not all categories are related to web-based services; for instance there are categories that cater to creative writing, financial and legal services. While these categories distinctly partition jobs on the website, similarities between these categories may vary (Ruef 2000; Kovacs and Hannan 2013). Tasks circumscribed in one job category can be perceived as being more or less similar or related to those in another. For example, "Menu Design" and "Logo Design" may be seen as being proximate but not identical job categories as opposed to the categories of "Website Programming" and "Business Plan Writing" perhaps representing more distant job categories.

In offline markets, staffing agents often mediate the relationship between a contractor and their employers and assist employers in identifying the appropriate relevant skills to best fit a job (Barley and Kunda 2004). To the extent that job categories lead to beliefs by employers of a particular candidate's suitability to a task, then a candidate who is able to craft a resume that demonstrates relevant experience and progress will be advantaged in securing subsequent work. So contractors and their brokers often modify and frame a freelancer's past experiences to convey relevance (Osnowitz 2010) leaving truly valid and uncensored information a "casualty of trade" (Barley and Knuda 2004: 133).

Here on Elance, to assist an employer in identifying appropriate freelancers, a freelancer's past completed jobs on the website is visible, organized chronologically, and identified by their job category. (See Figure 2 for an example of a freelancer's profile.) Note that a freelancer's past experiences are listing chronologically, with the four most recent jobs they have completed on the front page of their profile. Elance has institutionalized the presentation of

11

a freelancer's past work experiences by standardizing what gets displayed. A freelancer's list of jobs then poses as their online career which is displayed to employers to aid their hiring choices.

[insert Figure 2 about here]

HYPOTHESES

Category Spanning versus Relevance in Hiring

In choosing who to hire, an employer will be evaluating the underlying, difficult to infer, ability of a freelancer. Barring more obvious indicators of skill for the moment, a potentially salient indicator would be a freelancer's past job history. Working in a particular job category should suggest at least an ability to adequately perform that job again. Therefore, we should expect a freelancer with a job history focused in a particular job category to be more likely to be hired again in that category (Zuckerman et al 2003). However, it will be more difficult to infer the ability of a freelancer with more disparate past jobs because a broad background could be inferred in two, opposing, ways. On one hand, working broadly could indicate that a freelancer is multi-talented. So successfully completing a Website Design job and a Website Programming one should serve to indicate broader ability. However, another interpretation of such a history is that the freelancer is poorly skilled and therefore forced to move from failure to failure (Gibbons and Katz 1991). So working across several different job categories makes it difficult for an employer to be confident in your ability. We believe employers will generally shy away from such job seekers. Therefore, as past research has suggested, we expect as a 1st order prediction:

Hypothesis 1a: Greater spanning decreases chances of winning a job.

Because categories are socially constructed distinctions among potentially continuously varying characteristics of objects or items, an artificial division necessarily cleaves into two groups items which may share some characteristics with one another. For example, the genres of

12

action movies and adventure movies, while two distinct categories of films, likely share many characteristics. So the two categories that emerge are necessarily distinct from one another, but they may be more related to one another than a third category which resulted from a division among items with less characteristics in common. Following the above example, we would suggest that action and adventure movies share more characteristics than either do with romantic movies. Empirical regularities suggest this as well. Zuckerman and colleagues (2003) examination of movie actors identified how the cleave between comedy and drama genres was particularly sharp, thereby noting that those actors who began their careers in comedy, were much less likely to be able to move into drama (or any other genre), but this effect was not necessarily as strong for other movie genres. Presumably, this reflects beliefs by casting agents as to the greater differences between the skills required to act in a comedy from those required to succeed in other roles.

While the jobs on Elance are posted in distinct categories – we note two important characteristics. First, the job categories do not explicitly partition tasks into skill based groups. The jobs posted in the category of Website Programming, for example, may requires skills such as programming C#, or PhP, or Flash, all of them, some combination, or none. While there certainly are general distinctions of skill that are related to the categories (i.e. website programming and legal advice draw on distinct skills) there is overlap between the job categories. This leads to the second point.

Second, the categories themselves vary in how related or similar they are to one another because some jobs will share characteristics with one another while others may not. For example, the ability a freelancer possesses in order to complete a Menu Design job may be more relevant to their future success at completing a Logo Design job than a Database Programing job. To the

extent that a freelancer has experiences in a more relevant job category, then their experiences should serve to make them a more attractive candidate to an employer than a freelancer who has worked on less relevant jobs. If this were the case, we'd expect to see the following:

Hypothesis 1b: Greater relevance of past experiences to the bidding category increases chances of winning the job

Note that both relevance and spanning can both be indicators that an employer will likely rely on to infer the ability a freelancer possesses for their specific task. Employers use spanning as a proxy for the potential abilities a freelancer may possess. While relevance is a direct measure of the appropriateness of a freelancer's past jobs to the one they are bidding on. We therefore do not expect that these indicators are equally informative and we should see that relevance of past experiences overrides the discount received from spanning. Past research has confounded these two aspects by proposing that spanning serves as a proxy for underlying ability in the absence of more easily verifiable measures (Zuckerman et al 2003). Here, we suggest that the relevance of a freelancer's past experiences are these more easily verifiable measures and serve as a more pertinent indicator of ability (Merluzzi and Phillips 2014). In these instances, we should observe our measure of relevance to mediate the negative effect of spanning. Formally:

Hypothesis 1c: Relevance of past experiences mediates the relationship between spanning and winning a job.

Relevance effects on Spanning

Spanning has been demonstrated to continue to operate as a signal beyond a mere proxy for underlying quality or ability. Digressions could be seen as flaunting tradition because of the oppositional nature of the categories (Negro and Leung 2013), as demonstrated by the ratings penalties suffers by Italian winemakers who combined both Traditional and Modern winestyles.

Spanning could also suggest flexibility which was why more ambiguously classified software companies were more likely to secure funding from venture capitalists (Pontikes 2012). In labor markets, working progressively across a range of jobs may suggest a more committed employee (Leung 2014).

Decision theorists subscribe to a two-stage process in choosing among alternatives (Shocker et al 1991, Urban et al 1993) whereby the first stage of the process involves the winnowing down of available options to a shorter list. The second stage then involves choosing among this narrowed down list. Economic sociologists have embraced this notion and identified how the first stage decision criteria involve evaluations of a candidate's suitability for the job. Perhaps through demonstrations of categorical coherence (Zuckerman 1999) or proven ability in a particular area (Zuckerman et al 2003) – this first stage merely qualifies a candidate to be considered further. Paralleling this, because the average number of freelancers applying for a job posting on Elance is approximately twenty, the website has instituted a feature which allows employers to quickly note which bids to 'shortlist.' This allows employers to then sort these preferred bids and quickly browse and screen in freelancers who are obviously qualified, for example, by how relevant their past experiences are.

The second stage of the hiring decision then hinges on how the remaining freelancers can differentiate themselves from one another (Zuckerman 1999). Here, we suggest that freelancers who have accomplished a variety of jobs in the past will be seen as particularly attractive because the decision to hire extends beyond merely the skills specific to the job at hand. About 50% of employers on Elance, for example, are looking to hire freelancers for repeated jobs (Elance 2012). A parallel can be seen for hiring into a firm, which entails not only the entry level

position but possible subsequent promotions which would lead an employee to take on responsibilities which are unlike those at entry level.

Another reason this should be the case is that an employer may not know the precise skills or abilities will be necessary to complete a task. While past categorical experiences in a similar domain should serve as evidence of ability in that area, it assumes the task at hand precisely matches a previous one. Certainly, in product markets, where items may be commoditized, this is less of a concern. However, in labor markets it would be difficult to sustain the belief that one job is precisely identical to another one merely because it's labeled in the same way particularly because skills required may vary within a job category. In this case, demonstrations by freelancers of working across a broader range of domains increases the chances they will be able to bring to bear a broader skill set to any task. In sum, relevant freelancers will derive additional benefits from spanning. To the extent this social process is operating, we expect to see that as relevance to the focal category under consideration increases, freelancers spanning multiple job categories should become more attractive. Formally:

Hypothesis 2: Spanning is more (less) advantageous to winning a job as relevance to the focal bidding category increases (decreases).

Reputation effects on Spanning

Reputation is a public indicator of the quality of a social actor and is derived from records of their past actions. It acts as a signal of their future abilities (Rao 1994; Fombrun 1996; Podolny and Phillips 1996). It acts as a signal (Spence 1973) to would be buyers as to what they may expect in terms of quality. A good reputation therefore should be a positive asset to a social actor's future prospects. Rao (1994) conceptualizes reputation of nascent auto manufacturers as the legitimacy received from winning certification contests. He demonstrates how this improves

the organizations chances of survival. Roberts and Dowling (2002) identify how a good reputation can help an organization sustain superior profits over time. In online environment, where uncertainty of a transaction partner may be high due to the remote nature of the market, good reputations helped e-bay sellers increase the price they can charge (Resnick et al 2006) and close auctions more easily (Resnick and Zeckhauser 2002).

Reputations are constructed from several disparate domains and are often portrayed as diffuse in nature. Scholars of reputation have generally agreed that reputation is best defined in terms of 'overall assessments' or even 'generalized favorability' (Lange et al 2011). Because an actor's reputation is a perception of their quality, it should influence evaluations across diverse fields. For example, those with better reputations are more likely to associate with others who are considered to have a high reputation (Stuart et al 1999; Podolny and Phillips 1996). Better reputations also may induce potential buyers to pay more for goods or services because of the expectation of higher quality (Shapiro 1983). We therefore expect that a better reputation will increase a social actor's appeal.

Elance solicits and records a feedback score for each completed job on the website. After an employer and freelancer complete a job, the employer has the opportunity to leave a feedback score, which is a star rating on a scale from 1 to 5 stars reflecting how the freelancer performed on the job. Feedback is specifically solicited along multiple dimensions, which not only capture how skilled the freelancer was at this particular task, but broader indicators of ability. The rated dimensions include how well the freelancer ranked on Quality, Expertise, Cost, Schedule, Responsiveness, and Professionalism. These scores are all averaged into a single 'star' rating which is prominently displayed besides each freelancer's profile. The reputation system resembles one that makes use of *assessment* signals (Donath 1998) – these signals of quality are

17

not self-reports and have to be earned through both the delivery of high quality work and buyer satisfaction.

We expect these reputational cues, such as star ratings, to alter perceptions of spanning. Past demonstrations of a variety of activities, be it diversification into disparate industries by firm (Zuckerman 1999) or selling items across several different eBay categories (Hsu et al 2009), elicits one of two possible conclusions. First, that the social actor is highly skilled and is therefore demonstrating a broad range of abilities when they engage in a variety of domains. Or, second, that the social actor is not particularly skilled at all, and therefore moving from domain to domain in search of a suitable one. To the extent that reputation represents, at least in part, an overall assessment of quality of a candidate, then spanning should serve as an additional signal that would be congruent with broader ability. This is because it is more difficult to do many things well, so in demonstrating such a feat indicates more robust underlying qualities, a social actor will garner additional attention - what we call "idiosyncrasy credits."

Note, this may merely reflect the assumption that people hold of generalists necessarily being less-talented than specialists. Because of this assumption, an evaluator may come to expect generalists to have lower reputation, so they therefore care less about such a measure of quality when encountering it in a market setting –called the "generalist discount." Second, those that decide to transact with generalists may be looking to fill particular needs which may loom larger than mere reputation signals. For example, people shop at large department store not necessarily because they offer higher quality, but rather due to their convenience. Similarly, restaurants which attempt to cater to a broad range of tastes continue to persist perhaps due to their appeal to groups of heterogeneous patrons who cannot agree on a narrow cuisine. In doing so, those consumers may come to practically expect to see lower quality. These beliefs should lead to poor

reputational signals to bolster the lowered expectations a buyer already has with social actors who span categories. In sum, idiosyncrasy credits and the generalist discount should lead to the following prediction:

> *Hypothesis 3:* For those with a good (poor) reputation, spanning increases (decreases) ones likelihood of winning a job.

DATA AND METHODS

Empirical Setting

The research questions were tested with data on transactions from Elance, an online marketplace for freelancing services. Many characteristics of this setting stand out as apt for examination. First, Elance provides a convenient sandbox to examine external labor market dynamics. Examinations of external labor markets are tricky: sample boundaries are difficult to define; data collection is often challenging. Here, the online microcosm predetermines boundaries: seller-buyer transactions are here bounded by the current platform; yet the ecology resembles that of external labor markets in that employers and freelancers interact on short-term projects.

Second, because the setting is computer mediated assessment barriers become more salient; signals are especially important because the *actual quality* of the freelancer is difficult to ascertain *a priori* the award of a job contract. Aside from repeated transactions with a particular freelancer, employers have very few data points to infer quality apart from what is presented on their profile – which we observe. The salience of both the online reputation and job categorization system becomes evident as assessment signals of both the seller's ability, quality of work and identity (or lack thereof). As such, the platform provides strong, reliable signals side by side categorized historical records of past work. This organization of the online labor marketplace via categorical boundaries and reputation symbols mirror our natural instincts to

cluster and rank objects; this lends itself nicely to the independent variables of categorization and reputation.

Third, the setting dictates that buyers necessarily describe their needs from a particular job market category. They do so by composing a job posting which includes a text description of the work. These job descriptions provide rich, unstructured text data from which we are able to infer how job categories are utilized in discourse space. In particular, job descriptions allow a more nuanced conception of categories in two key ways. First, the definition of categories is not endogenously derived via post-hoc examinations of similarity in overlap measures as previously used. Rather, we directly record and aggregate employer expectations of how each category will be utilized. Second, job descriptions allow us to infer much more variation among the categories and not force us to treat categories as orthogonal to one another. Instead, as we detail below, we are able to represent categories as points in both discourse and time space.

Finally, the job bidding auction paradigm utilized on Elance allows us to construe of each job listing as a pseudo-game; we observe all information that is available to the employer, including the population of bidders for a particular job. This system thus provides strong counterfactuals to test our hypotheses as we observed both successful and unsuccessful contracts. Such a design will be considerably more difficult to implement in a face-to-face "real world" context.

Data Sampling

Our dataset consists of all activity on the website from January 2008 to December 2011. This includes all employer and freelancer activity such as each individual job description, the employer who posted it, the job category in which it belonged, and the precise time when it was

listed. Freelancer bidding information included the identification of the freelancer, the amount of their bid, and who was eventually awarded the job.

At the core of our design is the notion that each job auction can be viewed as a "pseudo-game:" each employer is presented with a choice of well-characterized freelancers from which he/she has to choose among. To approximate the pseudo-game structure, two data sets are formed. The first is a characterization data set used to operationalize both categorical positions and the past history of employers and freelancers. The second is the analysis data set used to test our hypotheses. We will describe these two data sets in turn. The characterization data set is a cumulative, dynamically updated data set. It accumulates the posted job descriptions over time, as well as buyer awarded feedback ratings and experience measures. These characterization variables are trained on all clean job descriptions and all awarded contracts in the full time window of the data (2008-2012) with minimum restrictions.

In contrast, for analysis, we only consider data that approximates the pseudo-game. As such, multiple criteria are needed to form the analysis subset. This subset is sampled at the level of the Job posting (henceforth "jobs"). A job is included if its characteristics meet the following criteria. First is the competition and singular winner criterion: i.e. a single freelancer was eventually hired and more than one freelancer was rejected. This eliminates jobs whereby multiple/all bids were accepted, no bids are accepted, or there were no bids at all. Second, we adopt a singular signal criterion: jobs which include even just one seller who registered multiple bids were removed from the analysis data set. This is done so as to limit potential confounding signals (e.g. multiple bids/rapid updating signalling strong interest). Third is a sufficient history criterion. Bidders need to register at least four prior jobs to stabilize their categorical positions. As such, completely new sellers fall out of our categorical measure. Jobs are thus sampled only if

all bidders have had four or more prior completed contracts. Finally, we left censor the data by two years to capture relevant prior actor histories at the beginning of the analysis window.

The complete data set for the duration of 2008-2012 registered 2,379,309 bids for 495,020 jobs. All 495,020 jobs are used for the characterization data set. In contrast, only 173,793 bids for 41,929 jobs fit the criteria imposed for the analysis data set.

Category Space

The process of categorization can be described as a social process that determines an actor's market legitimacy. As early as Weber (1924), social objects and orderings are considered to be legitimate if such items approximate the average orientation of "determinate maxims or rules;" categorization rules are thus implicit impositions of the market audience (employers). These audiences determine a candidate's (freelancers) membership in particular categories, and assesses the candidate's categorical make-up.

In this paper, we develop upon Hannan's restatement of categories as being fuzzy (e.g. Hannan et al. 2007). While earlier conceptions of categories (e.g. Zuckerman et al 2003; Philips and Zuckerman 2001) describe binary memberships as determined by an institutional "injector", the external labor market asserts no such schema: the free-form market allows freedom of categorical combinations and posits significant fuzziness in categorical sets. This is doubly asserted by empirical observations: although Elance attempts to categorize available talent into 168 subcategories, market buyers reject the equivalency of this categorization. Hiring activity is extremely lop-sided, with the top 30 categories making up 80 percent of market activity. Evidently, the absence of an expert or institutional arbiter poses a significant challenge to the attribution of legitimate categorical combinations.

While prior research in categories have attempted to discern legitimate *vis a vis* illegitimate categorical combinations via candidate-side similarity, these endogenous conclusions are post-hoc and assumes candidate compliance. We take advantage of the availability of audience discourse (namely job postings) in our setting: conceptions of labor skill categories as assessed by the external labor market will be directly measured via audience volunteered descriptions. We thus propose a novel measure of experience relevance and categorical spanning that utilizes audience-side employer contributed job descriptions. We consider these descriptions as time-dependent snapshots of market expectations, sentiments, and requirements of a particular job category. The granularity and richness of text descriptions allow us to abstract a cognitive categorical space upon which the 168 categories are posited. As such, the similarities and differences between categories are reified as distances in this space.

This paints a richer picture of categorical clustering. For instance, the job categories of Simple Website and Web Design will be relatively similar compared with the job categories of Simple Website and Academic Writing. In this categorical space, the distance between the two former categories will be smaller than that of the two latter categories. In addition, the time-dependent quality of job descriptions is reflected in the spatial dimension: the focal positions of the categories move in this space as the market updates its conception and definition of the categories. For instance, the categories of Flash Animation and Web Design would be more similar (i.e. closer in categorical space) in 2004 than they are today. Correspondingly their distance will have increased since their association in 2004.

To define the space, we first process the job description texts via the following steps. First, we strip the text of all stop-words, words that do not add any context specific meaning whatsoever (e.g. "is", "are", "and", "the"). Following which, each word is stemmed into stem-

tokens. This step implies that related words are collapsed into a single token. For instance, "consulting", "consultant", "consultation", and "consult" will all be stemmed into the token "consult." Similarly, "managing", "manager", management" etc. will all be stemmed into the token "manag." Following which, duplicate mentions of the words are purged and the word-order disregarded. At this stage, each job description is characterized as a "bag" of unique words. We call the sum of all bags of words the corpus.

To characterize the categorical space, we consider each unique word in the corpus as a dimension in the space. This is thus a binomial model: each word dimension is binary. We do not consider a multinomial model (whereby multiple counts of words are registered) as we are interested in the distances and clustering of word *meanings* and not word *intensity*.

In our corpus, we have a total of 27,080,872 words collected, of which 2,246,469 are unique. As a 2,246,469 dimensional space will be a significant computational challenge, we perform a selection of the most prominent features, capitalizing on the fact that word frequencies in language exhibits a power law (Newman 2005). To check this, we fit a power law distribution to our data. The resultant fit is both good and statistically significant (*α = 2.08, xmin = 8120, K-S.stat = 0.06, p = 0.01917*). As such, we selected the 8,695 most frequent words to characterize the categorical space. These 8,695 words comprise 95% of the total words in the corpus.

The relative frequency of word use in each category then denotes the category's position in this space at a particular time *t*. This is achieved through the characterization vector: an audience defined, magnitude normalized, and cumulative average of all jobs word vectors that fall into the category before time *t*:

$$\vec{r}_{Ct} = \frac{1}{\left\| \sum_{c \in C, \tau < t} \vec{x}_{c\tau} \right\|} \sum_{c \in C, \tau < t} \vec{x}_{c\tau}$$

Here, $\vec{r}_{Ct}$ denotes the characterization vector of category C at time t; $\vec{x}_{c\tau}$ is the word vector of each individual job in the category C at time τ as given by its bag of words (binomial word count). To check our characterization vectors for face validity, we calculated the Euclidean distances between the top 30 most transacted categories in Elance as below:

$$distance_{C_{it}C_{jt}} = \left\| \vec{r}_{C_{it}} - \vec{r}_{C_{jt}} \right\|.$$

Visually, a heatmap via the **heatmap()** function in R demonstrates the between-category distances; a hierarchical clustering algorithm reveals categorical groupings.

[Insert Figure 3 about here]

Figure 1 shows the heatmap of categorical distances coupled with the dendogram of hierarchical clusters. Amongst the 168 subcategories, related categories emerge through their discourse space positions: we see clusters that denote jobs related to broader categorical conceptions. For instance, in figure 1, subcategories that pertain to writing form a clear cluster ("Article Writing", "User Guides and Manuals", "Technical Writing", "Copywriting", "Academic Writing"). Clusters also form around jobs that pertain to design ("Website Design", "Other-Design", "Graphic Design") and web programming ("Other IT and Programming", "Web Programming", "Other-Engineering"). All in all, the categorical positions as posited in the discourse space have good face validity.

More importantly, this exercise shows that perceived relatedness and the legitimacy categorical combinations can be discern separately from candidate typicality. This point will be employed in the following exposition on categorical distance and spanning.

Categorical Distance and Spanning

We note that over time, the understanding and conception of each market category will change. For instance, Flash Animation was a defining skill in the 90's for Website Programming: today, Flash Animation is viewed as bloated and outmoded. While aggregate cross-sectional typicality ignores such nuances in category dynamics, the current time dependent updates of each categorization vector measures this: each seller occupies a time-dependent position in the category space depending on his/her past completed job categories; this position is determined by audience defined categorical vectors.

A center of mass measure indicates a seller's position as a point in categorical space:

$$\vec{R}_{it} = \frac{1}{M} \sum_{c \in C_{it}} \vec{r}_{ct}$$

Here, $\vec{R}_{it}$ indicates the position of seller i at time t. C_{it} is the set of categories that labels seller i's past 4 jobs prior to time t. We use the most recent four jobs as these are the jobs that are prominently featured in the seller profile. M is the number of all jobs with feedback prior to time t. $\vec{r}_{ct}$ is the characterization vector of category c at bid time t.

This position can then be compared to the current job applied for. The ability to affix categorical positions for both the freelancer and the current job category allows for direct assessments of categorical relevance: this is a step-up from candidate-side typicality measures that compares the focal candidate amongst his peers. Here, we believe that categorical similarity, as captured by our measure of distance, should serve to indicate how relevant a freelancer's past experiences are to the current bidding category. In essence, we compute relevance as a function of the distance of the candidate from the needs of the *audience*. This center of mass is compared with $\vec{r}_{c_{bid,t}}$, the current job of interest that the freelancer is bidding on at time t. The Euclidean

distance of these two vectors operationalizes freelancer relevance from the employer job category:

$$distance_{bid,it} = \left\| \vec{R}_{it} - \vec{r}_{c_{bid,t}} \right\|$$

A freelancer that has experience that is completely relevant to the job he/she is bidding on will have a distance score of zero. As this is a distance measure, relevance is thus reverse coded.

Similarly, employer audience defined categorization vectors allow us to determine freelancer candidate spanning as *perceived by the employer*. This is operationalized as the average of all Euclidean distances between all past job categorization vectors of a particular freelancer at time of bid. A freelancer who only took jobs from a single category will have a spanning score of 0. Note that unlike the relevance measure, the spanning variable is independent of the job category of the auction. Also note that unlike the relevance measure, the spanning measure is *not* reverse coded. Thus:

$$spanning_{it} = \frac{1}{M} \sum_{c_j, c_k \in C_{it}, j \neq k} \left\| \vec{r}_{c_j t} - \vec{r}_{c_k t} \right\|$$

We believe this distance measure of spanning improves upon traditional existing measures of generalism and specialism, such as the Herfindahl index, in at least two ways. Firstly, for practical purposes, given an employer will likely only perceive buyer categorical history in their past four jobs, the Herfindahl index only takes into account five possible combinations of job category spanning (1:1:1:1, 1:1:2, 1:3, 2:2, 4). In contrast, our spanning measure is asymptotically continuous. This is illustrated in Figure 4 where we present three plots; the calculated Herfindahl range against the continuous measure of spanning, the histogram of the Herfindahl, and the histogram of our spanning measure. Note, first, the spanning measure is certainly, as expected, correlated with the Herfindahl distance measure; second, as also

mathematically expected, a single Herfindahl measure corresponds to a multitude of spanning values; finally, the distribution of the spanning distance is considerably better behaved than that of Herfindahl scores.

Secondly, we employed a simple mediation analysis to compare the explanatory power of the two measures. In separate linear probability model (LPM) regressions with job fixed-effects, spanning distance and Herfindahl both have negative effects on a bidder's likelihood of winning (after Leung, 2014). However, when included in the same regression, the significance of the coefficient of spanning distance remains while the effect of Herfindahl becomes non-significant.

Reputation

The reputation of a freelancer depends upon the feedback ratings provided by employers upon job completion. Building upon our definition of reputation as the expected capability derived from past displays, we operationalize a freelancer's reputation score at time *t* as the cumulative average of all previous feedback scores. Should the employer choose not to give a feedback score, the observation is dropped for the purposes of the cumulative average; it is not entered into both the numerator and denominators of the reputation score. This measure reflects what is prominently displayed on the freelaner profile. Thus:

$$reputation_{it} = \frac{1}{N_t} \sum_{\tau < t} feedback_\tau$$

Where $reputation_{it}$ is the reputation score of freelancer *i* at time *t*; $feedback_\tau$ is the feedback score awarded to the freelancer at time τ; N_t is the number of feedback scores received by the freelancer up to time *t*.

Covariates

We attempt to address employer-side heterogeneity by utilizing job-fixed effects. However, we are concerned with two potential covariates which might correlate with our independent

variables of interest; these are included in the analysis. The first is a measure of overall freelancer experience. This is measured as the total number of prior contracts fulfilled by the freelancer. Experience is particularly necessary as our category discourse space measures stabilize with more freelancer information. The second is a monetary measure of the latest bid offered by the freelancer. This is measured in US dollars.

MODELS AND RESULTS

Modeling Methodology

The binary nature of the dependent variable lends itself to a logistic regression specification which estimates a dependent variable bounded by 0 and 1 (Long 1997). Here, we code 1 as a bid that wins a contract and 0 as one that does not. In addition, the online setting allows us to employ job level fixed effects to account for both buyer and time dependent heterogeneity.

The models run on the analysis data-set. As we are interested in reputation effects, we remove from the data all bids sellers with no prior reputation whatsoever. In addition, jobs which either fail to award any contracts or awarded contracts to *all* bidders are discarded to facilitate the within job design. This effectively removes jobs that exhibited *no seller competition* to obtain outcome variance within jobs. Note that these jobs are only discarded for the purposes of the regression model; they are included in the categorical definition measures.

Summary statistics of this data set are presented in Table 1. A table of correlations is presented in Table 2. As evidenced in Table 1, both covariates of experience and bid amount are long tail, pseudo-power law distributions. In addition, we find that the raw measurements of spanning and distance to exhibit significant positive skew. As such, transformations are necessary. We apply logarithmic transforms to the covariates of experience and bid amount (adding a small constant to the latter); we apply square root transforms for the categorical

variables of spanning and relevance. Upon transformation, the distributions of these variables approach Gaussian. Finally, these variables are then normalized to allow for ease of comparison.

[Insert Tables 1 and 2 about here]

Using a maximum likelihood estimator with job fixed effects, we estimated 7 different models. For these purposes, we used functions for the general linear mixed models R package glmmML (Broström, 2013). Cluster robust standard errors are reported.

We first check the covariates (Model 0). Following which, we test Hypothesis 1 by entering the variables without interactions in three models: the first (Model 1) examines the effect of Categorical Relevance on appeal; the second (Model 2) examines the effect of Categorical Spanning and the third (Model 3) enters both Relevance and Spanning into the regression. We then examine interactions between spanning and relevance (Model 4) and spanning and reputation (Model 5). These logit interaction models are particularly tricky to interpret and will require some care. A final saturated model enters both these interactions (Model 6); this model primarily checks for the independence of the two interaction effects and will not be interpreted.

Results

Estimated model coefficients are reported in Table 3. All models include all the control covariates. To interpret the results, we will first discuss models which employ no interaction terms first: these are models 1, 2 and 3. These models test the group of hypotheses that make up Hypothesis 1. Models 4 and 5 test Hypotheses 2 and 3.

[Insert Table 3 about here]

The first three models support hypothesis 1. Taken as separate constructs, spanning and distance operates on appeal in a manner that is consistent with extant literature and current

expectations. Model 1 shows that for an increase in 1 standard deviation in relevance, the probability of getting hired is increased by about 3.4%. Model 2 shows a smaller effect for spanning: an increase in 1 standard deviation in spanning results in a 0.9% reduction in probability of getting hired. When entered in the same regression, as in Model 3, we find that the significance of spanning goes away: thus supporting the multiple implications of our first hypothesis: spanning is mediated by relevance of the seller.

Having established the independent effects, we enter interaction terms[1]. In Model 4 we test and find support for our second hypothesis. By entering the interaction of spanning and relevance here, we see a positive and significant effect suggesting that as spanning increases along with relevance, the likelihood of being hired increases. We demonstrate this interaction graphically in Figure 5 which plots the change in the likelihood of winning a bid on the vertical axis as a function of the relevance of the freelancer on the horizontal axis. First, of note is the overall positive effect of relevance on winning a bid. Second, in order to demonstrate the interaction effect with spanning, we plot three lines which represent the effect of these two variables on three different levels of spanning. As we can see, the greater spanning, the steeper the slope of this relationship, suggesting that an additional benefit is garnered for those freelancers that exhibit both high spanning and high relevance. Interestingly enough, the discount/benefit of spanning varies for different levels of relevance: spanning actually benefits freelancers who are viewed relevant while, in contrast, it doubly hurts distant/less relevant freelancers. For example, as the level of relevance of a freelancer's past experiences increases 1 Std. Dev. from the mean, at the highest level of spanning (the red line, two Std. Dev. above the mean), there is an increase in the likelihood of winning the bid by ~6% however, for an identical

---

[1] We recognize the risks inherent in estimating interaction effects in a logistic regression (Ai and Norton 2003), we therefore corroborated all the findings utilizing a Linear Probability Model (LPM).

move in relevance for a low spanner (two Std. Dev below the mean) the likelihood of winning only increases by ~2%. Together, these observations of the results support Hypothesis 2.

[Insert Figure 5 about here]

In Model 5 we test and find support for our third hypothesis. By entering the interaction of spanning and reputation here, we see a positive and significant effect suggesting that as spanning increases along with reputation, the likelihood of being hired increases. We demonstrate this interaction graphically in Figure 6 which plots the change in the likelihood of winning a bid on the vertical axis as a function of the reputation of the freelancer on the horizontal axis. This functional form demonstrates how a higher reputation will increase the appeal of all sellers, but at a much higher rate for high spanning sellers. For example, as reputation increases one Std. Dev. above the mean, a high spanner (two Std. Dev. above the mean) increases their likelihood of winning by ~6%, whereas a low spanner (two Std. Dev. below the mean) increases their likelihood of winning by only 4%.

[Insert Figure 6 about here]

DISCUSSION

How we work is changing. As the labor force is increasingly embracing careers that are no longer tied to one organization (Bidwell and Briscoe 2010) or leaving permanent employment for self-employment as freelancers (Barley and Kunda 2004) – research which examines how employers and employees navigate these new career structures is increasingly warranted. This study brings to the fore what employees may face when they establish a set of experiences which may be seen as disparate to potential employers. Despite the burgeoning literature that has found penalties for spanning socially recognized boundaries across a range of circumstances, the so-called generalist discount, (Ferguson and Hasan 2013, Hsu et al 2009, Hsu 2006, Zuckerman

32

1999), we set out to identify the circumstance under which this behavior can be beneficial – an idiosyncrasy credit. In doing so we bolster the claims of more nascent work which has attempted to bring additional nuance to those original first order claims of a multi-category discount (Kovacs and Hannan 2013, Pontikes 2013, Leung 2014).

This paper contributes to the literature at the intersection of categories and labor markets by highlighting how the advent of contemporary markets which utilize technology to mediate the relationship between employers and employees may serve to reveal the additional challenges of assessing ability and skill *a priori*. Traditionally, the contract labor market was mediated by staffing agencies which added value by identifying the skills needed for each task and matching employers with employees. As these types of markets become increasingly virtual, thereby eliminating this role, employers are now left to learn how best to match their requirements with employees. Future work could continue to identify how hiring may be more or less similar to more traditional offline markets.

For example, it would be fruitful to demonstrate if findings in this external labor market setting are also reflected by the employment outcomes in internal labor markets. We may expect different results. Internal labor market participants differ from external ones because they are valued for the firm-specific skills they possess (Fernandez-Mateo 2009). To the extent that specializing in a domain, such as accounting, within a firm suggests the employer is accumulating a stock of human capital that is highly specific to the firm, such as social capital with other employees, familiarity with the firm's processes, and historic knowledge, then specialism should be highly prized (Fergusen and Hasan 2013). However, in an external labor market, such as Elance, the temporary and job specific nature of these transactions would suggest

that employers are only seeking general skills and therefore generalism is an indicator of perhaps both the skills a freelancer's possesses as well as once she can learn.

Empirical challenges remain. First, the non-linear interaction terms of the logistic regression proves to be extremely rich and complex, more could be done to unpack this effect. While we presented that the spanning reputation interaction term supports Hypothesis 3, a more accurate statement would be that the interaction term supports hypothesis for *certain subgroups of the population*. In additional analyses we find that the interaction effect is less significant for low likelihood of hire bids – that is, employers are not examining the backgrounds of those freelancers whom they are not likely to hire (for other reasons) anyway.

Our sample considers only experienced sellers with established job histories on the platform. These sellers have accrued reputation and history, and as such, reduced a certain amount of uncertainty about their quality. However, this strong sampling criterion might introduce certain biases. Firstly, the examination of survivors implies that sellers that have truly bad reputation will be left out of this sample. Our examination is therefore one of the sellers that survive on the platform and have demonstrated reasonable reputation. Secondly, we ignore first time entrants: candidates that poses the highest amount of quality uncertainty. Literature seems to suggest that vertical and horizontal differentiation signals will be most salient for this high uncertainty subgroup.

Taking our findings to their logical extremes would suggest that being an excellent ballerina as well as a highly esteem lumber jack would make such a job seeker highly desirable – an unrealistic position to support. We make two points to address this. First, note that the jobs proffered on this website don't span such extremely disparate domains and are those we expect to be carried out in an office environment (see Appendix). This drastically limits the amount of

spanning any worker on the website could engage in, thereby limiting the range of past experiences. Second, empirically, we explored the relationship between spanning on both relevance and reputation. We draw your attention to the scatterplot of each seller's individual average bid relevance and categorical spanning in Figure 7 (there was no discernable relationship between spanning and reputation). Note that there is a noticeable decline in relevance as spanning increases, suggesting there is a restriction in the range of past experiences one could engage in while still attempting to remain relevant. That the upper right quadrant, a zone where freelancers are highly likely to be hired according to our functional form, is barely populated corroborates this. What we believe is happening is that the optimal position to inhabit by a freelancer is to endure a slight drop in relevance and compensating for by a slight increase in spanning (Leung 2014). That demonstrations of movement among similar, but not identical jobs, is the most advantageous job history to accrue.

[Insert Figure 7 about here]

In addition to these theoretical contributions, we have also attempted several methodological advancements. First, we focused on the job descriptions posted by the buyers in the market to measure distance between categories. These descriptions are time-dependent snapshots of market expectations, sentiments and requirements of a particular job category. The advantage here is that we are able to depart from previous conceptions of spanning which assume categories are independent and equally distributed in space – as captured by the Herfindahl measure. Instead, our continuous measure of spanning, as well as relevance, allows us to define spanning by not only the number of different experiences, but also how broadly that reaches.

Second, we also believe this represents an improvement over extant attempts at calculating distance between categories which use co-occurrence measures (Kovacs and Hannan 2013, Leung 2014) because it more clearly captures how categories are actually used by buyers. Ultimately, categories are socially consequential to the extent buyers use them. Our direct examination of jobs descriptions allows us to do this. A second advancement is our use of text analysis tools to divine the socio-cognitive distance between two job categories. As social interactions becoming increasingly mediated by computers, the amount of text will increase commensurately. Led by much of the work done by computer scientists, we are more and more able to manipulate bodies to text to address social science questions. Our attempt here demonstrated the theoretical value to text analysis.

REFERENCES

Ai, C., and E.C. Norton
  2003 Interaction terms in logit and probit models. *Economics letters*, *80*(1): 123-129.
Arthur, M.B., and D.M. Rousseau, eds.
  1996 *The Boundaryless Career: A New Employment Principle for a New Organizational Era*. New York: Oxford University Press.
Barley, S.R. and G. Kunda
  2004 *Gurus, Hired Guns and Warm Bodies: Itinerant Experts in a Knowledge Economy*. Princeton, NJ: Princeton University Press.
Barnett, M.L. and Pollock, T.G., eds.
  2012 *The Oxford Handbook of Corporate Reputation*. Oxford, UK: Oxford University.
Bidwell, M. and F. Briscoe
  2010 "The Dynamics of Inter-Organizational Careers in the New Economy." *Organization Science*, 21(5): 1034-1053.
Broström, G.
  2013 glmmML: Generalized linear models with clustering. *R package version 1.0*.
Donath, J. S.
  1999 Identity and deception in the virtual community. *Communities in cyberspace*, *1996*, 29-59.
Faulkner, R.R.
  1983 *Music on Demand: Composers and Careers in the Hollywood Film Industry*. New Brunswick, N.J.: Transaction.
Ferguson, J-P and S. Hasan
  2013 "Specialization and Career Dynamics: Evidence from the Indian Administrative Service." *Administrative Science Quarterly,* 58(2): 233--256.

Fombrun, C. J.
    1996 *Reputation: Realizing value from the corporate image*. Boston: Harvard Business School Press.
Fiske, S.T., and S.E. Taylor
    1991 *Social Cognition*. New York: McGraw-Hill.
Ganz, H., M. Evans, and M. Jalland
    2000 In Maury A. Peiperi, Michael B. Arthur, Rob Goffee, and Tim Morris edited *Career Frontiers: New Conceptions of Working Lives*. Oxford: Oxford University Press, 24-53.
Gibbons, R. and L. Katz
    1991 "Layoffs and Lemons." *Journal of Labor Economics*, 9 (4): 351-380.
Hannan, M.T.
    2010 "Partiality of Memberships in Categories and Audiences." *Annual Review of Sociology*, 36: 159-181.
Hannan, M.T., L. Pólos, and G.R. Carroll
    2007 *Logics of Organization Theory: Audiences, Codes, and E cologies*. Princeton: Princeton University Press.
Hsu, G., M.T. Hannan, and Ö. Koçak
    2009 "Multiple Category Memberships in Markets: An Integrated Theory and Two Empirical Tests." *American Sociological Review*, 74: 150-169.
Hsu, G.
    2006 "Jacks of all trades and masters of none: Audiences' reactions to spanning genres in feature film production." *Administrative Science Quarterly*, 51: 420–50.
Hsu G. and M.T. Hannan
    2005 "Identities, genres, and organizational forms." *Organization Science*, 16(5): 474-490.
Jensen, M., H. Kim, and B.K. Kim
    2012 "Meeting Expectations: A Role-Theoretic Perspective on Reputation." In *The Oxford Handbook of Corporate Reputation*, edited by Michael L. Barnett and Timothy G. Pollock. Oxford University Press: 140-159.
Kovacs, B. and M.T. Hannan
    2012 "Category Spanning, Distance, and Appeal." *Working paper*, University of Lugano.
Kalleberg, A. L.
    2000 "Nonstandard Employment Relations: Part-time, Temporary and Contract Work." *Annual Review of Sociology*, 26: 341–65.
Kalleberg, A. L.
    2011 *Good Jobs, Bad Jobs: The Rise of Polarized and Precarious Employment Systems in the United States, 1970s-2000s*. New York: Russell Sage Foundation, American Sociological Association Rose Series in Sociology.
Kleinbaum, A.M.
    2012 Organizational Misfits and the Origins of Brokerage in Intrafirm Networks. *Administrative Science Quarterly*; 57: 407-452.
Lange, D., P. M. Lee and Y. Dai
    2011 'Organizational reputation: A review'. *Journal of Management,* 37:153-184.
Leung, M.D.
    2014 "Dilettante or Renaissance Person? How the order of job experience affects hiring in a virtual labor market." *American Sociological Review*, 79: 136-158.

Leung, M.D. and A.J. Sharkey

2013 "Out of Sight, Out of Mind? The audience-side effect of multi-category membership in markets" *Organization Science,* 25 (1): 171–184.

Long, J. Scott

1997 *Regression Models for Categorical and Limited-Dependent Variables*. Thousand Oaks, CA: Sage.

Merluzzi, J. and D. Phillips

2014. "The Specialist Discount: Negative returns for MBAs with focused profiles in investment banking." Tulane University working paper.

Moss, P., H. Salzman, and C. Tilly

2000 "Limits to Market-Mediated Employment: From Deconstruction to Reconstruction Of Internal Labor Markets." In *Nonstandard Work: The Nature and Challenges of Changing Employment Relationships*, 95-121.

Negro, G., M.T. Hannan, and H. Rao

2011 "Category reinterpretation and defection: Modernism and tradition in Italian wine making." *Organization Science*, 22(6): 1449–1463.

Negro, G. and M.D. Leung

2013 "'Actual' and Perceptual Effects of Category Spanning: An Analysis of Critical Ratings of Italian Wines," *Organization Science*, 24(3): 684-696.

Newman, M. E.

2005 Power laws, Pareto distributions and Zipf's law. *Contemporary physics*, *46*(5), 323-351.

O'Mahony, S. and B. Bechky

2006 "Stretchwork: Managing the Career Progression Paradox in External Labor Markets," *Academy of Management Journal* (49) 5: 918-941.

Osnowitz, D.

2010 *Marketing Expertise: Contract Professionals in the New Economy*. New York: Cornell University Press.

Phillips, D.J. and E.W. Zuckerman

2001 "Middle-Status Conformity: Theoretical Restatement and Empirical Demonstration in Two Markets." *The American Journal of Sociology*. Vol. 107, pp. 379-429.

Podolny, J. and D. Phillips

1996 "The Dynamics of Organizational Status." *Industrial and Corporate Change* 5(2): 453-471.

Polivka, A.E., S.R. Cohany, and S. Hipple

2000 "Definition, Composition, and Economic Consequences of the Nonstandard Work Force". In *Nonstandard Work: The Nature and Challenges of Changing Employment Arrangements*.

Pontikes, E.G.

2012 "Two Sides of the Same Coin: How Ambiguous Classification Affects Multiple Audiences' Evaluations." *Administrative Science Quarterly*, 57(1) 81-118.

Rao, H.

1994 The social construction of reputation: Certification contests, legitimation, and the survival of organizations in the American automobile industry. *Strategic Management Journal*, 15: 29-44.

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*, Vol. 108: 795–843.
Rao, H. P. Monin, and R. Durand
    2005 "Border crossing: Bricolage and erosion of categorical boundaries in French gastronomy." *American Sociological Review*, 70: 968-991.
Resnick, P. and R. Zeckhauser
    2002 Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. The Economics of the Internet and E-Commerce. Michael R. Baye, editor. Volume 11 of *Advances in Applied Microeconomics*. Amsterdam, Elsevier Science. 2002. pp. 127-157.
Resnick, P., R. Zeckhauser, J. Swanson, and K. Lockwood
    2006 The Value of Reputation on eBay: A Controlled Experiment. *Experimental Economics*. Vol. 9,2: 79-101.
Ruef, M.
    2000 "The Emergence of Organizational Forms: A Community Ecology Approach." *American Journal of Sociology*, 106 (3): 658–714.
Roberts, P.W., and G.R. Dowling
    2002 Corporate reputation and sustained superior performance. *Strategic Management Journal*, 23: 1077-1093.
Rosch, E.
    1978 "Principles of categorization." In E. Rosch, and B. B. Lloyd (eds.), *Cognition and Categorization*: 27-48. Hillsdale: Lawrence Erlbaum Associates.
Shapiro C.
    1983 "Premiums for high quality products as returns to reputations," *Quarterly Journal of Economics*, Vol. 98: 659–679.
Sharkey, A.
    2013 "Categories and Organizational Status: The Role of Industry Status in the Response to Organizational Deviance." *American Journal of Sociology.* Forthcoming.
Stuart T. E., H. Hoang, and R.C. Hybels
    1999 "Interorganizational Endorsements and the Performance of Entrepreneurial Ventures" *Administrative Science Quarterly*, Vol. 44, 1999**.**
Smith, E.B.
    2011 "Identities as Lenses: How Organizational Identity Affects Audiences' Evaluation of Organizational Performance." *Administrative Science Quarterly*, 56(1), 61-94.
Spence, M.
    1973 'Job market signaling'. *Quarterly Journal of Economics*, 87, pp. 355-375.
Zerubavel E.
    1997 *Social Mindscapes*. Cambridge, Massachusetts: Harvard University Press.
Zuckerman, E.W., T. Kim, K. Ukanwa, J. Rittmann
    2003 "Robust Identities or Nonentities? Typecasting in the Feature-Film Labor Market." *American Journal of Sociology*, 108 (5): 1018-1074.
Zuckerman, E.W.
    1999 "The Categorical Imperative: Securities Analysts and the Legitimacy Discount." *American Journal of Sociology*, 104:1398–1438.

Figure 1
Elance Job Listing

C# Web Programmer ☆ Add to Watch List

IT & Programming > Web Programming

Posted: Tue, Oct 01, 2013

Time Left: 12d, 3h

Location: Anywhere

Start: Immediately

Hourly Rate: $10 - $15 / hr

Hrs/wk: 40 | Duration: 1-2 weeks

Work View™ Payment Protection

U.S. freelancers must have W9

Less Detail ▲

Client Info | Peoria, Arizona, United States 🇺🇸

I am needing a C# programmer who is great with asp.net... You also need to be an awesome MSSQL guy. I am a programmer by trade and have a lot of extra work that I need help with. I have an application that is built in MS Access and needs to be fully converted to ASP.net and MSSQL. I have started this application but I need help finishing it. I do not hire companies and I only work with individuals. Please put "I love C#" inside your application. I need someone to start ASAP.

Desired Skills

.NET, C#, Microsoft SQL Server Programming

Job ID: 47509548

Report Violation

Figure 2
Freelancer Profile

Raj Kamal Singh

Web Developer. PHP, JQuery, HTML, CSS, Django

Hire Me »

India | Bangalore, Karnataka | 12:36 pm Local Time

Overview

I am from an Indian Institute of Technology and have an experience of more than 4 years in the field of web development. I have had diverse experience. I have developed and used to maintain a static site of about 50 pages for my college and I have also created and maintained several dynamic websites and mobile applications for a large number of clients as a free lancer in the past.

Read More »

Job History

Private Job
Dec 3, 2013 | Software Application | $0 | Working

Django/Python backend and frontend developers...
Nov 29, 2013 | Web Programming | $0 | Working

Python GUI for Data Entry into Postgresql Database
Nov 26, 2013 | Software Application | $0 | Completed

Scripy in python to download and save dump...
Nov 26, 2013 | Software Application | $0 | Working

View All »

My Snapshot

IT & Programming ▾

LEVEL
1

12 months | Lifetime

Jobs		
	3	Total
	1	Milestones
	0	Hours
Reviews	1	★★★★☆ 4.4
	N/A	Recommend
Clients	1	Total
	100%	Repeat
Earnings	$179	Total
	$179	Per Client

Identity

Username	rkssisodiya
Type	Individual

41

Figure 3
Heatmap of Top 30 Job Category Distances
(Darker means More Relevant)
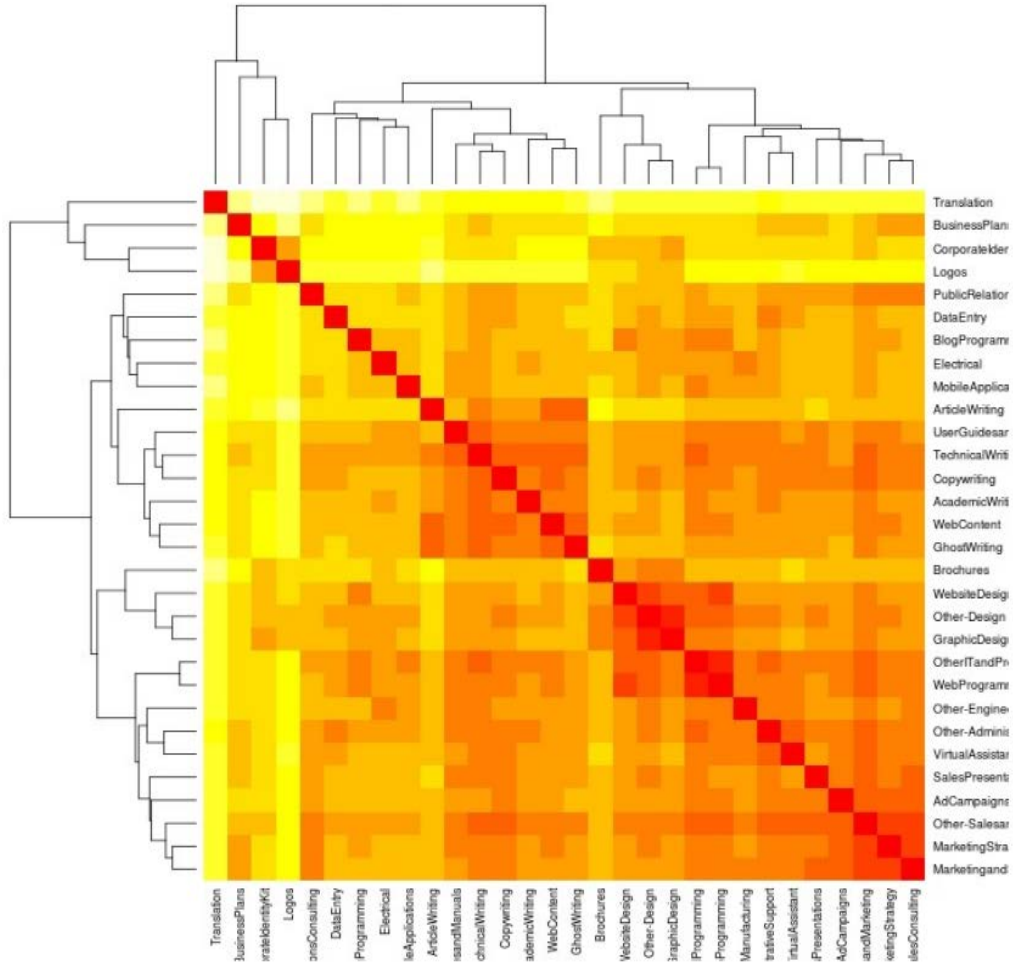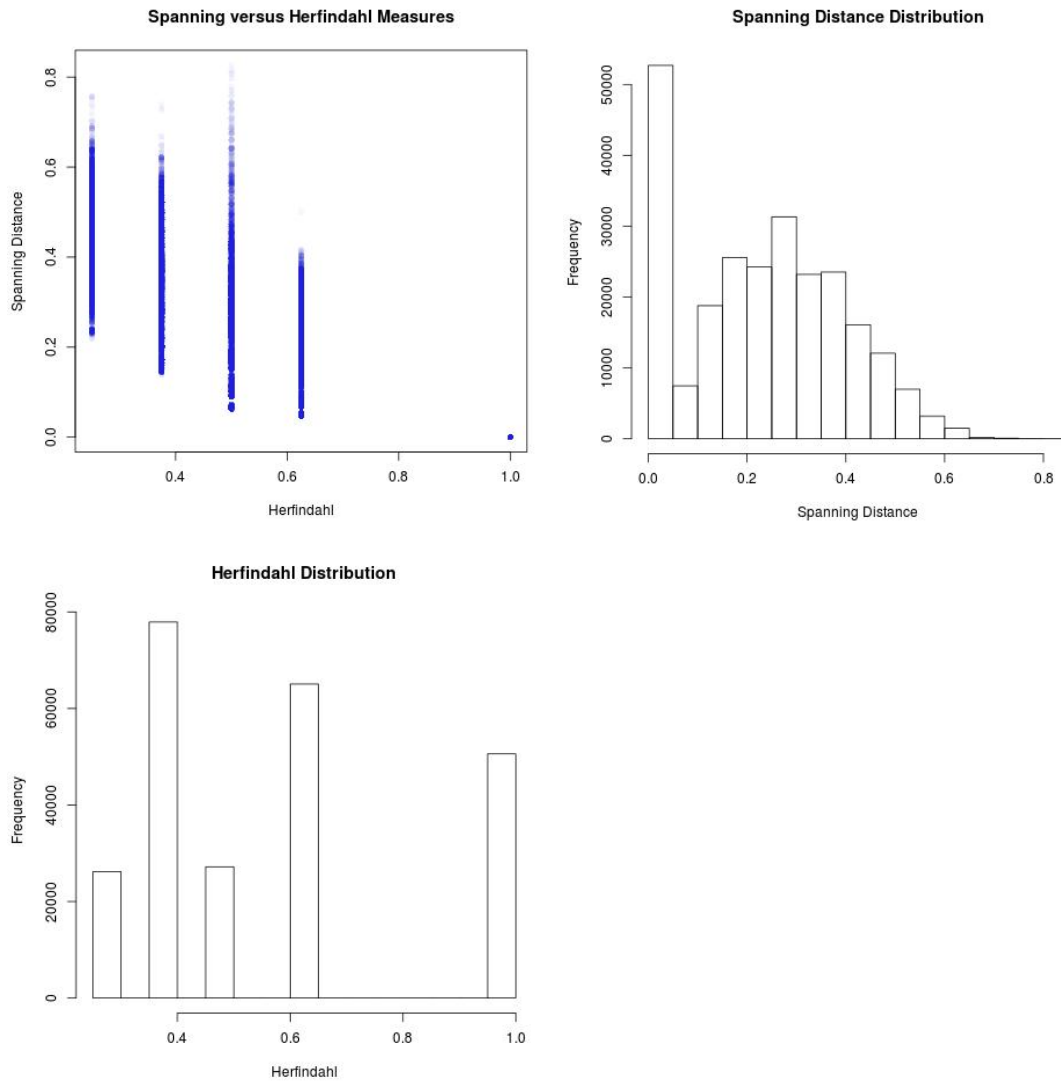
Figure 4
Herfindahl versus Continuous Spanning

Table 1
Summary Statistics

Variable	Mean	St.Dev.	Min	Max
Relevance	0.755	0.160	0	1
Spanning	0.224	0.160	0	0.882
Reputation	4.556	0.359	1	5
Experience	5.010	726	3	4036
Amount	438	3394	0	1127314

Table 2
Correlations

	Hired	Relevance	Spanning	Reputation	Experience
Relevance	0.022				
Spanning	-0.017	-0.467			
Reputation	-0.034	0.0001	0.006		
Experience	-0.099	-0.094	0.096	0.095	
Amount	0.156	-0.013	0.006	-0.036	-0.052

Table 3
Likelihood of Winning a Bid
(Fixed-Effects Logistic Regressions Estimates Grouped by Job)

	MODEL 0		MODEL 1		MODEL 2		MODEL 3	
Log. Experience	-0.191***	(0.007)	-0.188***	(0.005)	-0.182***	(0.005)	-0.183***	(0.005)
Log. Amount	0.165***	(0.005)	0.164***	(0.002)	0.165***	(0.002)	0.165***	(0.002)
Reputation	0.198***	(0.002)	0.197***	(0.007)	0.197***	(0.007)	0.197***	(0.008)
Spanning			-0.035***	(0.007)			0.005	(0.007)
Relevance					0.134***	(0.008)	0.105***	(0.009)
df	131861		131860		131860		131859	
AIC	312500		312500		312600		312700	

*Note: Standard Errors in Parentheses, * p < 0.05, ** p < 0.01, ***, p < 0.001*

Table 4
Likelihood of Winning a Bid
(Fixed-Effects Logistic Regressions Estimates Grouped by Job)

	MODEL 4		MODEL 5		MODEL 6	
Log. Experience	-0.181***	(0.005)	-0.183***	(0.005)	-0.181***	(0.005)
Log. Amount	0.165***	(0.002)	0.164***	(0.002)	0.165***	(0.002)
Reputation	0.197***	(0.007)	0.201***	(0.007)	0.200***	(0.007)
Relevance	0.154***	(0.013)	0.105***	(0.009)	0.154***	(0.013)
Spanning	0.012	(0.008)	0.005	(0.008)	0.012	(0.008)
Spanning x Relevance	0.044***	(0.008)			0.044***	(0.008)
Spanning x Reputation			0.018**	(0.006)	0.019**	(0.006)
df	131858		131858		131857	
AIC	312700		312700		312700	

*Note: Standard Errors in Parentheses, * p < 0.05, ** p < 0.01, ***, p < 0.001*

Figure 5
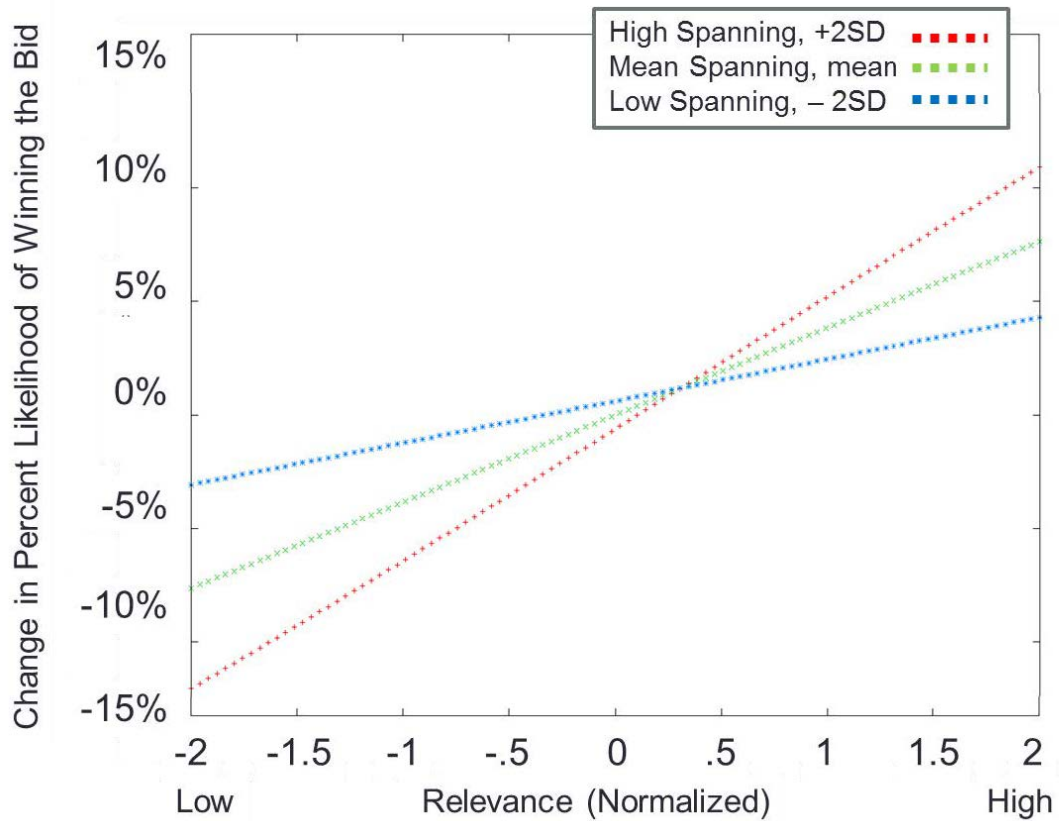Change in Likelihoo of Spanning and Relevance Interaction

Figure 6
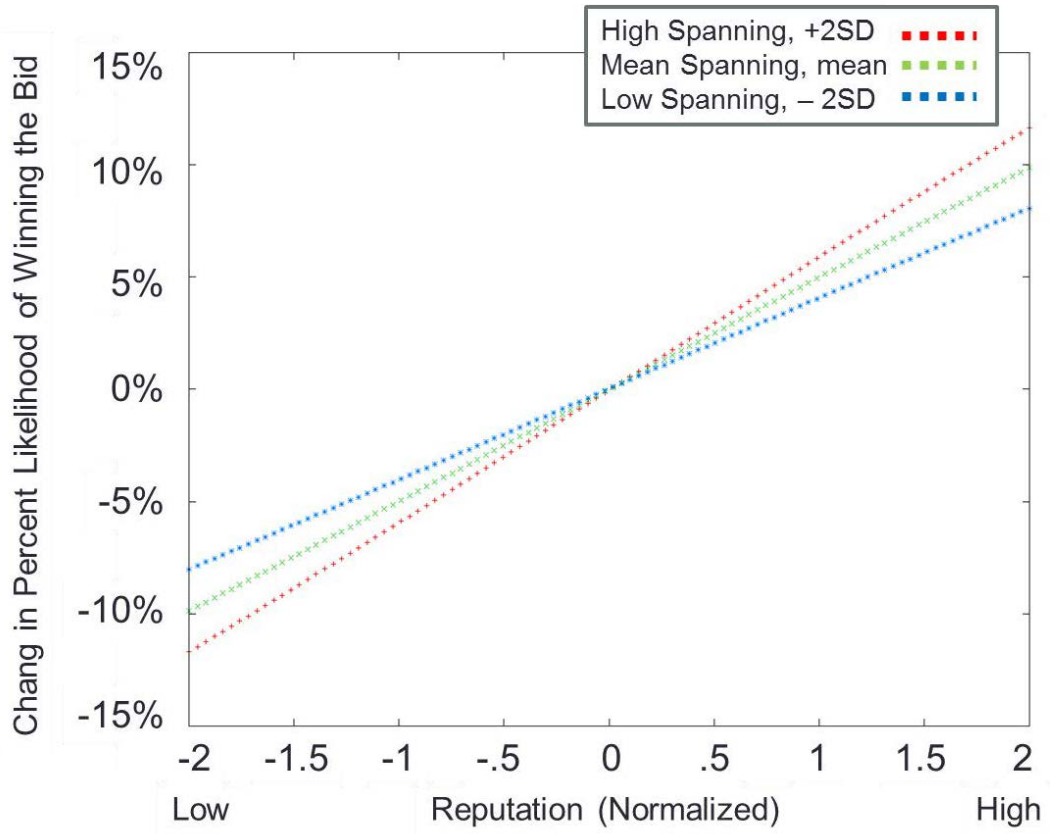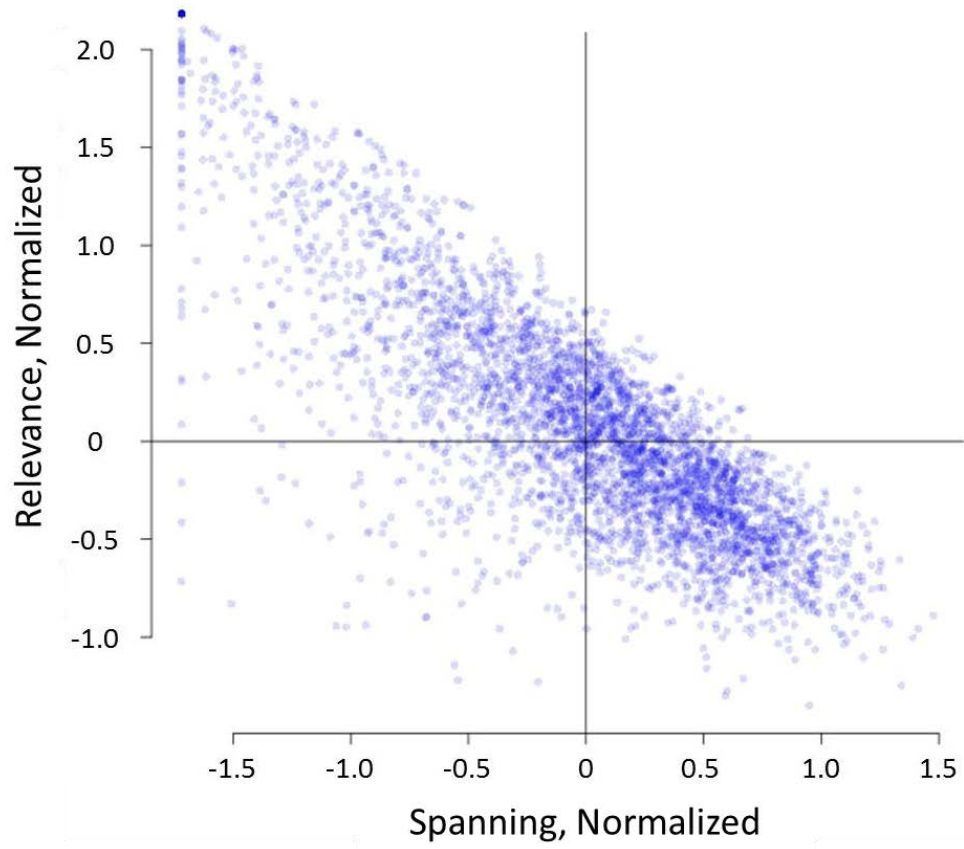Predicted Probabilities of Spanning and Reputation Interaction

Figure 7
Scatterplot of Bids, by Relevance and Spanning

Appendix – Elance Job Categories

ADMIN SUPPORT
Bulk Mailing
Customer Response
Data Entry
Event Planning
Fact Checking
Mailing List Development
Office Management
Other - Administrative Support
Presentation Formatting
Research
Transcription
Travel Planning
Virtual Assistant
Word Processing

DESIGN AND MULTIMEDIA
3D Graphics
Animation
Banner Ads
Brochures
Card Design
Cartoons and Comics
Catalogs
CD and DVD Covers
Commercials
Corporate Identity Kit
Digital Image Editing
Direct Mail
Displays and Signage
Emails and Newsletters
Embedded Video/Audio
Graphic Design
Illustration
Label and Package Design
Logos
Menu Design
Music
Other - Design
Other - Multimedia Services
Page and Book Design
Photography and Editing
Podcasts
Presentation Design
Print Ads
Radio Ads and Jingles
Report Design
Sketch Art
Stationery Design
Videography and Editing
Viral Videos
Voice Talent

ENGINEERING AND MANUFACTURING
Architecture
CAD
Civil and Structural
Contract Manufacturing
Electrical
Industrial Design
Interior Design
Mechanical
Other - Architecture and Engineering

FINANCE AND MANAGEMENT
Accounting and Bookkeeping
Billing and Collections
Budgeting and Forecasting
Cost Analysis and Reduction
Financial Planning
Financial Reporting
HR Policies and Plans
Management Consulting
Other - Management and
Finance
Outsourcing Consulting
Process Improvement
Stock Option Plans
Supply Chain Management
Tax

LEGAL
Bankruptcy
Business and Corporate
Contracts
Criminal
Family
Immigration
Incorporation
Landlord and Tenant
Litigation
Negligence
Other - Legal
Patent, Copyright and Trademarks
Personal Injury
Real Estate
Tax Law
Wills, Trusts and Estates

SALES AND MARKETING
Advertising
Branding
Business Plans
Business Skills
Business Software
Competitive Analysis
Corporate Training
Diversity Training
Email and Direct Marketing
Grassroots Marketing
Lead Generation
Management Training
Market Research and Surveys
Marketing and Sales Consulting
Marketing Collateral
Marketing Plans
Media Buying and Planning
Media Training
Other - Sales and Marketing
Other - Training and
Development
Policies and Manuals
Pricing
Product Research
Programming Languages
Project Management
Promotions
Public Relations

Retailing
Sales Presentations
Sales Training
Search and Online Marketing
Technical Training
Telemarketing
Tradeshows and Events

WEB AND PROGRAMMING
Application Development
Blogs
Database Development
Ecommerce Website
Enterprise Systems
Flash Animation
Handhelds and PDAs
HTML Emails
Network Administration
Online Forms
Other - Programming
Other - Website Development
Project Management
Quality Assurance
Scripts and Utilities
Security
SEO and SEM
Simple Website
System Administration
Technical Support
Usability Design
Web Design
Web Programming
Website QA
Wireless

WRITING AND TRANSLATION
Test Writing
Academic Writing
Article Writing
Children's Writing
Copywriting
Creative Writing
E-books and Blogs
Editing and Proofreading
Ghost Writing
Grant Writing
Newsletters
Other - Writing Services
Press Releases
Report Writing
Resumes and Cover Letters
Sales Writing
Speeches
Technical Writing
Translation
User Guides and Manuals
Web Content