Online Labor Markets: Reputation Transferability, Career Development Paths and Hiring Decisions

Marios Kokkodis
NYU Stern School of Business
mkokkodi@stern.nyu.edu

Abstract

In an online labor marketplace (*OLM*) employers post jobs, receive freelancer applications, and make hiring decisions. Once hired, freelancers complete the tasks online and receive their payment along with feedback about their performance. Because of the natural heterogeneity that appears in task categories, skills, and the latent abilities of freelancers and employers, these markets stuffer from a series of inefficiencies. In this work I focus on understanding these inefficiencies and propose solutions. In particular, I study three problems: (1) Reputation Transferability, (2) Skills Recommendations, and (3) Hiring Decisions. I start by proposing and evaluating different approaches that explain how freelancers' reputation transfers across different task categories. I then propose to examine the utility of skills in an *OLM*, given the level of expertise of the freelancer and the demand of each skill in the marketplace. Based on this analysis, I further propose to build a career development framework. Next, I build a series of explanatory and predictive models that describe employers' hiring decisions. In addition I study how each one of the available freelancers-employers characteristics affect hiring decisions. Finally, I conclude by highlighting the impact of my work on *OLM*s, and pose a series of questions that need to be addressed next.

Introduction

Online labor marketplaces (OLMs) such as oDesk.com and Freelancer.com allow employers to connect with freelancers around the globe to accomplish diverse tasks including web development, writing and translation, accounting, *etc.* These marketplaces are growing fast and the freelancer annual earnings are expected to grow from $1 billion in 2012 to $10 billion by 2020 (Agrawal et al. 2013). A typical scenario in these workplaces involves an employer posting a task, multiple freelancers bidding for it, some (or one) of them getting hired, completing the task online and finally receiving a payment.

In sync with offline workplaces, work in *OLM*s is an '*experience good*', meaning it is practically impossible to know the performance of a freelancer on a task in advance (Nelson 1970). To minimize this uncertainty, most of the online labor marketplaces have developed reputation systems: Freelancers get rated for the tasks they accomplish and these ratings become part of their online resumes. Employers can then get a better picture of the freelancers' past performance and make better informed hiring decisions. However, because of the diversity in task categories, past reputation might – or might not – be predictive of future performance when a freelancer switches to a new category. For example, what happens when a Web Developer decides to complete a task in Technical Writing? How this freelancer's past reputation in 'Web Development' transfers to 'Technical Writing'? Broader, are reputations transferable across categories and predictive of future performance?

On a different note, freelancers' value in *OLM*s resides in a combination of both observable and latent characteristics. The observed characteristics usually include a list of skills, the educational background, the work history and the certifications of the applicant. The latent characteristics include the freelancer's expertise and true ability on the listed qualifications. Very similar to the offline setting, the demand and supply distributions (and as a result the expected payoff) of each freelancer with a given set of skills and a given level of expertise are very heterogeneous; for example a Java expert might have a very different expected payoff than an expert in customer service support. Similarly, a $c\#$ expert might have a much higher expected payoff than a $c\#$ beginner, *etc.*. As a result, it would have been extremely useful, to both the freelancer and the marketplace, to know the value of each skill (or set of skills). But how can we quantify the value of each skill (or set of skills)? Furthermore, given a freelancer's current skillset and expertise, how can we tell which is the optimal career path to follow?

Although *OLM*s have provided to the employers a solution to the scarcity of local talent, they have not really changed the process that employers have to go through to source the ideal candidates for their tasks: an employer needs initially to describe the job opening requirements to which freelancers that are looking for opportunities can apply. Then, the employer has to (1) review all applicants by looking at their online profile information and/or by personally interviewing them, and (2) come up with a hiring decision. The existence of latent characteristics (*e.g.* freelancer quality), the heterogeneity that appears in the observed ones (*e.g.* freelancer's profile information, skills *etc.*.) as well as the interactions between the two make the matching process a very challenging task; Hiring decisions are based on manually shaped expectations of complicated similarities between job openings, employers and freelancers. However, it is unknown which – and how – freelancer's (or employer's) characteristics affects hiring decisions. In particular, how do employers choose freelancers? What attributes do employers value the most? In addition, is it possible to redistribute job applications to maximize both the number of hires as well as the long-term overall satisfaction of all the involved parties (freelancers, employers and the online marketplace)?

In this work I focus on addressing all these questions that naturally appear in *OLM*s, by proposing applicable solutions targeted to create a frictionless, mutually beneficial workplace for both the freelancers, the employers and the marketplace itself. More specifically, I start by proposing and evaluating different approaches that study freelancers' reputation transferability across different task categories. I then propose to examine the utility of skills in an *OLM*, given the level of expertise of the freelancer and the demand of each skill in the marketplace. Based on this analysis, I propose a career development framework for strategically recommending skills. Next, I build a series of explanatory and predictive models that describe employers' hiring decisions. I further study how each one of the available freelancers-employers characteristics affect hiring decisions. Finally, I conclude by highlighting the impact that my studies have towards a more efficient *OLM*, and pose a series of questions that need to be addressed next.

Brief Literature Review

Current research in Online Labor Markets (*OLM*s) spans across a variety of problems. A stream of work focuses on the validity of behavioral experiments in these markets and in particular on Amazon Mechanical Turk (Rand 2012). The general consensus of these studies is that online experiments appear to be as valid (both internally and externally) as laboratory field experiments. A different group of studies focuses on incentivizing freelancers as well as finding ways to manage the quality of their outcomes ((Mason and Watts 2010; Ipeirotis, Provost, and Wang 2010). These studies propose and evaluate a set of social and financial incentives, while the also provide sophisticated techniques that assure a certain level of outcome quality. Finally, an early version of the proposed study on reputation transferability is already published (Kokkodis and Ipeirotis 2013).

A lot of work has been done in the past that deals with skills' assessment and expert search. Hambleton *et al.* (Hambleton 1991) described the fundamental concepts of Item Response Theory (IRT), a theory widely used in Computer Adaptive Testing. Desmarais *et al.* (Desmarais, Maluf, and Liu 1995) proposed the creation of a network that captures implication relations among knowledge units (KU). Multiple studies have focused on using network analysis techniques to identify user expertise in online environments (Jurczyk and Agichtein 2007; Zhang, Ackerman, and Adamic 2007). Finally, (Saito et al. 2014) proposed a framework for developing micro-tasking skills.

In terms of hiring decisions, a lot of work has focused on gender and attractevness biases (Gilmore, Beehr, and Love 1986; Kawakami, Dovidio, and van Kamp 2005). The overall consensus of these studies is that gender and attractiveness have a strong effect on hiring decisions, but the type of the effect depends on the environment, the position, and the employer. Other characteristics that have been found to affect hiring decisions include applicant's clothing (Forsythe, Drake, and Cox 1985), demographics (Hu 2003), weak ties (Yakubovich 2005), 'cold-start problem' (Pallais 2012) and online profiles and political party affiliations (Acquisti and Fong 2013).

Research Models

Data

In my studies I use a unique transactional dataset from oDesk.com, that consists of 3,5 million job applications by 800,000 freelancers and 1 million hiring decisions (completed tasks/received feedback) by 200,000 employers. These transaction span 6 task categories: 'Software Development', 'Web development', 'Writing & translation', 'Sales & Marketing', 'Design & Multimedia' and 'Administrative'.

Study 1: Reputation transferability

In my first study I focus on whether and how freelancers' reputation is transferable across different task categories, as well as how current reputation systems can be benefited by incorporating information from other categories. I propose a set of predictive models to estimate the future performance of a user, based on category-specific past performance. Specifically, I assume that the category-specific qualities of a user are *latent and not directly observable*. However, these qualities are reflected into a set of other *measurable* characteristics, such as employer ratings for past projects. Based on these past ratings, I build models that are capable of connecting past performance *across categories* to predict performance in a new category for which the freelancer has either zero, or very few, past data points.

In all the proposed models I assume that there exists m categories of tasks, and that each user is endowed with a set of m category-specific, latent qualities which I denote with $q_{ij} \in [0,1]$(the quality of a user i in category j,$j \in \{1,\ldots,m\}$). The ultimate goal is to estimate q_{ij} by observing the user's past performance across all available categories.

Learning from past ratings, within category: I start with a very simple setting; I examine the case where a user is performing tasks only within a category j, and the performance rating on these tasks is strictly binary, either "good" or "bad". Given a past history of n tasks within the given category, and assuming that the current quality q_{ij} of the freelancer i in category j is known, I expect the number x of completed tasks rated as "good" to follow a binomial distribution. By using basic concepts of Bayesian statistics (Gelman et al. 2004), I infer q_{ij} based on the number of "good" and "bad" completed tasks (assuming a $Beta(\alpha,\beta)$ prior, we get a posterior $Beta(\alpha+x, n-x+\beta)$)

In practice, binary feedback is typically used for small tasks (e.g., on Amazon Mechanical Turk). For more complex tasks, we often see reputation systems that have multiple grades for feedback (*e.g.* 5-star ratings are common). To extend the previous model to account for a range of discrete outcomes, I use a multinomial distribution of K possible outcomes (instead of just two). By assuming a Dirichlet prior, the posterior is also Dirichlet, from which I draw an estimate for the freelancer's quality in a given category.

The implicit assumption so far was that the category-specific user qualities do not evolve over time. In reality, we would expect a more dynamic freelancer behavior: As users complete more and more tasks, their more recent tasks become more predictive than their initial ones. To capture this evolutionary behavior, I propose a linear dynamical system (LDS) (Bishop and others 2006). I assume that both the quality of the user q and the received feedback f follow normal distributions, whose means are linear functions of the states of their parents in the graph. As before, the goal here is to estimate the latent quality q based on the observed feedback f. Assuming that the freelancer at hand completes N tasks in the same category, the following equations hold for f and q:

$$p(q_1) = \mathcal{N}(\mu_0, p_0), \qquad p(q_n|q_{n-1}) = \mathcal{N}(aq_{n-1}, g)$$
$$p(f_n|q_n) = \mathcal{N}(cq_n, r), \qquad p(q_{n-1}|f_1, ..., f_{n-1}) = \mathcal{N}(\mu_{n-1}, v_{n-1})$$

I use this equations and estimate the next observed outcome by drawing from the conditional distribution $p(f_n|q_n)$.

Learning across categories: In practice, there is chance for the within-category history to be insufficient (*i.e.* very few data points); To overcome this, I propose that even if a freelancer have no experience in a given category, the past experience in some other, correlated categories might be predictive of future performance. Hence, the quality of a freelancer i for a category j (q_{ij}) can be estimated based on the knowledge of the history and values q_{ik}, for other categories (also see (Clemen and Winkler 1990)):

$$\mathrm{logit}(q_{ij}) = \sum_{k=1}^{m} \alpha_{jk}\mathrm{logit}(q_{ik}) + \varepsilon_{ij} \tag{1}$$

where α_{jk}, β_j are data-specific coefficients, ε_{ij} is a random disturbance, and logit is the standard logit function.

Baseline Models: To evaluate the performance of the proposed approaches I use two different baselines. The first one, averages the past reputation of the freelancers across categories. The second one draws on recommender systems, and predicts the outcome based on freelancers' similarities (user-user collaborative filtering (CF) (Shapira 2011)).

Study 2:Career Development Paths

The focus of this study is on estimating the value of each skill (or set of skills) given a level of expertise of the freelancer. The ultimate goal is to propose a framework that takes into account the current skillset of a freelancer and recommends new skills to be learned to maximize the freelancer's value in the marketplace.

The actual level of expertise of a given freelancer and a given skill is latent. However, in an *OLM* we observe multiple signals of this level of expertise (*e.g.* certifications, hiring rate, re-hires, profile information, reputation) . Given all these signals, our goal is to estimate the latent level of expertise. To achieve this, I propose to build a Hidden Markov Model (*HMM*). In particular, I assume that a skill's expertise is dynamic, and evolves over time. Next, I define a set of different discrete unobserved states of expertise. These states emit with different probabilities observations in a predefined observations set. Based on the emitted observations, I can estimate the conditional probability of the freelancer being at anyone of the available states.

Once I know the level of expertise of a freelancer, I can define the conditional utility of a skill. In particular, given a skill s, a freelancer with level of expertise on that particular skill $E_s = e$, and the number of jobs in the marketplace that require this skill D_s, I propose that the conditional utility of

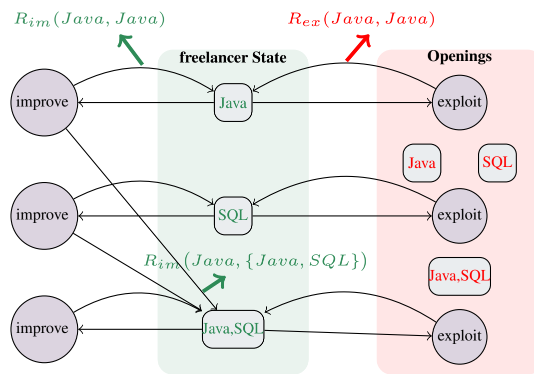

Figure 1: Markov Decision Process example.

this skill is given by the following:

$$\begin{aligned} U_{h|E_s=e,D_s} &= [\bar{W}_{E_s=e} - \bar{C}_{E_s=e}] \\ &- \frac{1}{|E_s|-1} \sum_{e'\in\{E_s-e\}} [\bar{W}_{E_s=e'} - \bar{C}_{E_s=e'}] \\ &= \bar{W}_{E_s=e} - \frac{1}{|E_s|-1} \sum_{e'\in\{E_s=e'\}} [\bar{W}_{E_s=e'}], \end{aligned} \tag{2}$$

where $\bar{W}_{E_s=e} = (\sum_{o\in O_{E_s D_s}} W_o)/(|O_{E_s,D_s}|)$ represents the average hourly wage on openings (o) that require the given skill s and hired a freelancer with level of expertise $E_s = e$, while $\bar{C}_{E_s=e'}$ is the average cost of freelancer's hourly effort at level of expertise e. The main assumption of the previous equation is that the cost of effort remains the same across levels of expertise and different skills. I further discuss this assumption in the last section of this proposal ('Challenges to discuss').

Once I estimate the utility of each skill (set of skills) given some level of expertise, I can proceed and build a framework for recommending skills. In particular, the question I study here focuses on freelancers' optimal decision making: at any point, freelancers have two options: (1) to exploit their current skillset and expertise by getting hired and completing a task, and (2), invest their time on improving/expanding their skillset and expect future increased returns. What is the optimal decision for each freelancer? Exploit or improve?

I propose to model this problem as a Markov Decision Process (MDP): each freelancer state will be a set of skills that a freelancer has some level of expertise. At every step of this process, freelancers make a decision: **exploit** their current state and apply for an opening, or try to **improve** their skillset and transition to a different state. Each decision leads to stochastic results with different expected rewards. We denote these rewards as $R_a(s, s')$, where $a \in$ {"improve","exploit"}, s is the current state, and s' the state that the freelancer transitions to after taking an action a. I present an example of this process in Figure 1.

Study 3: Hiring Decisions

The focus of this study is to understand hiring decisions in *OLM*s, and propose models that will decrease friction and benefit all involved parties (freelancers, employers and the marketplace). Specifically I propose a series of increasing

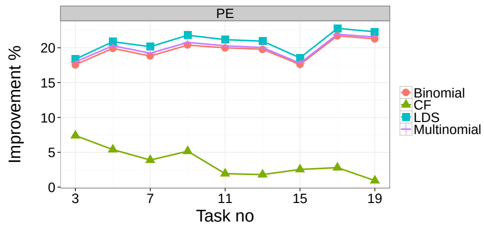

Figure 2: Study 1: Mean Absolute Error *improvement* over the average feedback baseline.

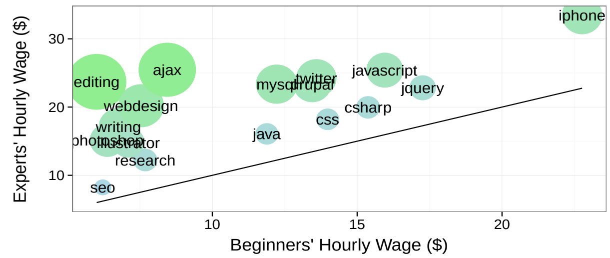

Figure 3: Utility of Skills

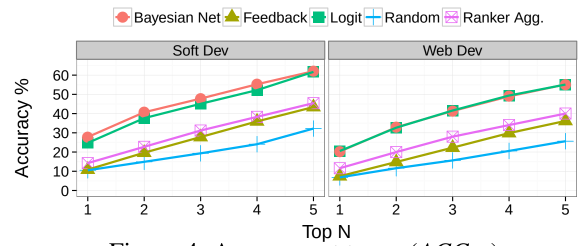

Figure 4: Accuracy at top-n (*ACC-n*).

complexity predictive models that describe employers' hiring decisions. I assume that employers are rational utility maximizers; Their utility is straightforwardly maximized along with the probability of selecting the best possible applicant for each specific opening. Based on this assumption, I first propose a ranking aggregator that ranks candidates in all the available dimensions and then aggregates these ranks to create a global ranking. Next, I draw on empirical economics and propose a Logit model and finally, I built a probabilistic graphical model (Bayesian network). Both the Logit model and the Bayesian network estimate the conditional hiring probability of an applicant given the applicant's, the employer's and the opening's characteristics. Next I compare these models with the vanilla reputation score baseline, where each employer ranks the available applicants based on their previously collected feedback score. I further perform an econometric analysis and a propensity score study and observe that the attributes that have the strongest positive effect on hiring probability are whether or not the freelancer and the employer have previously worked together, the available information on the freelancer's profile, the countries of the employer and the freelancer and the skillset of the freelancer. Finally, my analysis shows that the faster the freelancer applies to an opening, the higher is the probability to get hired.

Results

Study 1, Reputation Transferability: In Figure 2, I show the Mean Absolute Error (MAE) percentage improvement over the feedback baseline for all the proposed approaches (Binomial, Multinomial, LDS and Collaborative Filtering (CF)). The baseline is at zero (on the x-axis), and every positive value is an improvement. All approaches perform significantly better than the baseline, providing an improvement of up to 25%. Furthermore, the LDS performs better than the Multinomial, which performs better than the Binomial, which in turn performs better than Collaborative Filtering[1].

Study 2, Career Development Paths: For this study, I present only initial results for the first part, which deals with quantifying the utility of a skill given the freelancer's level of expertise (see Equation 2). In Figure 3, I show the utility of a set of skills in my dataset. I consider only two levels of expertise: 'expert' and 'beginner'. On the y-axis it's the experts' hourly wage, and on the X-axis, it's the beginners' hourly wage. The line is the 'equality line' (45 degrees), indicating that any point above that line has a positive utility. First, note that all the presented skills have positive utility, verifying intuition. Simply put, this shows that a freelancer has more value in the marketplace if the freelancer is an expert on a skill, than a beginner. Second, we observe that for some skills, the estimated utility is higher (e.g., 'editing','ajax', 'mysql', 'javascript', 'iphone' *etc.*) than others (*e.g.* seo, research, java, css, c#).

Study 3, Hiring Decisions: Recall that in this study I build ranking models that estimate the hiring probability of an applicant. In Figure 4 I show the results of Accuracy at top-n (*i.e.* whether the predicted ranks include a true positive (a hire) in the 1^{st}, 2^{nd}, ...,n^{th} position) for two of the six categories in my dataset ('Soft Dev' and 'Web Dev'). The Bayesian approach and the Logit model perform much better than the ranker aggregator and the feedback baseline in all categories. The actual probabilities of predicting hiring decisions are impressive; In the "Software Development" category, the Bayesian approach ranks the applicants in a way that 28% of the times, the applicant who ranks first is the one who gets hired. At random, this probability is 10%.

Towards a more efficient marketplace

All three studies have a series of very important implications for the marketplace and its users. My first study shows a clear and methodologically sound approach for analyzing the correlations between different task categories, and as a result, it provides a more accurate estimate of a freelancer's performance in a new category. This information is valuable among employers that participate in online labor markets, allowing them to make safer and better informed hiring decisions. Furthermore, the proposed approach provides a guideline for many other labor marketplaces. For example, TaskRabbit or LinkedIn can leverage this approach to infer correlations among job types. Even online marketplaces such as Ama-

[1] You can see an updated version of all studies in the following links: Reputation transferability, Career development paths and Hiring decisions

zon.com can use our approaches to improve the reputation scores displayed for merchants that are active across multiple product categories, and analyze the abilities of Amazon.com reviewers to provide helpful reviews across different product categories.

My second study provides a virtual assistant for developing new skills or improving current ones. By knowing the actual value of being an expert on a skill, and by taking into account the skill's demand and supply, the marketplace can strategically recommend skills (improve) or jobs (exploit). If freelancers start acquiring skills with high utility, both their demand and their income will increase. Since more freelancers will have skills that have high demand in the marketplace, both the job openings' closing rate and (as a result) the marketplace's revenue will increase.

Finally, regarding my third study, by developing approaches that estimate the applicants' hiring probabilities: (1) employers will be able to make better-informed and faster decisions based on the suggested applicants' rankings, (2) freelancers will save time by not applying to openings that have very low hiring probability and (3) the marketplace might identify weaknesses in freelancers' profiles (*e.g.* skills not reported, the profile description is not sufficient *etc.*) and suggest targeted profile improvements based on how each profile characteristic affects hiring decisions. As a result, the proposed approaches will create a more efficient marketplace, and increase both the marketplace's transaction volume as well as the overall satisfaction of the freelancers and the employers.

Challenges to discuss

My first study is almost complete. Any feedback about the approaches, the problem, or interesting extensions will be welcome and very appreciated. The other two studies are ongoing, and there are certain issues I would like to discuss. In particular for my second study, I would like to get feedback on the utility definition (Equation 2) and in particular, about my assumption regarding the cost of effort. I would also like to discuss ideas for estimating the exploit rewards (R_{ex} in Figure 1) of each freelancer state.

For the analysis I conducted in my third study, I only considered openings that lead to a single hire. As a result, my dataset does not capture the behavior of employers who decide (for whatever reason) not to hire anyone of the available applicants. As a result, the proposed approaches should not be used as-is to recommend potential candidates (*i.e.* freelancers that haven't applied yet) to employers. I intend to extend this work by incorporating openings that remain unfilled, and study why employers choose to not hire anyone of the available freelancers. I intend to use this information to built a framework that will recommend high quality freelancers to new openings. Once I have a complete model that provides good estimates of the hiring probability of each *active freelancer* (not applicant) on each *active opening*, I am planning to work towards maximizing the closing rate of openings in the marketplace. In particular, assume that we are given a bipartite graph, where the edges represent current applications from the available freelancers. How can we reallocate these edges in order to maximize not just the filling rate of jobs, but also, the expected payoff (*e.g. OLM* revenue) from the available openings. I would really appreciate the committee's feedback regarding these extensions.

References

[Acquisti and Fong 2013] Acquisti, A., and Fong, C. M. 2013. An experiment in hiring discrimination via online social networks. *Available at SSRN 2031979*.

[Agrawal et al. 2013] Agrawal, A.; Horton, J.; Lacetera, N.; and Lyons, E. 2013. *Digitization and the Contract Labor Market: A Research Agenda*. University of Chicago Press.

[Bishop and others 2006] Bishop, C., et al. 2006. *Pattern recognition and machine learning*, volume 4. springer New York.

[Clemen and Winkler 1990] Clemen, R. T., and Winkler, R. L. 1990. Unanimity and compromise among probability forecasts. *Management Science*.

[Desmarais, Maluf, and Liu 1995] Desmarais, M. C.; Maluf, A.; and Liu, J. 1995. User-expertise modeling with empirically derived probabilistic implication networks. *User modeling and user-adapted interaction* 5(3-4):283–315.

[Forsythe, Drake, and Cox 1985] Forsythe, S.; Drake, M. F.; and Cox, C. E. 1985. Influence of applicant's dress on interviewer's selection decisions. *Journal of Applied Psychology* 70(2):374.

[Gelman et al. 2004] Gelman, A.; Carlin, J.; Stern, H.; and Rubin, D. 2004. *Bayesian Data Analysis*. Chapman & Hall/CRC.

[Gilmore, Beehr, and Love 1986] Gilmore, D. C.; Beehr, T. A.; and Love, K. G. 1986. Effects of applicant sex, applicant physical attractiveness, type of rater and type of job on interview decisions*. *Journal of Occupational Psychology* 59(2):103–109.

[Hambleton 1991] Hambleton, R. K. 1991. *Fundamentals of item response theory*. Sage Publications, Incorporated.

[Hu 2003] Hu, L. 2003. The hiring decisions and compensation structures of large firms. *Industrial and Labor Relations Review* 663–681.

[Ipeirotis, Provost, and Wang 2010] Ipeirotis, P. G.; Provost, F.; and Wang, J. 2010. Quality management on amazon mechanical turk. In *Proceedings of the ACM SIGKDD workshop on human computation*, 64–67. ACM.

[Jurczyk and Agichtein 2007] Jurczyk, P., and Agichtein, E. 2007. Discovering authorities in question answer communities by using link analysis. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, 919–922. ACM.

[Kawakami, Dovidio, and van Kamp 2005] Kawakami, K.; Dovidio, J. F.; and van Kamp, S. 2005. Kicking the habit: Effects of nonstereotypic association training and correction processes on hiring decisions. *Journal of Experimental Social Psychology* 41(1):68–75.

[Kokkodis and Ipeirotis 2013] Kokkodis, M., and Ipeirotis, P. G. 2013. Have you done anything like that?: predicting performance using inter-category reputation. In *Proceedings of the sixth ACM international conference on Web search and data mining*, 435–444. ACM.

[Mason and Watts 2010] Mason, W., and Watts, D. J. 2010. Financial incentives and the performance of crowds. *ACM SigKDD Explorations Newsletter* 11(2):100–108.

[Nelson 1970] Nelson, P. 1970. Information and consumer behavior. *Management Science.*

[Pallais 2012] Pallais, A. 2012. Ineffiient hiring in entry-level labor markets. *Available at SSRN 2012131*.

[Rand 2012] Rand, D. G. 2012. The promise of mechanical turk: How online labor markets can help theorists run behavioral experiments. *Journal of theoretical biology* 299:172–179.

[Saito et al. 2014] Saito, S.; Watanabe, T.; Kobayashi, M.; and Takagi, H. 2014. Skill development framework for micro-tasking. In *Universal Access in Human-Computer Interaction. Universal Access to Information and Knowledge*. Springer. 400–409.

[Shapira 2011] Shapira, B. 2011. *Recommender systems handbook*. Springer.

[Yakubovich 2005] Yakubovich, V. 2005. Weak ties, information, and influence: How freelancers find jobs in a local russian labor market. *American sociological review* 70(3):408–421.

[Zhang, Ackerman, and Adamic 2007] Zhang, J.; Ackerman, M. S.; and Adamic, L. 2007. Expertise networks in online communities: structure and algorithms. In *Proceedings of the 16th international conference on World Wide Web*, 221–230. ACM.