What's the Inside Scoop? Challenges in the Supply and Demand for Information about Job Attributes*

Jason Sockin
University of Pennsylvania

Aaron Sojourner
University of Minnesota

September 11, 2020

Abstract

Workers struggle to understand prospective jobs and employers. Glassdoor is an online platform that offers jobseekers information about prospective employers from other workers' volunteered reviews. Analyzing Glassdoor data reveals how jobseekers share and use this information. Jobseekers rate reviews of employers more helpful if they contain more-negative information, but such information is relatively scarce. Volunteers supplying negative information are more likely to conceal aspects of their identity, degrading the supplied information's value. Concealment is more likely in reviews for smaller firms and from current employees, where retaliation risk is higher. While workers demand information about some workplace attributes more than others, supply and demand for such information is imbalanced. Across firms, not all hard-to-observe yet desirable attributes improve with easier-to-observe pay, providing rationale for why jobseekers value firm-specific information. Reputation institutions provide valuable but partial solutions to workers' information problems.

JEL: J3, J28, D83, J62.

Keywords: Imperfect Information, Job Search, Job Satisfaction.

*Thanks to Alan Benson, Jonathan Borowsky, Andrew Chamberlain, Laura Giuliano, John Horton, Suresh Naidu, Tara Sinclair, Bryan Stuart, Abigail Wozniak, and Daniel Zhao, as well as seminar participants at UMaryland, UMinnesota, and UPenn for productive discussions about these ideas. The paper uses data from Glassdoor, where Sockin previously worked. Glassdoor has no control rights over the manuscript.

1 Introduction

Harnessing data from an institution that facilitates the flow of information between workers about job quality and job attributes, Glassdoor, we develop evidence to better understand the nature of workers' supply of this kind of information, their demand, and potential imbalances. Broadly, Glassdoor is similar to other third-party review sites like Yelp or Tripadvisor, except it helps workers share information about employers rather than consumers share information about sellers. They all rely on volunteers to supply their private information toward the public good of publicly-available reputation for counter-parties.

First, we highlight a challenge in promoting the flow of information about employers, that concern about potential retaliation may dissuade workers with negative information from sharing it. Evidence in Marinescu et al. (2018) is consistent with this story. If a jobseeker browses volunteers' reviews long enough, Glassdoor limits access to additional reviews until after the jobseeker contributes their own review (a Give-to-Get policy). For this group, average firm ratings were lower than for others. The marginal rating is more negative than the average rating. This is consistent with positive bias in volunteered reviews relative to the population of potential reviews and potentially higher implicit cost to workers of volunteering more negative reviews.

Our analysis provides novel evidence that volunteers are more likely to conceal aspects of their identity with their reviews—their job title and location—when reporting more negative information about a job at a particular firm than when reporting more positive information. This is consistent with concern about employer retaliation. Corroborating evidence comes by checking whether, conditional on supplying negative information, volunteers were more likely to conceal aspects of their identity when they face higher retaliation risks, either reporting on a current (rather than former) employer or working at a smaller firm where they blend into a smaller pool of likely suspects. Even looking across multiple reviews from the same volunteer, the person is more likely to conceal aspects of their identity when leaving a more negative review, when leaving a review about a current rather than former employer, and when reviewing a smaller rather than larger employer. Further, the (inverse) relationship between review rating and rates of identity concealment is stronger in contexts where there is a greater risk of retaliation.

The concealment of identifying information from volunteers tends to degrade the value of the supplied information to jobseekers, suggesting a fundamental challenge for information flow in the labor market. Jobseekers tend to vote reviews less helpful when volunteers conceal aspects of their identity, specifically their job title or location, suggesting job seekers are particularly interested in information specific to their labor market preferences. Taken together, this suggests a challenge in the design of information institutions. Without a market-clearing mechanism, there's little reason to expect the supply and demand for private, firm-specific information to meet.

While Glassdoor, Yelp, TripAdvisor, and many other public reputation systems are built on similar principles, the stakes differ in the labor market. On the information-supply side, workers with negative information have little incentive to reveal it. Worker supplying accurate, private information about an employer creates a positive externality on other workers and it will tend to be undersupplied via volunteerism. Weil and Pyles (2005) model and discuss these externalities in regards to workers' decisions to report employer violations of their rights to public enforcement agencies and many of the same dynamics are at play here. However, an employee deciding to report a rights violation to an enforcement agency may be motivated by the hope of recovering damages, but the prospect of that benefit is not relevant in the decision to share information with other workers so the incentive to supply may be even weaker. The externality on firms will be mixed and depend primarily on if the information will tend to increase or decrease labor supply to the firm. This gives firms incentives to intervene in the information supply process. The employer that would be hurt by the negative information has a concentrated interest in the worker not supplying it, whereas the set of workers and competitors with any interest in the information coming out each have only a minor interest, thereby making it difficult to coordinate (Olson, 2009).

When considering whether to volunteer a negative review, potential volunteers may worry much more about retaliation risk from an employer rather than from a restauranteur or hotelier. If a consumer gives a bad review to a restaurant on Yelp or to a printer on Amazon, they can neglect almost any risk of negative repercussions. The customer's always right. The seller is unlikely to know enough to identify the consumer or to care enough to try to impose costs. A very small piece of any consumer's economic life is at stake. Employment is different.[1] Workers deciding whether to volunteer negative information about an employer may feel it risky to do so. If enough workers believe the risk is real, negative information will be under-supplied. Is it worth the possibility of antagonizing your boss, losing your job, and causing a lifetime of retaliatory references in order to help jobseekers—whom you do not know—make more informed employment decisions? Firms have many ways to retaliate against current employees. And while they have fewer ways to retaliate

[1]As Hart (1989) puts it, "the reason an employee is likely to be more responsive to what his employer wants than a grocer is to what his customer wants is that the employer has much more leverage over his employee than the customer has over his grocer. In particular, the employer can deprive the employee of the assets he works with and hire another employee to work with these assets, while the customer can only deprive the grocer of his custom and as long as the customer is small, it is presumably not very difficult for the grocer to find another customer."

against former employees, a past firm has the power to deter a worker's career.[2] There is anecdotal and survey evidence that workers risk employer retaliation when sharing negative information about working conditions. Women defying pressure and retaliation threats from their bosses to share information about sexual harassment at work gave rise to the Me Too movement. Firms use broad interpretations of nondisclosure and nondisparagement clauses to threaten lawsuits to prevent workers from sharing negative, private information about jobs there.[3] Over a third of U.S. workers report being bound by nondisclosure agreements (NDAs) (Starr et al., 2019). According to Lobel (2019), "NDAs regularly include information beyond traditionally defined secrets under trade secrecy laws, including general know-how, skills, client lists, and salary information. They also include provisions prohibiting the employee from disparaging the company." Anecdotal evidence of employers threatening retaliation against workers for disclosing negative information about working conditions is easy to find.[4] Employers' interest in suppressing negative information and willingness to retaliate against employees who share it provide the rationale for many whistle-blower protections laws and procedures (Weil and Pyles, 2005; U.S. Occupational Safety and Health Administration, 2017). Systematic evidence about the chilling role of retaliation fears comes from a recent study showing that stronger protection of workers who blow the whistle on rights violations by employers increases the willingness of workers to report such violations (Johnson et al., 2020).

Second, we establish that jobseekers do not value new information that guides them towards better, higher-rated employers similarly as that which guides them away from worse, lower-rated ones. Jobseekers do not value information across the full spectrum of quality equally. A relative under-supply of negative information would predict that jobseekers

[2] According to Weil and Pyles (2005), "Public law groups and other organizations representing low-wage workers note that many employee complaints... are filed after a worker has been fired by an employer, often for other causes (thereby lowering the cost of complaining at that point)."

[3] Silver-Greenberg and Kitroeff (2020) say, "Employees who are fired or resign in frustration are often pushed to sign contracts that prohibit them from in any way disparaging the company, several of the former employees said in interviews. Those pacts bar the employees from even acknowledging the existence of the agreements.... The Times spoke to 13 former Bloomberg employees... who said they wanted to be released from their exit agreements so that they could speak openly about the culture at the company... If they were free to talk, some of the former employees said, they would describe a company that, while it provides generous pay and benefits, can be an uncomfortable place to work, especially for women." Benner (2017) says, "Nondisparagement clauses are not limited to legal settlements. They are increasingly found in standard employment contracts in many industries, sometimes in a simple offer letter that helps to create a blanket of silence around a company."

[4] Eidelson (2020) writes, "In the past few months, U.S. businesses have been on a silencing spree. Hundreds of U.S. employers across a wide range of industries have told workers not to share information about Covid-19 cases or even raise concerns about the virus, or have retaliated against workers for doing those things, according to workplace complaints filed with the National Labor Relations Board (NLRB) and the Occupational Safety and Health Administration (OSHA)."

express a stronger demand for negative information about prospective jobs than positive information. This contrasts with findings from consumer review sites and thus does not appear to be a fundamental property of reputational information flow. Hong et al. (2017) conduct a meta-analysis of the factors that contribute to product review helpfulness and find that the review rating has no significant bearing on the helpfulness of product reviews. We know of no prior investigations of what constitutes helpful information in employer reviews.

The finding that jobseekers vote reviews delivering more negative information more helpful than reviews delivering positive information, holds using many uni-dimensional, vertically-differentiated measures of how positively/negatively volunteers evaluate their firms, such as the one-to-five star overall rating, the share of Pros versus Cons text written in the review, and whether the volunteer would recommend the firm to a friend. This stronger relative demand for negative information holds across firms of different ages, sizes, review counts, and pay premia. Where we can precisely track an individual volunteers' votes, we find that the increased demand for more negative information holds true regardless of the volunteers' experience with their own employer, highlighting the universal desire for avoiding a future match with a low-rated employer. The only exception we found is when the jobseeker votes for a review of their own employer: those who rate their own employer highly also tend to vote positive reviews of their employer as more helpful. This could be sincere or strategic.

Moving beyond uni-dimensional ratings, we find that jobseekers demand for more negative information is not restricted to concern about a few attributes but holds true across all job attributes. It's not that jobseekers find negative reviews more helpful because they tend to focus on particular job attributes about which jobseekers care more. To do this, we introduce novel measures of volunteers' evaluation of a reviewed job on a set of hard-to-observe attributes by exploiting the structure of volunteers' text reviews in separate Pros and Cons fields and modeling latent attributes of jobs described therein. This decomposition of text into attributes provides a unique way to capture how positive or negative a volunteer assesses their job along many dimensions.[5] Jobseekers find reviews delivering more negative information on every attribute more helpful conditional on the review's evaluation of all other attributes.

We do not claim that an under-supply of negative information is the only factor that

[5]Other research has explored Glassdoor reviews and used them to conduct text-based topic analysis. Marchetti (2019) applies a similar latent Dirichlet allocation (LDA) model for partitioning the text of Glassdoor reviews to explore how firm culture relates to the synergy of mergers and company acquisitions but does not exploit the Pro/Con structure to generate evaluations. While Marchetti (2019) focuses on the interpretation of these topics, our analysis refrains from assigning interpretations to topics, exploiting the LDA algorithm as simply an effective method by which to partition text into latent groups of related text.

drives the premium jobseekers place on it. Workers tend to be risk averse and as such, will value a signal that shifts their posterior belief about the expected value of a potential job down by ϵ more highly than a signal that shifts it up by ϵ. In addition, positive information about a potential employer may be available outside a reputational information system. In fact, we establish that the text description within job postings more closely resembles the text of positive employer reviews than negative ones, a new though perhaps not surprising finding. Given employer preference for supplying information and worker risk aversion, it would be surprising if jobseekers expressed no preference for negative information from volunteers even without volunteers censoring or degrading of negative information. But volunteers' tendency to do so likely contributes to making jobseekers' even hungrier for this kind of information.

Third, jobseekers' preference for vertically-differentiating, negative information does not diminish when the model includes the potential for reviews to provide information that horizontally differentiates in attributes as well. In a subsample where we can observe individual jobseekers' votes and their characteristics, we form jobseeker-review pairs. We measure the correlation between (1) what attributes each jobseeker cares about as measured by the emphasis they devote to it in their own job reviews and (2) what attributes are emphasized in the reviews on which the jobseekers vote. Neither (1) nor (2) uses vertically-differentiating information, just the shares of attention devoted to each attribute. When there's a stronger positive correlation between what the jobseeker cares about and what the review emphasizes, the probability that the jobseeker votes the review helpful increases. However, inclusion of this channel does nothing to diminish the importance of vertically-differentiating information in the empirical model. Jobseekers seem to value both vertically- and horizontally-differentiating information.

Fourth, we estimate the relative importance of job attributes in workers' average preferences, in order to understand what kinds of information they would value more. Consider the relative importance of attributes in explaining jobseeker ratings of reviews as helpful or unhelpful, as described in the discussion of the second contribution above. We interpret these coefficients as containing information about the relative value jobseekers place on the attributes, given that the value of information about an attribute should grow with the change in expected utility jobseekers derive from improving their understanding of the likely level of that attribute at a prospective firm. We explore this interpretation by comparing the coefficients against those derived from a model of volunteers' job satisfaction as a function of her own evaluation of her job's attributes.

This model is very similar conceptually to that used by Maestas et al. (2018) and Mas and Pallais (2017), though aspects of design differ.[6] Among two different set of workers making

[6]For our own work and that of Maestas et al. (2018), preferences are expressed on an ordinal scale as

two different kinds of decisions but each of which hinges on the worker's preferences over attributes, we find a very high level of correlation in relative importance across attributes. Given this, we use these as a measure of relative demand for information across attributes.

Next, we compare relative supply and demand for information across attributes to assess balance. The relative attention that the corpus of volunteered reviews devotes to each attribute serves as a measure of relative supply. Attribute coefficients from the jobseeker helpfulness model serves as a measure of relative demand. We find that the most-demanded half of attributes account for 64 percent of total demand weight but only 52 percent of supply weight. Again, there is little reason to expect balance without a price mechanism but this does highlight and quantify a problem with the flow of information about jobs.

Lastly, we find that the attribute about which jobseekers demand information most strongly is an attribute that both strongly drives worker job satisfaction (high marginal utility) and tends to be negatively correlated with firms' pay premia, meaning that easy-to-observe pay is a bad proxy for this hard-to-observe attribute. Extracting employer fixed effects from an Abowd et al. (1999) (AKM)-style approach to pay and attributes, we document variation in attributes across firms and relate firm fixed effects on pay to their fixed effects on satisfaction and each job attribute. A one standard deviation increase in firm-specific wage premium correlates with a roughly one-sixth standard deviation increase in the employee's overall rating of the firm. Many but not all firm job attributes are increasing with pay, consistent with the operation of both Rosen and Mortensen motives (Sorkin, 2018).

2 Literature

Workers face information problems in choosing between employers that appear the same but actually differ in unobservable ways. Jobs are complex, difficult to fully characterize and subject to change. Non-wage attributes matter greatly in determining the value of jobs to workers (Mas and Pallais, 2017; Maestas et al., 2018). An incomplete list of job attributes include wages; aspects of health insurance quality and cost; criteria for and schedule of potential raises; opportunities for career development and advancement; risk of occupational illness, injury, or fatality; degrees of autonomy and micro-management; personal and professional (dis)courtesy paid by one's supervisor and peers; presence of sexual harassers; layoff risk in a downturn; whether one is routinely asked to work overtime and what consequences

a function of job attributes. They have survey outcomes from a nationally-representative sample about pairs of hypothetical jobs with randomly-assigned attributes. We have survey outcomes from a large, convenience sample about pairs of real jobs that they have held. Mas and Pallais (2017) offered real jobs with randomly-assigned attributes to a convenience sample and observed choices. They both used specified, explicit attributes. We use latent attributes derived from analysis of review text.

would follow from refusal; and ease in scheduling time off to take a child to the doctor. Attributes' starting levels and possible future paths can matter.

Given this complexity, jobseekers value information that helps them better understand how a prospective employer treats workers, especially in hard-to-observe and hard-to-contract dimensions. Nelson (1970) described experience goods as those whose quality cannot be learned before transaction and evidence points to jobs having a large experience good component. He wrote that, for experience goods, "Information about quality differs from information about price because the former is usually more expensive to buy than the latter." For jobs, information about the complex bundle of amenities is more difficult to acquire than information about wages. Menzio and Shi (2011) develop a directed search model that permits analysis of the extent to which prospective jobs are inspection goods, where match quality and the value of a prospective job is known to the jobseeker before accepting it, versus experience goods, where the jobseeker cannot tell the job's type before deciding whether to accept. They find that the experience-good model better explains cyclical labor-market dynamics and estimate that this difficult-to-observe match quality explains a huge share of variance in job productivity. Moving from the 10^{th} to the 90^{th} percentile of match quality almost doubles job productivity, suggesting significant scope to help workers better recognize match quality ex ante.

Workers' private information about jobs at a particular firm does not flow easily to other workers. Price mechanisms to govern information flow are limited. Hungry to understand the job quality one can expect from a given firm, jobseekers seek the inside scoop from current and former employees of the firm—who have private and otherwise inaccessible knowledge of how the firm treats its staff. Carmichael (1984) developed a theory of employer reputation in the labor market. He argues, "Since a searching worker does not typically get to observe a firm very closely before he joins it, it does not seem sensible to assume he has intimate knowledge of its technology or the tastes of its owners." He points to an employer's public reputation as a way jobseekers can deal with this information problem because "if the worker himself is not very different from other workers the firm has hired in the past, then he may do very well just by assuming that the firm will treat him as it has treated everyone else." Empirically, Brown and Matsa (2016) found that jobseeker applications fall as public information about the prospects for firm survival diminish and, in the online labor market Amazon Mechanical Turk, Benson et al. (2020) provide the first field experimental evidence that an employer with a better crowdsourced, public reputation enjoys increased labor supply.

The flow of experienced information about jobs is not only suppressed due to a lack of markets and mechanisms. Even when presented with the opportunity to do so, employees simply may not be willing to volunteer such information, especially when the information is damaging to the employer and there are real risks of retaliation. Cortina and Magley (2003) found in a survey of public-sector employees, where incentives for managers to retalitate may be weaker than in the private sector, that only 27 percent of respondents who experienced some recent interpersonal mistreatment in the workplace voiced concern over their mistreatment. Among those who did, 66 percent reported being the subject of work-related or antisocial retaliatory behavior. Workers who report workplace violations might face retaliation that lowers their income or sours their job satisfaction vis-a-vis reduced pay, fewer hours, task reassignment, etc.[7] The consequences though can be even more dire, resulting in more extreme responses such as deportation.[8] Although volunteers are told that their review will be posted anonymously on Glassdoor, still respondents seek out additional layers of anonymity by concealing potentially identifying information. We contribute to the literature on worker disclosure by adding an empirical foundation for workers' concerns with firm retaliatory behavior—as volunteers reporting negative sentiment are more likely to conceal identifying aspects of their review when it is easier for them to be identified by their employer.

Better evidence on workers' use of information about job quality and attributes would improve our understanding of labor markets but these processes have been difficult to measure until recently. Due primarily to data limitations rather than the credibility of the assumption, compensating-differential models typically assume workers know with certainty attributes at prospective jobs (Rosen, 1986; Hwang et al., 1998; Sorkin, 2018), implicitly imposing a subjective distribution that's degenerate at the truth. This abstracts away challenges workers face in learning about jobs and the operation or design of institutions that may affect workers' beliefs or information sets.[9] The introduction of online platforms such as LinkedIn, Glassdoor, and Indeed has enabled workers to search for, learn about, and apply to jobs online. Before these platforms existed, workers had to rely on ads, word of

[7] Covert (2020) discusses retaliation against McDonald's workers reporting sexual harassment, "Instead of the harassers facing discipline, punishment was often meted out to the victims... assigned difficult or uncomfortable tasks—working the grill all day or being stuck at the drive-through window for an entire shift. She was also disciplined for minor infractions, had her hours cut, was demoted, and got suspended for two weeks. She was eventually fired." This is an awful example of workers expressing dissatisfaction with workplace conditions and bearing the negative consequences from their current employer for doing so.

[8] For example, when undocumented workers in Minnesota complained to their employer about working conditions and said they would complain to others, the employer reported them to U.S. Immigration & Customs Enforcement, which deported them (Walsh, 2018; Chen, 2018).

[9] Economists have devoted considerable attention to firms' strategies for dealing with information problems in distinguishing between workers who appear the same but differ in unobserved type (Spence, 1976; Altonji and Pierret, 2001) or effort (Shapiro and Stiglitz, 1984). Personnel economics has taught us a lot about how firms use performance contracts, educational credentials, referral-based hiring, and many other strategies to deal with their informational challenges (Oyer et al., 2011). The analogous problem for workers—to distinguish between employers that look the same but differ in unobserved ways—has received less attention.

mouth, personal networking, and best-places-to-work stories in magazines to learn about potential employers. The movement toward online job search has changed how workers gather, process, and act on job-related information and what economists can observe.

3 Setting and data

A volunteer can submit one employer review per employer-year, which she is free to update, but can review multiple employers within the same year.[10] Each review-r pertains to a job at a specific firm-f. Each has an associated volunteer $v(fr)$ and creation time $t(fr)$. In a separate type of report, workers can report their salary at a firm.

Figure 1: Blank Volunteer Review Form

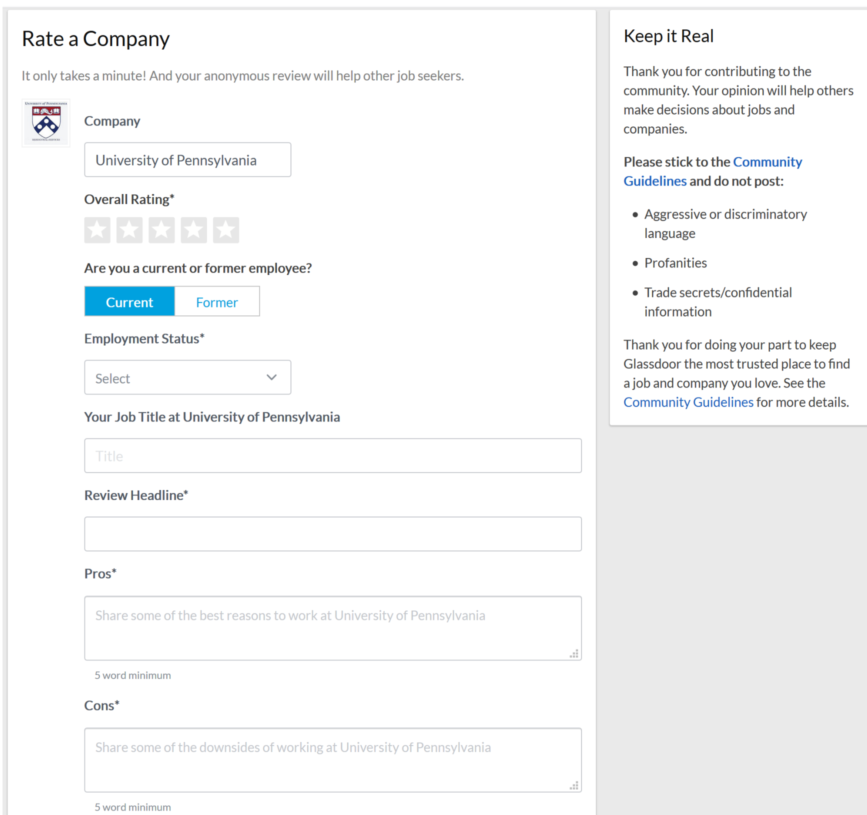

Rate a Company

It only takes a minute! And your anonymous review will help other job seekers.

Company

University of Pennsylvania

Overall Rating*

Are you a current or former employee?

Current | Former

Employment Status*

Select

Your Job Title at University of Pennsylvania

Title

Review Headline*

Pros*

Share some of the best reasons to work at University of Pennsylvania

5 word minimum

Cons*

Share some of the downsides of working at University of Pennsylvania

5 word minimum

Keep it Real

Thank you for contributing to the community. Your opinion will help others make decisions about jobs and companies.

Please stick to the Community Guidelines and do not post:

- Aggressive or discriminatory language
- Profanities
- Trade secrets/confidential information

Thank you for doing your part to keep Glassdoor the most trusted place to find a job and company you love. See the Community Guidelines for more details.

Notes: Figure is a screenshot of the survey users are asked to fill out when submitting an employer review to Glassdoor. An asterisk indicates that the field is required to submit the survey. Overall rating is restricted to between one and five stars, with only integer ratings possible. Once text is added to the "Review Headline," "Pros," or "Cons" sections, users are asked to provide additional information (not shown), which includes the location of employment with the firm.

[10] To create some accountability on volunteers, Glassdoor requires them to have a verified active email address or a valid social network account, assesses the content of each submitted review, and suppresses those outside their guidelines. Assessment guidelines are here: https://help.glassdoor.com/article/Community-Guidelines/en_US.

Each review contains many kinds of information. Figure 1 displays a blank review form for the University of Pennsylvania. Volunteers are asked to provide feedback on their employer, through both supplying a star rating for the firm and submitting free response text for the "Pros" and the "Cons" of working for that firm. In order to complete the review, volunteers are required to supply the following information: an overall rating, employment status, review headline title, "Pros" text, and "Cons" text. They also have the option to voluntarily supply information about their job title, tenure at the firm, and location of employment. The volunteer's employment status, job title, and geographic location (if available) are displayed with their review to site visitors.

Each volunteer assigns the firm an overall rating on a one-to-five-star Likert scale. Let R_{fr} denote the volunteer's overall rating R in review-r. R is a vertically differentiated, scalar summary of the volunteer's overall evaluation of the employer and is meant to inform other workers about the quality of employment provided by firm-f. We sometimes refer to reviews that include one- or two-star ratings as negative reviews and those including four- or five-star ratings as positive reviews, or sometimes just use "more negative reviews" to refer to reviews that include a lower rating. Broadly, the site contains about twice as many positive reviews as negative reviews.

The volunteer can also write text freely in "Pros" and "Cons" text fields. We will leverage this uniquely rich source of information about what workers think of their jobs through this pro-con field structure. To get multidimensional ratings, we leverage volunteers' free-text responses in the "Pros" and "Cons" fields of each review. This allows us to model difficult-to-observe aspects of jobs.[11] Using an unsupervised topic-modeling approach on 6.81 million reviews reported through August 2019, we extract the most important latent attributes for explaining the review text.[12] The researcher has freedom over the number of latent attributes to extract, and we choose 20. Any text document can then be scored according to these attributes. The scores are probability weights and so measure the extent to which the document (review-side, or review x pros/cons) discusses each attribute. We score each review's "Pros" and "Cons" sections with each of the 20 latent attributes. For each review-side rs, this yields estimates of what share of the text's attention pertains to each attribute, $p_{rsa} \in [0, 1]$ for $\forall a = \{1, 2, ..., 20\}$, and these shares approximately sum to one within each review-side, $\sum_a p_{rsa} \approx 1$.[13] We measure the extent to which any review-r is

[11] Figure A-1 provides a few examples from employees of the University of Minnesota.

[12] We use latent Dirichlet allocation (LDA) modeling with the review text, merging the "Pros" and "Cons" sections into a single document and analyzing the highest-incidence one-, two-, and three-word phrases. This dimension-reduction technique yields estimates that enable scoring the original, high-dimensional text space into a lower-dimensional attribute space (Blei et al., 2003; Lopez-Lira, 2019; Marchetti, 2019). For a detailed description of how we implement the LDA algorithm, see Appendix Section A.2.

[13] Restricting our LDA algorithm to the highest-incidence one-, two-, and three-word phrases means that

about attribute-a as the average probability mass across its two sides, weighted by the share of review-r characters devoted to that side ($c_{r,pro}$): $p_{ra} \equiv p_{r,pro,a}c_{r,pro} + p_{r,con,a}(1 - c_{r,pro})$. Figure A-3 displays how much the average review discusses each attribute and how much weight is attributable to the pro and con sides of each attribute.[14]

In any review, we measure the volunteer's evaluation of each attribute on the job as the difference between the pros and cons sections' scores on that attribute. The review-r evaluation of attribute-k is denoted by $A_{rk} = p_{r,pro,k} - p_{r,con,k}$, giving equal weight to each side. This normalizes within side to control for the volunteer's overall propensity toward positivity or negativity. For instance, in a review, attribute 3 may dominate the "Pros" section and barely show up in the "Cons" section ($p_{r,pro,3} > p_{r,con,3}$). Then, A_{r3} is highly positive and review-r of this job is measured as highly favorable from the volunteer's perspective with respect to attribute 3. In contrast, attribute 5 may be discussed with equal weight in the pro and con sections ($p_{r,pro,5} = p_{r,con,5}$). In that case, we measure the job as being neutral with respect to attribute 5. Appendix Table A-1 provides summary statistics for the evaluation of each attribute across reviews. These evaluation variables necessarily range from -1 to 1, tend to average near 0 across reviews, and have standard deviations ranging from 0.053 to 0.122 across the 20 attributes.

Once a volunteer's review is added to the website, people who go to the website to learn about that employer can see the review and vote it as helpful. Until a few years ago, users could also vote it as unhelpful (Appendix A.1 discusses in greater detail). Via these (un)helpful votes, jobseekers express what kinds of private information about employers they value, so we use them as a measure of jobseeker demand for different kinds of information about jobs with employers.

Worried volunteers may thus try to conceal aspects of their identity when leaving a review. We measure whether a volunteer attempts to conceal their identity by how they fill in the job title and location fields of the review. Forty-four percent of reviews' volunteers choose to leave the job title field blank, a decision we interpret as potentially motivated by a desire for anonymity. We separately measure reviews that actively, rather than passively, express a desire for job-title anonymity. These reviews are ones where the word "anonymous" is included in the job title, or where the job title ends in the phrase "employee." Just over half

not every phrase with a review-side is necessarily mapped to any attribute. On average, 95.2% of a review's text is apportioned to attributes, with the left-out share roughly even across the pro and con sections.

[14] The Appendix contains further description of these variables. Figure A-4 displays the words and phrases most strongly associated with each of the attributes. This study does not focus on assigning distinct meaning to each latent attribute but simply takes them as representing variations in attributes workers describe to one another about their jobs. Future work could explore developing richer interpretation of latent attributes. Marchetti (2019) finds that matching between firms' work cultures, as measured by the match of latent factors, predicts merger success.

a percent of volunteers actively report anonymously. Volunteers can also hide identifying information by leaving the geographic location field blank. About 42 percent of volunteers do not report a location. These pieces of volunteer identity information, or lack thereof, pipe into the display of the review that jobseekers see.

To understand jobseeker demand for information in a review, we measure the helpfulness of any review with the share of jobseeker votes that are helpful rather than unhelpful. When a jobseeker reads a review, she could click to classify the review as either helpful or unhelpful or could refrain from expressing an opinion. We measure any review-r's helpfulness to jobseekers as the share of total jobseeker classifications that are helpful ($H_r = \#\text{helpful}_r/[\#\text{helpful}_r + \#\text{unhelpful}_r]$), which we call helpful share.[15] When analyzing helpful share, we analyze the 685,505 reviews posted before 2015 that got at least one of the 1.90 million helpful or 0.23 million unhelpful votes. These reviews averaged 2.76 helpful votes, 0.34 unhelpful votes, and a helpful share of 0.86 with a standard deviation of 0.31. Using the helpful share, rather than the helpful count, allows us to control for the number of jobseekers who saw a review but results are similar in any case. Other reviews posted during this time got no votes at all (731,989) and are excluded because we lack a way to control for how (in)frequently jobseekers viewed them or a measure of how meaningful the review was to jobseekers who did in fact view them.

Glassdoor has additional information about reviewed firms through an employer lookup table. For each firm, there is a single entry that contains the following information (when available): the industry of the firm, the most-recent employment total for the firm (meaning that our measure of firm employment has no time variation), and the year in which the firm was founded.[16] Firm age is then calculated as the difference between the year in which the review is submitted and the year the firm was founded. We incorporate each of these three firm characteristics in order to study differences in identity concealment, review helpfulness and job satisfaction across firms.

[15] Analyses with this outcome will always control for the year-month the review was posted. In 2015, Glassdoor phased out the option for jobseekers on its website to classify reviews as unhelpful, leaving the helpful and refrain options. For an in-depth description of the timeline, see Appendix A.1. For the share of unhelpful classifications submitted for employer reviews over time, see Appendix Figure A-2.

[16] There are 25 broad industries in the data: Accounting & Legal, Aerospace & Defense, Agriculture, Arts & Recreation, Biotech & Pharmaceuticals, Business Services, Construction, Consumer Services, Education, Energy, Finance, Government, Health Care, Hospitality, Information Technology, Insurance, Manufacturing, Media, Mining, Non-Profit, Real Estate, Restaurants, Retail, Telecommunications, and Transportation.

4 Results

4.0.1 Identity concealment and information degradation

If workers worry about potential retaliation from leaving negative reviews, then they might be more likely to leave negative reviews when it is harder for employers to infer who wrote them. In fact, volunteers leaving more negative reviews are more likely than other volunteers to conceal aspects of their identity (Table 1: Columns 1–2). This holds in the 6.6 million reviews across volunteers (Column 1). Reviews by volunteers who passively conceal their job title by leaving the field blank tend to rate the employer 0.04 star lower overall on average controlling for the year-month of the review and the employing firm. Reviews by volunteers who actively anonymize their job title tend to rate their job –0.38 stars lower. Leaving location blank is not significantly associated with rating in this specification.

These relationships are not driven by individual differences across volunteers in their tendency to review negatively and to withhold information about identity. The relationships hold across reviews when including volunteer fixed effects among those who volunteer multiple reviews (Column 2). The same people are more likely to conceal each aspect of their identity when leaving more negative reviews.

If fear of retaliation contributes to concealment behavior, it would be less common in contexts where this risk is less of a concern. What kinds of workers should worry less? First, former employees should worry less than current employees. Retaliation threats that apply to former employees (e.g., the prospect of a bad reference in the future) also apply to current ones but current employees face additional risks (e.g., undesirable schedule or task assignments, demotion, personal harassment, or firing). Second, workers in firms with more employees might worry less about the boss inferring their identity, because they can blend into a larger crowd.

Consistent with these predictions, volunteers leaving negative reviews (overall rating of one or two stars) are less likely to conceal each of these identity aspects if they are not a current employee and if their firm is larger (Columns 3–5).[17] For these specifications, we incorporate volunteer fixed effects, meaning that our estimates are identified off of the same worker leaving multiple reviews of 1 or 2 stars on the website for firms of different sizes and for whom reviews toggled between current and former employee status.[18] Current employees are 3.2 and 4.1 percentage points more likely, and respondents working at twice-as-large firms are 2.1 and 1.4 percentage points less likely, to leave the job title or location blank,

[17] We have only a cross-sectional measure of firm size from 2019, used in log form.

[18] Semi-parametric specifications for Columns 3–5 are presented in Figure A-5, where concealing job title comprises leaving the job title blank or anonymizing the job title.

Table 1: Overall Evaluation and Review Helpfulness by Degree of Identity Concealment

	Star rating		𝟙{Conceal info | 1 or 2 stars} job title blank	job title anonymized	location blank	Share of review votes helpful
Job title blank	-0.039***	-0.070***				-0.019***
	(0.003)	(0.005)				(0.001)
Job title anonymized	-0.384***	-0.238***				-0.026***
	(0.009)	(0.021)				(0.003)
Location blank	-0.003	-0.085***				-0.014***
	(0.002)	(0.005)				(0.001)
Current employee			0.032***	0.004***	0.041***	
			(0.003)	(0.001)	(0.003)	
Log(employment)			-0.021***	-0.001***	-0.014***	
			(0.001)	(0.000)	(0.001)	
Star rating						-0.088***
						(0.002)
Sample mean	3.34	2.99	0.399	0.012	0.347	0.86
Year-month FE	✓	✓	✓	✓	✓	✓
Employer FE	✓	✓				✓
Volunteer FE		✓	✓	✓	✓	
Industry FE			✓	✓	✓	
N	6606743	945845	262072	262072	262072	657691
Adjusted R^2	0.14	0.42	0.29	0.08	0.27	0.24

Notes: When submitting a review, users are asked to provide their job title and location, both of which are optional. If these fields are not completed but the review is still submitted, they are labeled as "blank." Additionally, respondents can choose to leave a job title but anonymize their position. A job title is classified as anonymized if either the word "anonymous" is included in the job title or the job title ends with "employee." Firm employment is based on a fixed employer lookup table from 2019 and so does not vary over the sample period. Standard errors are clustered at the employer level. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. For a summary of the dataset, see Table A-1.

respectively. Although actively anonymized job titles are rare, they have 0.4 percentage point higher incidence among current employees, or one-third the regression sample average.

Finally, when volunteers conceal aspects of their identity, it may degrade the value of the information they supply, since doing so limit the ability of jobseekers to judge its relevance to their own decisions. For instance, a jobseeker deciding whether to apply for a particular job posting may derive the greatest value from reviews by volunteers with the same job title and from the same location as the posting. To test this, we predict the review's helpful share by whether the volunteer concealed aspects of their identity. Indeed, jobseekers tend to classify reviews from volunteers who concealed their identity as less helpful conditional on the reviewed firm, the review's year-month, and the volunteer's overall rating of the firm (Column 6). Reviews for which the job title is not identifiable are on average 1.9-2.6 percentage points less helpful and those for which the location is unavailable are about 1.4 percentage points less helpful.

Here, we also see the first evidence that jobseekers consider more negative information supplied by volunteers as more helpful. Jobseekers tend to give reviews with a one-star higher overall rating an 8.8 percentage point lower helpful share. The next section will develop this point more fully.

The result that volunteers in situations with higher retaliation risk (current employee or in a smaller firm) are more likely to conceal aspects of their identity conditional on leaving a negative review (Table 1: Columns 3–5) generalizes across the full sample and is weaker among those facing lower retaliation risk, consistent with the mechanism posited here. If the fear of retaliation suppresses the supply of negative information, workers facing higher retaliation risk should seek anonymity when relaying negative information more so than when relaying positive information. We test this with interactions between volunteer's rating of the firm and measures of retaliation risk in predicting volunteer identity concealment. Figures 2 and 3 present flexible, descriptive evidence and Table A-4 reports analogous inference.

Figure 2: Share of Volunteers Concealing Aspects of Their Identity by Their Rating of Firm, Employment Status, and Firm Size Decile

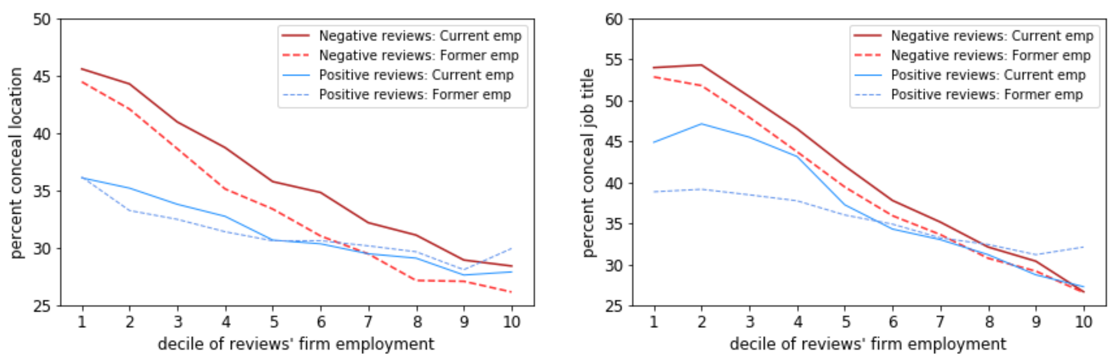

(a) Conceal location (b) Conceal job title

Notes: The figures above detail the rate at which volunteers conceal potentially identifying information depending upon the size of the employer and whether the volunteer is still currently employed at the firm. The sample of positive reviews reflects four- and five-star reviews, while the sample of negative reviews reflects one- and two-star reviews. Sample of volunteers is restricted to those who leave multiple reviews on the website. Firm size deciles are defined across reviews.

To assess the robustness and generality of this result, we first present simple summary evidence from an expanded sample that includes not only negative reviews as in Columns 3–5 of Table 1 but also includes positive reviews (four- or five-star). We stratify the sample by measures of retaliation risk—whether the volunteer was a current or former employee and the size decile of the firm reviewed. In Figure 2, reviews are grouped into decile based on firm size with reviews of larger firms further right on the horizontal axis. Reviews by current (former) employees are represented by solid (dashed) lines. Those with negative (positive) ratings are represented by thicker red (thinner blue) lines. First, those leaving negative reviews are more likely to conceal their location and job title than those leaving positive reviews, particularly in smaller firms. The difference in concealment probabilities between negative rating and positive rating reviews is larger in smaller firms than in larger firms.[19] Third, volunteers leaving negative reviews conceal more when they are current employees than former employees. For negative reviews, the additional wedge in concealment between current and former employees varies in magnitude but persists across the spectrum of firm size.

If fear of retaliation drives identity concealment—and thus likely works to suppress negative feedback more broadly—we should observe that reviews offering negative ratings are more likely to have identifying aspects concealed than those relaying positive ratings, especially among volunteers facing higher retaliation risk. To illustrate this, we consider a model that relates identity concealment to the interaction between leaving a negative rating and our two measures of retaliation risk. Let $1(Conceal)_{ir}^{a}$ indicate if volunteer-i concealed aspect-a of their identity in review-r. Consider linear probability models for $a \in \{location, jobtitle\}$,

$$\begin{aligned} 1(Conceal)_{ir}^{a} =& \beta_1^a Rating_{ir} + \beta_2^a log(FirmSize)_{ir} + \beta_3^a 1(CurrentEmployee)_{ir} \\ &+ 1(Rating \leq 2)_{ir}[\beta_4^a log(FirmSize)_{ir} + \beta_5^a 1(CurrentEmployee)_{ir}] \\ &+ \delta_{ir}^{a,Industry} + \delta_{ir}^{a,Year-Month} + \gamma_i^a + \epsilon_{ir}^a \end{aligned} \tag{1}$$

With reviews left by volunteers who left multiple reviews, we investigate whether a worker's concealment probability is higher when leaving a negative review and being either at a smaller firm ($\beta_4^a < 0$) or a current employee ($\beta_5^a > 0$). Estimates of this full model are displayed in Table A-4 and are consistent with the predicted signs.

Figure 3 presents more flexible evidence that allows for but does not find possible non-linearities in the key relationships. The top-left (top-right) panel considers how the difference in the probability of concealing location when leaving a negative rather than a positive review relates to firm size (current employee status). The former (latter) gives a visual representation of the semi-parametric relationship expressed in β_4 (β_5) for $a = location$. We

[19] Another way to say this is that the decline in concealment probability with firm size appears steeper for negative reviews than positive reviews. Compared with workers leaving negative sentiment for the largest firms, employees supplying negative information for the smallest firms are on average about 2x more likely to conceal their job title and 1.6x more likely to conceal their location. For positive reviews, the multipliers are still positive but smaller at 1.5x and 1.3x for concealing the job title and location, respectively.

Figure 3: Relationships Between Probability of Identity-Aspect Concealment and Leaving a More-Negative Review by Measures of Retaliation Risk

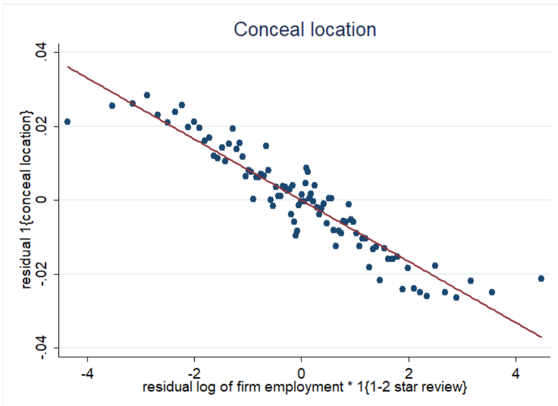

(a) Location by firm employment

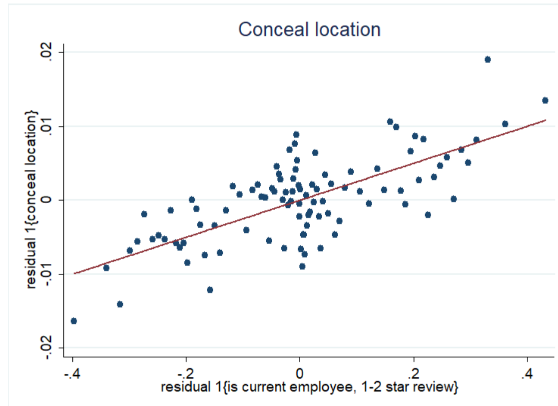

(b) Location by current employee

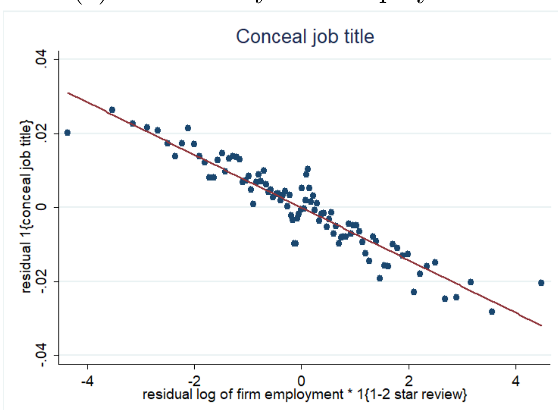

(c) Job title by firm employment

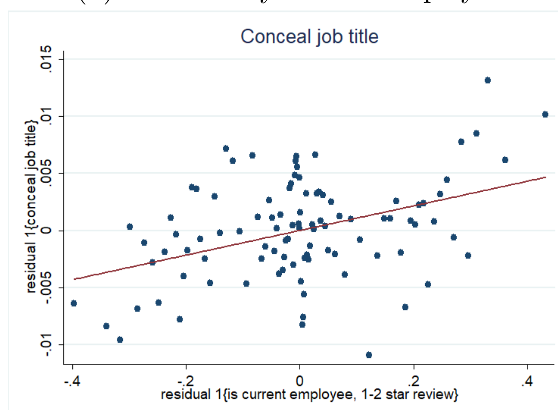

(d) Job title by current employee

Notes: The figures above describe the average residual rate at which volunteers conceal aspects of their identity (location in top row and job title in bottom) depending on measures of retaliation risk (number of employees in firm in left column and whether volunteer is a current versus former employee in right) interacted with an indicator of leaving a negative (one- or two-star) rating instead of a higher rating. Concealing job title reflects reviews that leave the job title blank or anonymize the job title. Concealing location reflects reviews that leave the location blank. The sample is restricted to volunteers who leave multiple reviews. In each panel, the residuals come from a regression like equation 1 but excludes the interaction of the horizontal retaliation-risk variable with an indicator of whether the rating is negative. Solid red line reflects a linear line of best fit. Residuals clustered into bins using the *binscatter* function in Stata.

residualize an indicator for concealing location on the main effects of review rating, log firm size, a current employee indicator, and fixed effects for the volunteer, the year–month, and the firm's industry. We also include the interaction of the other retaliation risk measure, current employee status, with an indicator of negative review. Only the interaction term between an indicator of a negative review and the focal measure of retaliation risk, log firm size, is excluded from the model. Observations are binned by the residualized interaction between log firm size and the negative review indicator and scattered in the top-left panel.

Evidently, the relation appears linear and is strongly and significantly negative, pointing to systematically higher rates of concealment when supplying negative information at smaller firms than at larger firms. The top-right panel repeats the exercise but swapping the roles of the retaliation risk measures, looking instead at current employee status. Here again, concealment is higher in the context of negative reviews and higher retaliation risk. The lower panels of Figure Figure 3 repeat this analysis for concealing job title rather than location. The results are similar.

Retaliation risk may raise the implicit cost of supplying negative information above the cost tied to positive information. Given evidence that negative information may be undersupplied, existing evidence that workers are less likely to volunteer negative information than positive information (Marinescu et al., 2018) and that many firms create legal obstacles to prevent workers from disclosing negative information (Starr et al., 2019; Lobel, 2019) and the new evidence from this section that workers are more likely to conceal aspects of their identity when volunteering negative information in ways consistent with fear of retaliation risk, we next turn to understanding workers' demand for information. Table 1 showed initial evidence that, conditional on the degree of volunteer identity concealment, jobseekers found reviews attached to more negative overall ratings more helpful.

4.1 Jobseekers demand negative information

The section presents evidence that stronger jobseeker demand for negative information is very robust, a new empirical result. Scarcity of negative information from volunteers should contribute to raising jobseeker relative demand for negative information. We also discuss non-mutually exclusive channels that also likely contribute, such as worker risk aversion and employers supplying positive information. A failure to find this relationship would have cast strong doubt on the idea that missing or degraded negative information matters to jobseekers. Because of that, finding it provides some corroboration.

Our analysis rests on the idea that, when jobseekers view a volunteer's job review, they are more likely to classify the review as helpful if it contains the kinds of information they want to learn and are more likely to classify it as unhelpful if contains kinds of information they do not want. First, we focus on whether the positive or negative overall information content relates to jobseekers' evaluation of a review's helpfulness. Second, drawing on novel measures of typically-hard-to-observe job attributes described in each review, we allow for the possibility that jobseekers' preferences for information differ depending on the attribute of the job on which it focuses. Third, we see how jobseeker evaluation of information helpfulness varies depending upon individual preferences for attributes.

Jobseekers find reviews containing negative information about employers unequivocally more helpful than those containing positive information. Straightforward evidence comes from comparing the distributions of helpful and unhelpful votes depending on the overall rating value attached to each review (Figure 4). Not only are helpful votes predominantly concentrated among the lowest two rating options (one or two stars), but unhelpful votes—a clear indication of relative informational value—are predominantly concentrated among the highest two rating options (four or five stars). Reviews where the volunteer gave their firm a one-star rating account for 38 percent of helpful votes but only 13 percent of unhelpful votes. For the other extreme (five-star ratings), the vote shares flip. And although reviews with less extreme values (two, three, or four stars) may be less biased, they are each deemed less helpful than the most negative (one star) reviews.

Figure 4: Marginal Distributions of Jobseeker (Un)Helpful Votes by Review Rating

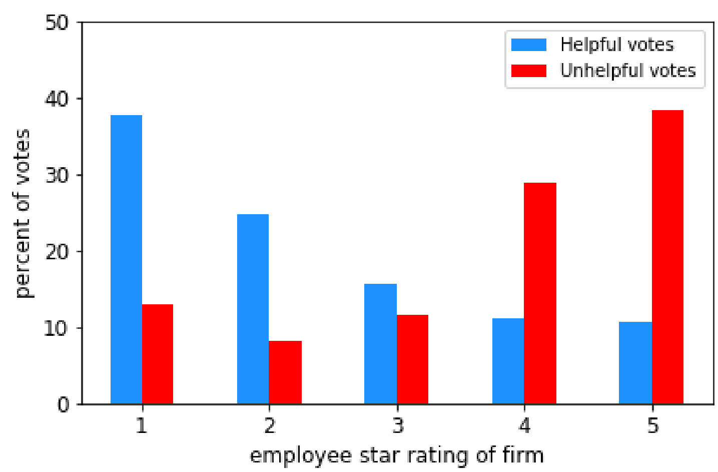

Notes: The figure above shows the distribution of all helpful votes and unhelpful votes submitted by users of the website for employer reviews submitted before 2015 (see Appendix A.1), where reviews are partitioned according to the employer rating.

The value of information, in particular negative information, might reasonably depend upon the firm in question. For example, information may be less in demand for more-established firms with well-understood off-line reputations or for firms that already have more established on-line reputations, with many reviews on Glassdoor. With this in mind, we examine whether the value of positive and negative information to jobseekers changes across firm characteristics.

Although the magnitudes vary, in every type of firm observed, jobseekers rate more positive reviews as less helpful than other reviews. Figure 5 shows this is true across firm age, overall rating for the firm at the time the review is submitted, the number of prior reviews on the website, and firm-specific wage premia as proxied by a fixed effect in base

Figure 5: Share Helpful Votes by Review's Rating and Firm Demographics

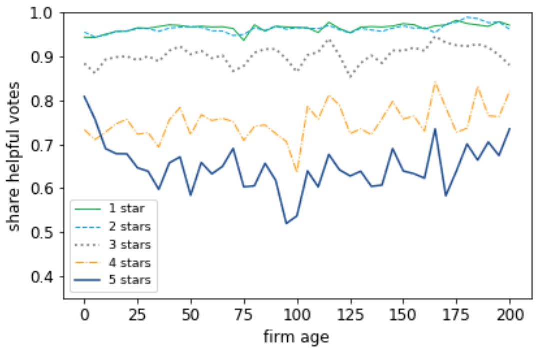

(a) Firm age

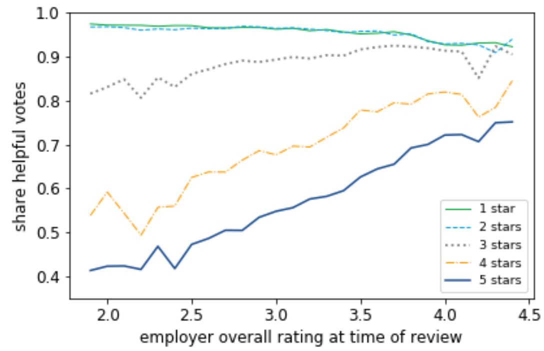

(b) Firm overall rating

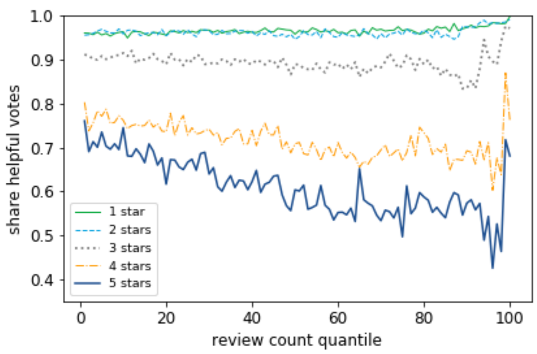

(c) Firm review count

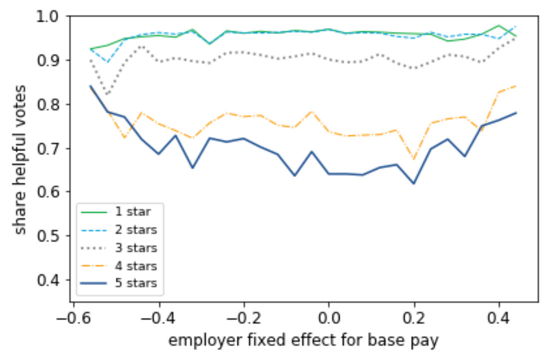

(d) Firm FE for base pay

Notes: Figures display the average share helpful votes across firms partitioned by the star rating of the review. Firm age is calculated as the difference between the year in which the review is submitted and the year in which the firm was founded (which is recorded in a fixed employer lookup table from 2019). Sample is restricted to reviews submitted before 2015 (see Appendix A.1). Sample for panels (b) and (c) restricted to reviews submitted for firms that have had at least 10 reviews submitted prior. Firm FE for base pay, derived using Glassdoor pay data, are from a regression of annualized log base pay on a quadratic in years of specific experience and fixed effects for year, state, gender, educational attainment, pay frequency, industry-job title pairing, and employer (see section 4.3.1). Firm age grouped into five-year bins. Firm overall rating grouped into bins rounded to the nearest tenth. Firm FE for base pay grouped into 0.04 bins. For each panel, bins in which one of the five ratings represents fewer than 150 unique reviews are excluded.

pay.[20] Consider sub-figure (a). Firm years of age at the time of the review in five-year bins are spread across the horizontal axis. Conditional on that, the reviews are partitioned into 5 groups corresponding to the star rating attached to the review, capturing the extent of negative or positive information in the review. Averaged across reviews within each age-stars group, the share of jobseekers' votes that corresponded to helpful is graphed. For every firm age, one star reviews are rated helpful by almost all jobseekers who vote whereas

[20] This is described in greater detail in section 4.3.1. For detailed analysis of Glassdoor pay data, see Sockin and Sockin (2019a,b). For a comparison of Glassdoor pay data to surveys of income commonly used in the literature, see Liu et al. (2017).

five-star reviews are rated helpful by only about 60 percent of jobseekers who vote. Across all four degrees of variation we explore across firms, four- and five-star reviews have the lowest average helpful vote shares. Negative information about the firm is most sought after by jobseekers regardless of how long the firm has been in operation, how highly rated it has been, how many prior reviews it has, or how much it pays its workers relative to other firms. The ranking of helpfulness across review ratings is qualitatively highly stable. The average helpful vote share for one- and two-star reviews consistently hovers around 95 percent, whereas the average helpful vote share for four- and five-star reviews consistently falls below eighty percent.

Additional evidence comes from a regression of a review's helpful share on the review's star rating of the firm, controlling for other factors that might influence helpfulness. In particular, we control for firm fixed effects; year-month fixed effects; and measures of the nature of the review text, including the length of the review's text, the sophistication of its language, and the subjectivity of its tone.[21] Reviews communicating a one-star-higher rating were classified as helpful by a smaller share of voting jobseekers, resuling in an 8.5 percentage points lower helpful share (Table 2: Column 1). This result is robust to predicting a review's helpful share with alternative measures for the volunteer's evaluation of her job in lieu of the employer rating. One such alternative is whether the volunteer would recommend the employer to a friend (59.6 percent would per Appendix Table A-1). On average, the helpful share of would-recommend reviews is 22.8 percentage points lower that that of would-not-recommend reviews (Column 2). Reviews where the volunteer approves of the CEO witness a 17.3-percentage-points-lower helpful share (Column 3), while those that report a positive business outlook for the firm see a 17.7-percentage-points-lower helpful share (Column 4). An additional alternative—which is also a more continuous measure of a review's positivity—is the share of text characters spent discussing pros of the job, rather than cons. The pro share of review text (which averages 46.8 percent) gauges the volunteer's overall evaluation of her job, avoiding the coarseness of the discrete five-star rating system. Reviews with a 10-percentage-points-higher pro share of review text result on average in a 3.8-percentage-points-lower helpful share (Column 5). Finally, we apply sentiment analysis to the review text and measure the positive versus negative emotional polarity of the text and similarly find that more negative sentiment improves helpful share (Column 6). When we include all of these measures together, they all negatively predict helpfulness conditional on one another:

[21] Although not the main takeaway from Table 2, the helpful share regressions indicate that longer, better-written, and more objective reviews are viewed as more helpful by jobseekers. While the quality and subjectivity of the text have relatively muted effects, doubling the length of a review is associated with a helpful share that is about 4 percentage points higher, conditional on the review's attitude to the firm.

Evidently, this finding is extremely robust.[22]

Table 2: Predicting Helpfulness of Volunteer's Review

	Share of review votes helpful						
Star rating	-0.085*** (0.001)						-0.044*** (0.001)
Would recommend employer to a friend		-0.228*** (0.004)					-0.075*** (0.003)
Approves of the CEO			-0.173*** (0.003)				-0.028*** (0.002)
Positive business outlook for the firm				-0.177*** (0.003)			-0.041*** (0.002)
Pro share of review text					-0.383*** (0.006)		-0.039*** (0.004)
Polarity of text						-0.165*** (0.003)	-0.017*** (0.003)
Log length of review (total characters)	0.044*** (0.001)	0.043*** (0.001)	0.057*** (0.001)	0.049*** (0.001)	0.040*** (0.001)	0.062*** (0.001)	0.034*** (0.001)
Flesch kincaid reading grade	0.002*** (0.000)	0.001*** (0.000)	0.002*** (0.000)	0.001*** (0.000)	0.002*** (0.000)	0.001*** (0.000)	0.002*** (0.000)
Subjectivity of text	-0.015*** (0.003)	-0.020*** (0.003)	-0.029*** (0.003)	-0.020*** (0.004)	-0.021*** (0.003)	0.007** (0.003)	-0.008** (0.004)
Employer, year-month FE	✓	✓	✓	✓	✓	✓	✓
N	530502	458916	407704	355688	530502	530502	291785
Adjusted R^2	0.26	0.26	0.20	0.21	0.19	0.16	0.28

Notes: The dependent variable, share helpful votes, is defined as the ratio of helpful votes to the sum of helpful and unhelpful votes. Sample is restricted to reviews submitted before 2015 (see Appendix A.1) as well as reviews for which the Flesch Kincaid reading grade is non-negative and no greater than 20. Polarity and subjectivity of each review are measured through natural langugage processing using the *TextBlob* library in Python. Standard errors clustered at the employer level. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. For a summary of the dataset, see Table A-1.

Next, we consider whether there is heterogeneity in demand for negative information depending on the voting jobseeker's satisfaction with their own job. We split the sample between jobseekers who expressed dissatisfaction with their own employer (one or two stars), moderation (three stars), or satisfaction (four or five stars). This requires us to restrict the analysis to a subsample of (un)helpful votes where the jobseeker reviewed her own employer. To try to clarify the influence of potential insincere voting by "jobseekers," we further split each subsample between jobseekers voting on the helpfulness of a review who (a) do not work at the reviewed firm and are, therefore, presumably more likely to be engaged in genuine job search versus (b) do work at the reviewed firm and, therefore, may be more likely to be engaged in insincere sock puppetry. This partitions observations into six (un)helpful voter

[22] An alternative specification—where the count of helpful votes is predicted rather than the helpful share—reaffirms this robust finding that negative information is most helpful. Results from these count-helpful regressions are displayed in Table A-2. This sample of reviews is not restricted to the pre-2015 period.

subsamples, by the voter's own expressed satisfaction level at their own employer crossed by whether or not their employer is the firm in the review on which they are voting.

In each subsample, we predict whether a jobseeker votes the review helpful rather than unhelpful as a function of the review's overall rating, to capture if jobseekers show a stronger relative demand for negative versus positive information (Table 3). For jobseekers voting on reviews of firms where they do not work, demand for negative information is very similar regardless of their own job satisfaction. Among those who rated their own firm as one- or two-star, reviews with one-star higher overall rating by the volunteer are 7.9 percent less likely to be voted helpful by the jobseeker. Among those who rated their own firm as four- or five-star, the analogous estimate is 7.8 percent less likely.

Table 3: Conditional Probability of Helpful Vote by Voter Rating and Firm Co-incidence

	Voter's rating of employer from own review					
	Review firm ≠ Voter's firm			Review firm = Voter's firm		
	1–2 stars	3–4 stars	5 stars	1–2 stars	3–4 stars	5 stars
Review's rating of employer	-0.079***	-0.063***	-0.078***	-0.181***	-0.057***	0.123***
	(0.005)	(0.008)	(0.011)	(0.005)	(0.009)	(0.038)
N	77990	24537	9509	51888	8155	5515
Adjusted R^2	0.42	0.43	0.37	0.51	0.28	0.42

Notes: Sample consists of a panel of (un)helpful votes for different employer reviews. The dependent variable is a dummy variable that the user up-voted the review helpful. Because the dataset consists only of helpful and unhelpful votes—meaning it excludes decisions where no vote was given—this dummy is conditional on submitting a vote. Sample is restricted to voting users who submitted at least one of their own employer reviews on the website prior to submitting the (un)helpful vote. Each regression includes fixed effects for the reviewed employer. Standard errors are clustered by voter. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. For a summary of the dataset, see Table A-1.

In contrast, among jobseekers who work at the reviewed firm, there's a sharp difference in voting behavior depending on own job satisfaction. Those who express dissatisfaction with their own job are more likely to vote negative reviews as helpful. However, those who express satisfaction with their own firm are more likely to vote positive reviews as helpful. This is the only exception to the pattern of more negative information being more demanded. This could be due to sock puppetry, a boss rewarding employees for supplying positive reviews of the firm and upvoting other positive reviews of the firm. It could also be sincere. I like my firm and I think reviews by others who like the firm will be more helpful to jobseekers.

Jobseeker preference for negative information extends from unidimensional, overall evaluations of the employer to multidimensional evaluations based on multiple attributes detailed in the review's text. For every single dimension conditional on the others, jobseekers rate reviews with more negative volunteer evaluations as more helpful. We predict each review's helpful share with the review's vector of 20 attribute evaluations. Every coefficient is negative and significant looking across firms (Table 4: Column 1) and within firms (Column 2).[23] In the cross-section, certain volunteers are simply just more helpful than others. If this is correlated with their evaluations of attributes, it could create bias. To minimize this threat, we look at consecutive pairs of reviews by the same volunteer, although this restricts our sample to volunteers who left multiple reviews. For each of the 31,984 same-volunteer review pairs before 2015, we study how first-differences in attribute evaluations predict differences in helpful share. Within a given volunteer, we regress differences in helpful share on differences in the vector of attribute evaluations and find the same basic result. On every attribute, the coefficients are smaller in magnitude, but reviews with more negative information are still voted more helpful by jobseekers (Column 3).

Theory suggests that a jobseeker will find a review more helpful if it contains information that shifts her beliefs on a particularly steep part of the expected utility function. If a jobseeker doesn't really care about a specific job attribute but reads a review that focuses heavily on it, then the jobseeker should be unlikely to classify that review as helpful. We test this theory by allowing the review's helpfulness for a jobseeker to differ depending on the jobseeker's own preferences. To test this, we need to see the identities of jobseekers who classified a review as (un)helpful and to measure the strength of each one's preferences over each attribute. Although identities of jobseekers are not available for every vote on every review, for a subsample of jobseekers who were logged into the website when classifying reviews, anonymous identifiers are available. To measure heterogeneity in the strength of their preferences across attributes, we consider any reviews that a voting jobseeker volunteered on their own employers in the past. We assume that if a jobseeker devoted a greater share of her total review text to an attribute, then she cares more about that attribute.

Instead of predicting each review's helpful share based on a share among all voting jobseekers, as in Tables 2 and 4, for this subsample, we predict whether an individual jobseeker voted a review as helpful versus unhelpful, as in Table 3. In this review-voter pair sample where we have a measure of voter preferences, we reproduce the basic finding from Table 2, that jobseekers find reviews with more negative overall ratings significantly more helpful (Table 5: Column 1). Then, we look for evidence on whether jobseekers find a review increasingly helpful when it provides more information about the attributes they value more highly. To do so, we calculate the distribution of mass across the attributes ($\{p_{ra}\}_{a=1}^{20}$) for

[23]Without loss of generality, we ordered the attributes by declining magnitude of their coefficients in this regression (Column 2). Attribute 1 is the attribute where a unit variation in A_{rk} share has the largest magnitude of conditional association with helpful share.

Table 4: Predicting Helpfulness of Volunteer's Review by Its Evaluation of Job Attributes

	Share helpful		First difference share helpful
Pro minus con: attribute 1	-0.316***	-0.361***	-0.221***
	(0.007)	(0.007)	(0.018)
Pro minus con: attribute 2	-0.277***	-0.322***	-0.249***
	(0.007)	(0.007)	(0.015)
Pro minus con: attribute 3	-0.321***	-0.320***	-0.203***
	(0.007)	(0.007)	(0.025)
Pro minus con: attribute 4	-0.296***	-0.306***	-0.173***
	(0.006)	(0.005)	(0.014)
Pro minus con: attribute 5	-0.249***	-0.286***	-0.140***
	(0.008)	(0.008)	(0.029)
Pro minus con: attribute 6	-0.247***	-0.275***	-0.156***
	(0.005)	(0.005)	(0.014)
Pro minus con: attribute 7	-0.267***	-0.274***	-0.171***
	(0.008)	(0.008)	(0.027)
Pro minus con: attribute 8	-0.262***	-0.263***	-0.201***
	(0.010)	(0.008)	(0.028)
Pro minus con: attribute 9	-0.214***	-0.218***	-0.163***
	(0.009)	(0.008)	(0.033)
Pro minus con: attribute 10	-0.192***	-0.217***	-0.138***
	(0.005)	(0.007)	(0.019)
Pro minus con: attribute 11	-0.180***	-0.215***	-0.135***
	(0.004)	(0.005)	(0.012)
Pro minus con: attribute 12	-0.187***	-0.214***	-0.112***
	(0.005)	(0.004)	(0.010)
Pro minus con: attribute 13	-0.194***	-0.208***	-0.119***
	(0.009)	(0.009)	(0.031)
Pro minus con: attribute 14	-0.162***	-0.171***	-0.095**
	(0.012)	(0.012)	(0.047)
Pro minus con: attribute 15	-0.133***	-0.164***	-0.106***
	(0.007)	(0.007)	(0.022)
Pro minus con: attribute 16	-0.132***	-0.141***	-0.082***
	(0.006)	(0.005)	(0.019)
Pro minus con: attribute 17	-0.129***	-0.140***	-0.125***
	(0.006)	(0.005)	(0.020)
Pro minus con: attribute 18	-0.061***	-0.077***	-0.071***
	(0.006)	(0.005)	(0.017)
Pro minus con: attribute 19	-0.062***	-0.058***	-0.001
	(0.008)	(0.008)	(0.035)
Pro minus con: attribute 20	-0.037***	-0.056***	-0.005
	(0.008)	(0.009)	(0.028)
Year-month FE	✓	✓	
Employer FE		✓	
Former x latter industry, year-month FE			✓
N	685500	657688	31984
Adjusted R^2	0.11	0.19	0.08

Notes: The (first difference in) share helpful is (the change in) the ratio of helpful votes to the sum of helpful and unhelpful votes (between consecutive reviews for the same user). Sample restricted to reviews submitted before 2015 (see Appendix A.1). Topic distributions assigned separately to "Pros" and "Cons" of each review, using an LDA algorithm. For each topic, "pro minus con" is the review-specific difference in these two distributions. Standard errors clustered at the employer (former x latter industry) level. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. For a summary of the dataset, see Table A-1.

both the review being voted on and the review written by the voting user. We measure the correlation in mass across the 20 job attributes for each review-voter pair. The correlation equals 1 when the review on which the jobseeker is voting divides its attention across the attributes in the same shares as the jobseeker did in her own past review of a different job.[24] In these cases, the review prioritizes discussion of job attributes in the same way that the jobseeker did.

Table 5: Coincidence of Preferences: Individual (Un)Helpful Votes

	$\mathbb{1}$\{Voted helpful | voted helpful or unhelpful\}				
Review's rating of employer	-0.085***	-0.082***	-0.086***	-0.099***	
	(0.005)	(0.005)	(0.004)	(0.004)	
Correlation of review's attributes with voter's own		0.051***	0.055***	0.053***	0.006***
		(0.008)	(0.007)	(0.006)	(0.002)
Voter FE			✓	✓	
Reviewed employer FE				✓	
Review being voted on FE					✓
N	182410	182410	181219	179623	130527
Adjusted R^2	0.11	0.11	0.35	0.46	0.92

Notes: Sample consists of a panel of (un)helpful votes for different employer reviews. The dependent variable is a dummy variable that the user up-voted the review helpful. Because the dataset consists only of helpful and unhelpful votes—meaning it excludes decisions where no vote was given—this dummy is conditional on submitting a vote. Sample is restricted to voting users who submitted at least one of their own employer reviews on the website prior to submitting the (un)helpful vote. For each voters' own review and the review being voted on, the review text (combination of "Pros" and "Cons" sections) is apportioned into 20 topics. Correlation of review's topics with voter's own reflects the correlation of these two distributions across the 20 topics. Standard errors are clustered at the voting-user level. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. For a summary of the dataset, see Table A-1.

Conditional on the review's overall employer rating, jobseekers tend to find reviews more helpful when the reviews supply information about the attributes they prioritize. The result is robust to adding fixed effects for the jobseeker and comparing her votes across multiple reviews, as well as to adding fixed effects for the firm being reviewed (Columns 2–4). Comparing across reviews controlling for the firm and for the voter (Column 4), a review with a one-star higher employer overall rating lowers the probability of the review's being determined helpful rather than unhelpful by 9.9 percentage points. If the review's focus aligns perfectly with jobseekers' preferences, then jobseekers are 5.3 percentage points more likely to vote the review as helpful—rather than unhelpful—compared with a jobseeker whose preferences are uncorrelated with the review's prioritization across attributes. Finally, we

[24] Index review shares and voter preferences for pairing i by r and v, respectively, and $\{p_{ia}\}_{a=1}^{20}$ represents the extent to which i divides attention across attributes. Let σ_i represent the standard deviation of p_{ia} across a. Then, the review-voter pair correlation is $[(19)\sigma_r\sigma_v]^{-1}\sum_{a=1}^{20}(p_{ra}-1/20)(p_{va}-1/20)$.

focus on the subsample of reviews with multiple voters, include fixed effects for the review being voted on, and find that voters whose preferences line up more closely with the review's prioritization of attributes are more likely to vote it as helpful. However, the relationship is an order of magnitude weaker, highlighting that some attributes are universally more helpful than others, regardless of jobseeker priorites; still, individual-specific attribute rankings matter somewhat. This is some evidence against a theory that all the action is idiosyncratic matching, rather than ranking.

4.2 Why is information on some attributes more helpful?

As reported in Table 4, information about some attributes of firms play a more important role in explaining variation in jobseekers' votes of review helpfulness than information on other attributes. We conjecture that these are attributes about which jobseekers care more, those with greater marginal utilities. In this section, we provide corroborating evidence for that conjecture by analyzing which attributes play a more important role in explaining variation in volunteers' job satisfaction ratings. This is a different sample of workers making a different kind of decision. In theory, the relative importance of different attributes for worker average preferences should play an important role in both. To what extent is this the case?

To narrow the scope of potential confounders across employers and job opportunities, we focus our analysis across multiple reviews within each firm and look at how job satisfaction varies for workers within the same job as a function of their evaluation of the jobs' attributes. We use the full sample of volunteered reviews and study how the overall rating of a job varies with the attributes of each job. Controlling for the firm and the job associated with each review implicitly captures differences in satisfaction that might arise along firm and worker unobservables, such as implicit worker sorting patterns across jobs, and differences in pay. The coefficient on each attribute a_{ir} in predicting the overall rating worker-i assigns in review-r (R_{ir}) captures the extent to which a a worker's job satisfaction, relative to the firm average, changes as their evaluation of attribute-a changes by a unit. For robustness, we consider an additional specification in which we consider how, for the same worker, the change in a volunteers' attribute evauations relates to the change in their overall rating across employers.

The relative importance of the 20 attributes in explaining volunteers' overall job satisfaction is highly correlated with their relative importance in explaining which reviews jobseekers find most helpful. Figure 6 shows the latter on the horizontal and former on the vertical.[25] The correlation of the magnitudes of attributes' coefficients in explaining two different kinds

[25] Table A-3 reports full estimates for volunteer job-satisfaction models paralleling the analysis in Table 2 for review helpful share models for jobseekers.

of outcomes across two sets of workers is 0.89. This bundle of job attributes does have explanatory power for volunteer job satisfaction in a way that largely parallels their explanatory power for the helpfulness of information to jobseekers.

Figure 6: Correlation between Attributes' Coefficients from Different Models

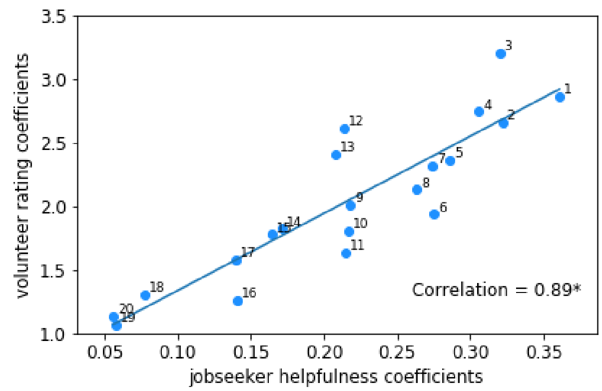

Notes: Coefficients for each topic reflect the absolute value of the "pro minus con" coefficients. Coefficients for jobseeker helpfulness correspond to Column 2 of Table 4. Coefficients for volunteer rating correspond to Column 2 of Table A-3. Topics are ordered in descending rank according to their coefficients for jobseeker helpfulness. Asterisk reflects that Pearson's correlation coefficient is statistically significant at the 1% level.

4.2.1 Imbalance in supply and demand for information?

Given evidence about the distribution across attributes of the supply of information by volunteers (Figure A-3) and evidently-reliable measures of the relative value workers place on information about these attributes (Figure 6), we summarize how well demand for and supply of information across attributes balance using two figures.

To measure relative demand for information across attributes, we focus on coefficients from the jobseeker-helpfulness model.[26] The relative importance of each attribute is displayed along the horizontal axis of panel (a) of Figure 7. If all 20 attributes were equally important, they would each have 0.05 weight. In fact, there's a lot of variation in relative demand, per the coefficients in Column 2 of Table 4.

The relative supply of each attribute is calculated as each attribute's share of total review text (Pros + Cons) that is assigned to an attribute. The vertical axis expresses each attribute's share of information supplied in reviews. The correlation between what's demanded and what's supplied is not statistically significant at 0.20, suggesting only weak alignment between the attributes jobseekers tend to want to understand and the attributes volunteers tend to describe.

[26] Results are similar using the volunteer job satisfaction model or attribute coefficients averaged across the two models.

Figure 7: Balance between Demand and Supply of Information across Attributes

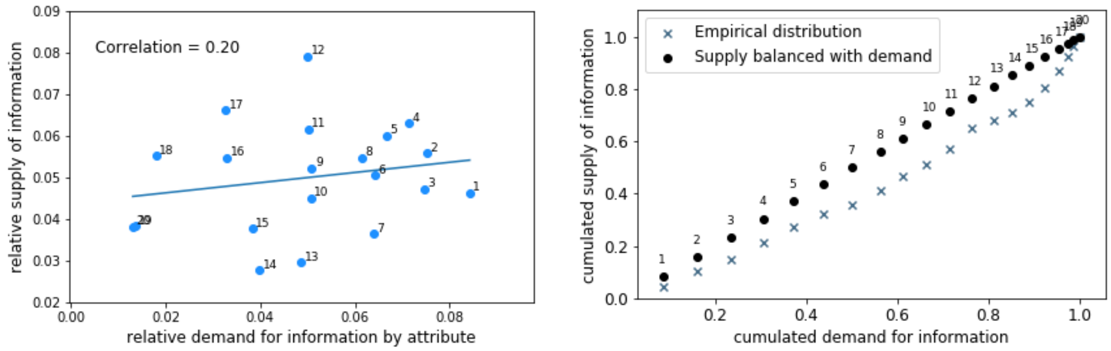

Relative supply and demand
(a) for information on each attribute

Cumulative supply and demand
(b) for information across attributes

Notes: The relative demand for each attribute is defined as the coefficients for jobseeker helpfulness (Table 4: Column 2) normalized to sum to 1. The relative supply for each attribute is defined as each attribute's share of total review text (Pros + Cons) that is assigned to an attribute (normalization of Figure A-3). Black dots in panel (b) form a 45-degree line indicating perfect balance between supply and demand. Topics are ordered in descending rank according to their relative demand.

Panel (b) of Figure 7 uses a Lorenz curve to clarify more precisely the magnitude and sites of imbalance. Attributes are ranked from most to least demanded. The most-demanded attribute's relative weight is expressed as its horizontal location and its share of information supplied in reviews by its vertical position. The 45-degree line represents balance between supply and demand, and a point's position relative to this expresses the extent of imbalance. The most-demanded attribute receives 8.4 percent of weight in demand but only 4.6 percent of weight in supply. Here, there is only a small degree of under-supply. Next, we cumulatively add the second most-demanded attribute, so the new horizontal and vertical positions express the sum of it and the more-demanded attribute. Here, there is a larger degree of under-supply. We continue this iterative process across all 20 attributes. The most-demanded half of attributes account for 66 percent of total demand weight but only 51 percent of supply weight. In contrast, the least-demanded attributes (represented furthest right and more closely clustered horizontally) have the steepest increases in cumulative supply, meaning that these are the most over-supplied. This reveals some scope for institutional changes to promote a better match between the supply of and demand for information.

4.3 Why the stronger demand for negative information?

One reason jobseekers might have a stronger demand for more-negative information is that workers undersupply it, but the other channels we contemplate—risk aversion and the tendency of employers to supply positive information—would also push this way. This section discusses and offers some evidence on these mechanisms but does not try to quantify relative contributions. The point is only to build an understanding about the process and acknowledge that, given employer preference for supplying information and worker risk aversion, it would be surprising if jobseekers expressed no preference for negative information from volunteers even without volunteers censoring or degrading of negative information. Volunteers' tendency to do so likely contributes to making jobseekers' even hungrier for this kind of information.

First, consider risk aversion. Consider a risk-neutral jobseeker for whom each new review is an independent, mean-zero signal about the true quality of an employer and who so far has seen n ratings of an employer with average $\bar{R}_n$. Consider the $n+1^{th}$ signal: R_{n+1}. The value of this signal to the jobseeker does not depend on the particular value of R_{n+1} or the difference $R_{n+1}-\bar{R}_n$. A signal leads to an updating of the prior and increases belief precision. If it does this by confirming that the prior average was correct, this is as valuable as one that leads to an update up or down. On the other hand, if the jobseeker is risk averse, then the value of a signal that shifts her mean posterior belief down ϵ is more valuable than one that shifts it up by the same amount. Avoiding an ϵ-worse outcome is worth more than gaining an ϵ-better one. That workers tend to be risk averse seems uncontroversial and likely to be part of the explanation.

Second, an obvious channel by which jobseekers get information about a job is through communication with the hiring firm. From Glassdoor, we obtained a sample of 413,846 job postings for which we are able to examine the job description text associated with each posting.[27] These job descriptions can be interpreted as advertisements for the firm. As one might expect, firms write job postings full of positive language. Sentiment analysis of these job descriptions confirms this. Job description text is less positive than the text in reviews' "Pros" field but far more positive than the text of reviews' "Cons" field (Figure 8). Almost no employer advertises their vacancies with negatively-charged text: less than 1 percent of job descriptions are interpreted as relaying negative sentiment. It's in the firm's interest to supply positive, not negative, information.[28] Without a reputation system in place where

[27] Glassdoor collects job postings from three main sources—online job boards, applicant tracking systems, and company websites—and captures about 81 percent of total U.S. job openings, as measured in the Job Openings and Labor Turnover Survey conducted by the U.S. Bureau of Labor Statistics (Chamberlain and Zhao, 2019).

[28] Firms have incentives to supply information that horizontally differentiates them from competitors in workers' eyes. This can improve fit. They do not have incentives to supply information that vertically differentiates them in a negative way. High-road firms' challenge is to credibly differentiate themselves from others. Institutions, including Glassdoor, may help communicate both kinds of information. Further, note that what distinguishes a dimension of horizontal versus vertical differentiation is primarily the correlation in workers' tastes. More highly correlated tastes imply greater vertical differentiation. Less correlated tastes imply greater horizontal differentiation.

jobseekers can obtain information about the firm from experienced sources, presumably the only employer-specific information non-referral workers would have access to would be job advertisements and public media coverage. The former, as evidenced by Figure 8, contains little if any negative content, while the latter would likely be limited in scope to only large, well-known employers.

Figure 8: The Distribution of Emotional Polarity for Job Posting Descriptions, Cons Section of Reviews, and Pros Section of Reviews

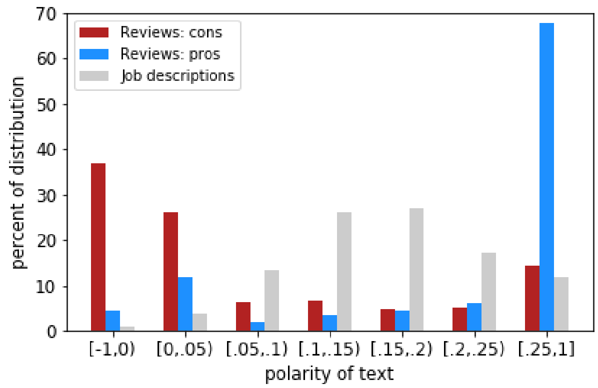

Notes: Polarity of the "Pros" section for employer reviews, "Cons" section for employer reviews, and job posting descriptions are measured through natural langugage processing using the *TextBlob* library in Python. The polarity measure ranges from -1 to 1, with more postive (negative) values relecting more postively (negatively) charged text, and is partitioned into seven bins. The leftmost, center, and rightmost bars reflect the share of reviews within a polarity bin for the "Cons" field, "Pros" field, and job descriptions, respectively. Bars sum to 1 for each text category.

4.3.1 How well does wage alone proxy for attributes?

Analyzing Glassdoor employer reviews and pay reports helps illuminate the extent to which wage and non-wage attributes of jobs co-vary and the extent to which wage alone is a good proxy for job value to workers. Job opportunities differ across firm- and job-specific attributes, affecting the quality of the match. These attributes are not unidimensional (Table A-3) and differ in their relative importance to jobseekers and to those evaluating their current match (Figure 6). While certain attributes appear more important on average, all of the attributes we extract from the review text appear important to workers to some extent.

Suppose jobseekers are choosing from a menu of J job opportunities $\{(w, \vec{a})\}_J$, where w is a wage offer and $\vec{a}$ a vector of attributes for each job. In a compensating-differentials model, workers choose the job that maximizes their utility $u(w, \vec{a})$. We assume u increases in w and each component of $\vec{a}$ because we define our measures of $\vec{a}$ as the difference between pros and cons. Workers are willing to trade-off between inputs in ways governed by u's curvature. Workers might happily give up some w to enjoy improved attributes. Employers could then strategically post jobs in which wages and attributes move inversely, depending on the firm's marginal costs of providing amenities (Rosen, 1986). In contrast, consistent with Mortensen (2003) and Maestas et al. (2018), firms may also compensate workers along multiple attribute dimensions so that jobs with higher wages also offer better attributes. As Sorkin (2018) poses it, the Rosen motive might push for negative correlation between wage and non-wage amenities, while a countervailing Mortensen motive might push toward a positive correlation.

To understand the relationship between wages and attributes within and across firms, we estimate employer fixed effects from models in which employee wages and attribute evaluations are regressed on industry–job title and employer fixed effects in a fashion similar to (Abowd et al., 1999) (AKM).[29] For 100,827 firms, we estimate employer-specific wage premia w_j^{FE} by extracting employer fixed effects from base pay and employer-specific attribute premia a_j^{FE} by extracting employer fixed effects from each review's overall rating, pro share of review text, and pro-minus-con evaluation of each attribute. Looking across employers and relating w_j^{FE} to a_j^{FE} provides a measure of correlation between wages and attributes across firms. To what extent are non-wage aspects of jobs correlated with firm tendency to pay more than others?

Table 6 summarizes the results. Firm fixed effects of log base salary summarize a firm-specific wage premium and vary across firms with a standard deviation of 0.18 log point (Column 2), corresponding to 20 percent higher wages, and with a 0.48 log point difference between the 90th and 10th percentiles (Column 3). Firm fixed effects for each of the other variables will be related to the wage fixed effects and reported in the subsequent columns.

Across firms, fixed effects for two measures of overall job quality, which might be considered proxies for worker utility from the job, are positively correlated with firm wage premia. Consider the distribution of firm fixed effects in overall ratings. The standard deviation is 0.7 star, and the 90–10 range is 2.0 stars. Firms' overall rating fixed effects are correlated 0.17 with their wage fixed effects (Column 4), and a simple regression of the former on the latter yields a slope coefficient of 0.61. Firms with a one standard deviation higher wage premium tend to have 0.11 star higher firm fixed effects in overall rating. This is 0.16 of a standard deviation of overall-rating firm fixed effects. This is a strong, positive relationship,

[29]Employee wages are not submitted as part of an employer review, but rather are available separately through Glassdoor pay reports data. Because the employer reviews and pay reports data have rich cross-sections but thin panels of repeat worker responses, we incorporate industry-job title fixed effects in lieu of employee fixed effects when estimating employer fixed effects. The sample of pay reports is restricted to full-time U.S. private-sector workers and wages constitute annualized base pay which excludes performance-based and overtime compensation.

Table 6: Firm-Specific Premiums in Compensation and Amenities

	AKM	Distribution of FE		Relation to base pay FE		
Firm fixed effects	R^2	standard deviation	90th–10th percentiles	correlation	slope	standard error
Log base pay	0.86	0.18	0.48	1	–	–
Employee rating of firm	0.22	0.7	2.0	0.17	0.61***	0.01
Pro share of text	0.16	8.3	24.3	0.05	2.21***	0.14
Pro minus con : attribute 1	0.10	3.0	7.6	–0.09	–1.44***	0.05
Pro minus con : attribute 2	0.13	3.8	10.0	0.02	0.47***	0.07
Pro minus con : attribute 3	0.08	2.5	6.4	0.08	1.16***	0.04
Pro minus con : attribute 4	0.10	3.5	9.9	0.05	0.98***	0.06
Pro minus con : attribute 5	0.08	2.8	7.5	–0.03	–0.45***	0.05
Pro minus con : attribute 6	0.08	2.5	6.0	–0.13	–1.81***	0.04
Pro minus con : attribute 7	0.07	2.0	4.9	0.08	0.87***	0.04
Pro minus con : attribute 8	0.08	3.0	7.5	0.06	1.01***	0.05
Pro minus con : attribute 9	0.08	2.6	6.3	0.10	1.37***	0.04
Pro minus con : attribute 10	0.08	2.5	6.1	0.08	1.05***	0.04
Pro minus con : attribute 11	0.09	3.5	8.8	–0.03	–0.55***	0.06
Pro minus con : attribute 12	0.11	4.2	11.5	0.12	2.71***	0.07
Pro minus con : attribute 13	0.09	1.9	4.1	0.00	0.01	0.03
Pro minus con : attribute 14	0.08	1.7	4.0	0.01	0.08***	0.03
Pro minus con : attribute 15	0.09	2.0	4.7	0.03	0.29***	0.04
Pro minus con : attribute 16	0.08	2.3	5.1	0.05	0.59***	0.04
Pro minus con : attribute 17	0.11	3.8	10.0	–0.09	–1.77***	0.07
Pro minus con : attribute 18	0.08	2.9	7.6	–0.03	–0.49***	0.05
Pro minus con : attribute 19	0.08	1.9	4.2	0.04	0.41***	0.03
Pro minus con : attribute 20	0.09	1.7	3.8	–0.02	–0.15***	0.03

Notes: There are 100,827 firms for which pay and attribute fixed effects are estimated. The mean number of employer reviews and pay reports for each firm is 45 and 35, respectively. Correlations and regressions between attribute and pay fixed effects are weighted by the square root of the geometric mean of each firm's employer review total and pay report total. Given the relative paucity of users leaving multiple pay reports or employer reviews, firm fixed effects for each measure are from an AKM-style specification, where industry-job title fixed effects are incorporated in lieu of worker fixed effects. For employer reviews, fixed effects for metro and year-month are included; for base pay, a quadratic in years of specific experience and fixed effects for year, metro, gender, educational attainment, and pay frequency are included.

but much of the firm variation in overall ratings, proxying for worker utility, is not explained by wage premia. Similar results arise when using the pro-share of volunteers' text reviews as the overall quality measure. Firm fixed effects in this have different units (percentage points as opposed to stars), and the correlation with wage fixed effects is weaker. Overall job quality depends on both wage and non-wage attributes. Dispersion across firms in measures of overall job quality reflect dispersion in both their tendency to pay differently and their tendency to provide other desirable attributes.

Next, we estimate firm fixed effects for each of the 20 attributes and relate these to firms' wage fixed effects. While firms that pay more also tend to be better on many non-wage attributes, some attributes tend to be worse at higher-paying firms. In particular, consider attribute 1, the attribute that jobseekers implicitly revealed as most important in predicting the helpfulness of a review. Higher-paying firms tend to be lower on attribute 1 such that firms with a one-standard-deviation-higher wage premium tend to be 0.38 [=(0.18 * –1.44)/3.0] standard deviation lower along attribute 1. The attribute jobseekers most wanted to learn about is one negatively correlated with wage. Furthermore, volunteers implicitly rank attribute 1 the second most important in driving overall firm rating (Figure 6).

These results are consistent with both Rosen and Mortensen motives operating. Higher-paying firms offer better job attributes in some dimensions, but the wage premium could also reflect a need to compensate employees with higher wages in order for them to accept being worse off with respect to other attributes. What kind of information would be most helpful for jobseekers? Precisely that which illuminates some aspect of a job that is both negatively correlated with an easily observed and ranked component of job offers (pay) and crucial for driving overall job satisfaction. Attribute 1 fits both of those conditions.

5 Conclusion

We add new evidence that workers face higher expected costs of supplying negative information about a current or former employer than positive information. This asymmetry leads to a relative under-supply of negative information—which is in high demand—as well as the degradation of the information supplied. Threats to each worker's exercise of voice can increases other workers' uncertainty about job quality at prospective employers and reduce their expected value of exit. Institutions that improve the information flow to workers can, in theory, increase aggregate labor supply as well (Benson et al. (2020)). We harness new data to generate some new empirical findings about basic relationships.

In addition to the widespread use of nondisclosure agreements (Starr et al., 2019), the evidence points toward a problem with the flow of private information between workers about employers. Because jobs matter so much to workers' lives, they have incentives to guard against risks of employer retaliation. This may lead them to refrain altogether from supplying negative information or to conceal information in ways that degrade the value of the negative information supplied. The evidence developed here—along with that in Marinescu et al. (2018)—suggests that the negative information that is actually supplied may be just the tip of an iceberg of negative information that remains largely out of view, owing to firms' and workers' reluctance to supply it as volunteers. It suggests the possible value of improved institutions to better elicit, aggregate, and distribute workers' private information and raises the question how much workers value the marginal information about jobs. This area is ripe for fruitful links between theory and empirics.

Difficulty in obtaining information creates labor market friction. Lack of credible information can prevent workers from pursuing or even accepting job offers that would have higher value to the worker than their better-understood current job. In search models, this is frequently represented by uncertainty about uni-dimensional match quality about which workers receive unbiased signals. The worker will have a distribution of subjective beliefs about how this prospective job will turn out. Even if beliefs are centered on the truth, many beliefs would generate lower expected utility than the current situation and make a risk-averse worker reluctant to pursue an alternative. Even if a worker knows her current job is bad, she may not trust another employer's promising offer. The devil you knows is better than the devil you don't. Increased uncertainty is comparable to a larger mobility cost. Institutions that reduce uncertainty should increase efficiency and reduce mobility costs. If workers face adverse selection in choosing employers such that there are low-road employers that break their promises but successfully pool with high-road employers that fulfill their promises, the worker should discount the promise and have mean expectations below the promise. In this case, institutions that help workers distinguish firms that treat workers well versus ill would have value (Carmichael, 1984; Benson et al., 2020).

The incompleteness of employment contracts means that conditional on present terms, workers' beliefs about how future adjustments at the firm would be made is critical to their evaluation of a firm's job offer. A convenient way to index this is allowing the parameter governing the worker's share of a job's value, the bargaining-power parameter β, to vary across firms. For instance, it might vary across firms because of variation in bosses' altruism towards workers, the strength of workers' unions, or different regulatory regimes. In many standard search models, the value of a job offer is perfectly correlated with the wage offer because firms are assumed to share a common β. However, if β varies across firms, equal wage offers could come from a firm with higher β and lower idiosyncratic productivity shock or one with a lower β and higher shock. The wage isn't sufficient to rank offers. A worker with two equal wage offers might prefer the former scenario because she expects faster productivity growth and looks to capture a greater share of the returns to this growth.[30] Information that helped her infer which type of firm β_f was making the offer would be useful in assessing the value of the offer.

[30] Consider a standard wage-setting equation in a search model with idiosyncratic productivity shocks, as in equation (2.10) of Pissarides (2000), $w(x) = (1-\beta)z + \beta p(x + c\theta)$. The offered wage conditional on the idiosyncratic productivity shock $w(x)$ depends on the interaction between the worker's share of job value β and the productivity shock x. A worker's knowledge of w does not reveal (β, x).

Board and Meyer-ter Vehn (2013) show how reputational incentives depend on the particular signal arrival process. Our finding that negative signals may be more likely to be under-supplied or supplied in a noisier, degraded form suggests that employer reputation may evolve more by "good news" than "bad news" arrival. This would predict a work-shirk equilibrium where employers work to keep their reputations above a threshold level but then shirk so as not to push it far above, rather than a shirk-work equilibrium where employers with reputations above a threshold keep investing to build their reputations but those below it do not.

The difficulty in establishing a market for firm reputation would certainly not be due to the information lacking aggregate value. Rather, at least in part, the sheer number of employers, workers, and the high degrees of differentiation in tastes, productivities, and amenities between them makes the information space required to characterize all jobs for all workers quite complex. Demand for and supply of information is highly differentiated and thinly distributed. There is no mass market for a piece of labor-market information. It's all long tails, with only a tiny share of people interested in any one piece of information and only a tiny share of others possessing it. Lots of small transactions would have to be coordinated between strangers.

In terms of information problems in the labor market, labor economics generally and personnel economics particularly have focused heavily on the manager's information problem of choosing among workers who appear the same but actually differ in unobservable ways. Workers differ in unobservable type ex ante (adverse selection, unobserved productivity) and in unobservable strategy ex post (moral hazard, unobserved effort). Abundant theory and empirical knowledge have developed how these asymmetries create inefficiencies and how employers attempt to deal with these challenges using screening mechanisms (credentials, interviews, monitoring) and incentive mechanisms (pay-for-performance, promotion tournaments, job security). As Oyer et al. (2011) note, "Personnel Economics has grown up largely within leading business schools... Because many researchers in this field must take their insights into MBA classrooms and offer advice to future managers, Personnel Economists are typically interested in how firms can solve human resource management problems and how the solutions to HR problems are related to firms' broader strategic contexts."

Our knowledge about the worker's information problem remains less developed than the manager's problem. Economic theory depicts workers as choosing among possible employment opportunities based on known compensation (and perhaps amenities), but employers also differ in unobservable type and strategy. With the rise of online labor market platforms— where employer information can be readily shared, aggregated, and consumed— information flow has improved but there's reason to believe it remains quite imperfect. A market for information about jobs is missing, making the full quality of jobs difficult if not impossible to observe. Nonetheless, the rise of new institutions that make information exchange visible to economists creates new opportunities for improving our understanding of the worker's information problem.

References

J. M. Abowd, F. Kramarz, and D. N. Margolis. High wage workers and high wage firms. Econometrica, 67(2):251–333, 1999. [Cited on pages 6 and 32.]

J. G. Altonji and C. R. Pierret. Employer learning and statistical discrimination. The Quarterly Journal of Economics, 116(1):313–350, 2001. [Cited on page 8.]

K. Benner. Abuses hide in the silence of nondisparagement agreements. https://www.nytimes.com/2017/07/21/technology/silicon-valley-sexual-harassment-non-disparagement-agreements.html, 2017. [Cited on page 3.]

A. Benson, A. Sojourner, and A. Umyarov. Can reputation discipline the gig economy?: Experimental evidence from an online labor market. Management Science, 66:1802–1825, 2020. [Cited on pages 7, 34, and 35.]

D. Blei, A. Ng, and M. Jordan. Latent dirichlet allocation. Journal of Machine Learning Research, 3:993–1022, 05 2003. [Cited on page 10.]

S. Board and M. Meyer-ter Vehn. Reputation for quality. Econometrica, 81(6):2381–2462, 2013. [Cited on page 35.]

J. Brown and D. A. Matsa. Boarding a sinking ship? an investigation of job applications to distressed firms. The Journal of Finance, 71(2):507–550, 2016. [Cited on page 7.]

H. L. Carmichael. Reputations in the labor market. The American Economic Review, 74(4):713–725, 1984. ISSN 00028282. [Cited on pages 7 and 35.]

A. Chamberlain and D. Zhao. Methodology: Glasdoor job market report. Technical report, Glassdoor, April 2019. [Cited on page 30.]

M. Chen. This labor-trafficking case exposes the twin cities' seedy subcontracting underbelly. https://www.thenation.com/article/archive/this-labor-trafficking-case-exposes-the-twin-cities-seedy-subcontracting-underbelly/, 2018. [Cited on page 8.]

L. M. Cortina and V. J. Magley. Raising voice, risking retaliation: Events following interpersonal mistreatment in the workplace. Journal of Occupational Health Psychology, 8(4):247–265, 2003. [Cited on page 8.]

B. Covert. Mcdonald's has a real sexual harassment problem. The Nation, 2020. [Cited on page 8.]

J. Eidelson. Covid gag rules at u.s. companies are putting everyone at risk. https://www.bloomberg.com/news/features/2020-08-27/covid-pandemic-u-s-businesses-issue-gag-rules-to-stop-workers-from-talking?sref=a5mft8DW, 2020. [Cited on page 3.]

O. Hart. An economist's perspective on the theory of the firm. Colum. L. Rev., 89:1757, 1989. [Cited on page 2.]

H. Hong, D. Xu, G. A. Wang, and W. Fan. Understanding the determinants of online review helpfulness: A meta-analytic investigation. Decision Support Systems, 102:1 – 11, 2017. [Cited on page 4.]

H.-s. Hwang, D. T. Mortensen, and W. R. Reed. Hedonic wages and labor market search. Journal of Labor Economics, 16(4):815–847, 1998. [Cited on page 8.]

M. S. Johnson, D. Schwab, and P. Koval. Legal protection against retaliatory firing improves workplace safety. Available at SSRN 3619384, 2020. [Cited on page 3.]

T. K. T. Liu, C. Makridis, P. P. Ouimet, and E. Simintzi. Is cash still king: Why firms offer non-wage compensation and the implications for shareholder value. 2017. [Cited on page 20.]

O. Lobel. Exit, voice & innovation: How human capital policy impacts equality (& how inequality hurts growth). Hous. L. Rev., 57:781, 2019. [Cited on pages 3 and 18.]

A. Lopez-Lira. Risk factors that matter: Textual analysis of risk disclosures for the cross-section of returns. Available at SSRN 3313663, 2019. [Cited on pages 10 and 43.]

N. Maestas, K. J. Mullen, D. Powell, T. von Wachter, and J. B. Wenger. The value of working conditions in the united states and implications for the structure of wages. Technical Report 25204, National Bureau of Economic Research, October 2018. [Cited on pages 5, 6, and 32.]

A. Marchetti. Firms of a feather merge together: The role of acquirer-target culture compatibility in technology acquisitions. Job Market Paper, 2019. [Cited on pages 4, 10, and 11.]

I. Marinescu, N. Klein, A. Chamberlain, and M. Smart. Incentives can reduce bias in online employer reviews. Academy of Management Proceedings, 2018(1):11189, 2018. [Cited on pages 1, 18, and 34.]

A. Mas and A. Pallais. Valuing alternative work arrangements. American Economic Review, 107 (12):3722–59, 2017. [Cited on pages 5 and 6.]

G. Menzio and S. Shi. Efficient search on the job and the business cycle. Journal of Political Economy, 119(3):468–510, 2011. [Cited on page 7.]

D. Mortensen. Wage dispersion: why are similar workers paid differently? MIT press, 2003. [Cited on page 32.]

P. Nelson. Information and consumer behavior. Journal of political economy, 78(2):311–329, 1970. [Cited on page 7.]

M. Olson. The Logic of Collective Action: Public Goods and the Theory of Groups, Second Printing with a New Preface and Appendix, volume 124. Harvard University Press, 2009. [Cited on page 2.]

P. Oyer, S. Schaefer, et al. Personnel economics: Hiring and incentives. Handbook of Labor Economics, 4:1769–1823, 2011. [Cited on pages 8 and 36.]

C. A. Pissarides. Equilibrium unemployment theory. MIT press, 2000. [Cited on page 35.]

S. Rosen. The theory of equalizing differences. Handbook of Labor Economics, 1:641–692, 1986. [Cited on pages 8 and 32.]

C. Shapiro and J. E. Stiglitz. Equilibrium unemployment as a worker discipline device. The American Economic Review, 74(3):433–444, 1984. [Cited on page 8.]

J. Silver-Greenberg and N. Kitroeff. How bloomberg buys the silence of unhappy employees. New York Times, 2020. [Cited on page 3.]

J. Sockin and M. Sockin. Job characteristics, employee demographics, and the cross-section of performance pay. mimeo UPenn and UT Austin McCombs School of Business, 2019a. Working paper. [Cited on page 20.]

J. Sockin and M. Sockin. A pay scale of their own: Gender differences in performance pay. mimeo UPenn and UT Austin McCombs School of Business, 2019b. Working paper. [Cited on page 20.]

I. Sorkin. Ranking firms using revealed preference. The Quarterly Journal of Economics, 133(3): 1331–1393, 2018. [Cited on pages 6, 8, and 32.]

M. Spence. Competition in salaries, credentials, and signaling prerequisites for jobs. The quarterly journal of economics, 90(1):51–74, 1976. [Cited on page 8.]

E. Starr, J. J. Prescott, and N. Bishara. Noncompetes in the U.S. Labor Force. U of Michigan Law & Econ Research Paper, (18-013), 2019. [Cited on pages 3, 18, and 34.]

U.S. Occupational Safety and Health Administration. Recommended practices for anti-retaliation programs. osha.gov/Publications/OSHA3905.pdf, 2017. [Cited on page 3.]

P. Walsh. Charges: Twin cities contractor threatened to report his undocumented workers if they complained. https://www.startribune.com/charges-twin-citis-contractor-threatened-to-report-his-undocumented-workers-if-they-complai
494386221/, 2018. [Cited on page 8.]

D. Weil and A. Pyles. Why complain-complaints, compliance, and the problem of enforcement in the us workplace. Comp. Lab. L. & Pol'y. J., 27:59, 2005. [Cited on pages 2 and 3.]

6 Appendix

A.1 Review helpfulness

The first employer reviews uploaded on the Glassdoor website are from mid-2008. At that time, users of the website could up-vote an employer review as helpful or down-vote the review as unhelpful, if they wanted to provide feedback for other users on the site about how they valued the information. The combination of these two measures—helpful and unhelpful votes—provides a review-specific metric of how helpful the review's content was to jobseekers. Starting in 2015, however, the option to down-vote an employer review as unhelpful was phased out on the website. Looking at the share of total helpfulness votes submitted each year that were attributable to unhelpful votes shows a clear structural break at the end of 2014 (Figure A-2). This motivates the decision to restrict the sample of reviews for which the share helpful, $\frac{\#helpful}{\#helpful+\#unhelpful}$, is the measure of interest to pre-2015.

The ability to submit unhelpful votes was not completely eliminated from the platform, however. As evidenced in Figure A-2, unhelpful votes still accounted for a sliver of total helpfulness votes submitted (under 2 percent) from 2015 through the first half of 2017. These unhelpful votes were generated through users submitting down-votes through Glassdoor's mobile application, which is accessible via phone or tablet. The dataset that tracks individuals' helpulness votes (used in Tables 3 and 5) stems from this period. By the end of 2017, the option to down-vote an employer review as unhelpful was completely phased out.

A.2 Implementing LDA algorithm

Implementing a Latent Dirichlet Allocation (LDA) algorithm of text-based documents typically requires some data-cleaning processes, as well as making decisions that are left to the discretion of the researcher. First, we append the "Pros" and "Cons" sections of each review to form a single document. The data cleaning processes we then implement are the following: converting each document to lower case, removing any punctuation, tokenizing the text to split each review into a vector of words, dropping tokens that are less than three characters, dropping well-accepted stopwords that are ubiquitous but typically unmeaningful (with the exception of the word "off" in order to capture attributes like paid time off), and lemmatizing tokens such that the algorithm recognizes commonality in words that are highly similar but differ in their suffixes or tenses. Additionally, for each review, we remove the name of the employer from the review text by dropping any tokens that are equivalent to the name of the employer stored in Glassdoor's employer lookup table.

Once we have these tokenized documents, we impose the following additional conditions.

First, we combine unigrams (singular tokens) into bigrams (pairs of tokens) if two tokens occur in repetition for at least 1,000 unique instances across reviews. Then, we combine unigrams and bigrams into trigrams (triples of tokens) using the same threshold of repetition. We then create a dictionary from the universe of tokens, restricting our attention to the 100,000 most common tokens. We impose some additional filters to remove outliers by excluding tokens that occur in fewer than 10 reviews but also any tokens that occur in more than 25 percent of reviews. This results in a final dictionary of 76,659 unique tokens. Finally, we implement an LDA algorithm for 20 attributes and apportion each review's text for the "Pros" and "Cons" sections separately across the 20 attributes.

A.3 Additional tables

Table A-1: Summary Statistics for Employer Reviews Dataset

Measure of Interest	Summary Statistics				
	N	mean	standard deviation	p10	p90
A. Dependent variables					
Star rating	6,809,319	3.35	1.40	1	5
Helpful votes	6,809,319	1.24	3.99	0	4
Helpful votes pre-2015	685,505	2.76	6.59	1	6
Unhelpful votes pre-2015	685,505	0.34	1.09	0	1
Share helpful votes pre-2015	685,505	0.86	0.31	0.29	1
B. Firm characteristics					
Firm age (years)	5,334,194	56.3	51.8	10	131
Firm employment (1000s)	6,408,571	53.8	199.6	37.0	140.0
C. Panel of users with multiple reviews					
Next star rating	441,570	2.91	1.60	1	5
Delta star rating	441,570	0.002	1.84	-2	2
Delta helpful votes pre-2015	47,981	0.812	7.41	-3	5
Delta unhelpful votes pre-2015	47,981	-0.002	0.98	-1	0
Delta share helpful pre-2015	18,592	2.79	34.5	-20.0	40.0
D. Dummy variables					
Would recommend to a friend	5,621,590	0.5961	–	–	–
Approves of the CEO	4,157,888	0.4822	–	–	–
Postitive business outlook	5,065,925	0.4815	–	–	–
Is current employee	6,809,319	0.5229	–	–	–
Job title blank	6,809,319	0.4401	–	–	–
Job title anonymized	6,809,319	0.0057	–	–	–
Location blank	6,809,319	0.4222	–	–	-
E. Individual (un)helpful votes					
Voted helpful	182,410	0.852	–	–	–
Review's rating of employer	182,410	2.34	1.39	1	5
Correlation of review's topics with voter's own	182,410	0.198	0.318	-0.181	0.669

Measure of Interest	Summary Statistics				
	N	mean	standard deviation	p10	p90
F. Text-based variables					
Pro share of text	6,809,319	0.468	0.196	0.191	0.723
Character length	6,809,217	336.5	420.8	78	706
Polarity of text	6,809,217	0.22	0.24	-0.04	0.52
Subjectivity of text	6,809,217	0.54	0.17	0.35	0.75
Flesch Kincaid grade	6,158,522	9.2	4.1	4.4	15.1
Pro minus con : 1	6,809,217	-0.018	0.084	-0.110	0.006
Pro minus con : 2	6,809,217	0.019	0.098	-0.013	0.128
Pro minus con : 3	6,809,217	-0.010	0.081	-0.090	0.027
Pro minus con : 4	6,809,217	-0.039	0.104	-0.187	0.005
Pro minus con : 5	6,809,217	0.023	0.094	-0.007	0.139
Pro minus con : 6	6,809,217	-0.016	0.088	-0.112	0.007
Pro minus con : 7	6,809,217	0.004	0.069	-0.015	0.053
Pro minus con : 8	6,809,217	0.016	0.092	-0.008	0.116
Pro minus con : 9	6,809,217	0.013	0.088	-0.010	0.108
Pro minus con : 10	6,809,217	-0.012	0.081	-0.093	0.007
Pro minus con : 11	6,809,217	-0.006	0.100	-0.096	0.056
Pro minus con : 12	6,809,217	-0.042	0.122	-0.217	0.005
Pro minus con : 13	6,809,217	0.002	0.057	-0.014	0.022
Pro minus con : 14	6,809,217	0.007	0.053	-0.006	0.045
Pro minus con : 15	6,809,217	-0.006	0.068	-0.059	0.009
Pro minus con : 16	6,809,217	-0.015	0.091	-0.111	0.008
Pro minus con : 17	6,809,217	0.011	0.108	-0.069	0.123
Pro minus con : 18	6,809,217	0.006	0.090	-0.061	0.094
Pro minus con : 19	6,809,217	-0.004	0.070	-0.047	0.020
Pro minus con : 20	6,809,217	0.004	0.069	-0.010	0.039

1. Cross-section of employer reviews	
Dimension of Interest	Sample Coverage
Horizon window	2008-23 to 2019-35
Reviews (pre-2015)	6,809,319 (1,417,494)
Users	6,201,251
Employers	463,968
Job titles	400,279
Helpful votes (pre-2015)	8,459,920 (1,895,376)
Unhelpful votes pre-2015	231,837

2. Individual (un)helpful votes	
Dimension of Interest	Sample Coverage
Horizon window	Jan 2015–Sep 2017
Voted on reviews	92,320
Reviewed employers	10,117
Voting users	6,396
Helpful votes	155,490
Unhelpful votes	26,920

Notes: This table reports summary statistics and group counts for user-submitted employer reviews on the website Glassdoor. Panel of users with multiple reviews excludes job stayers as well as job switchers who submit both reviews in the same year-week. Distribution of review text into 20 topics is assigned using an unsupervised LDA model (see Lopez-Lira (2019) for an in-depth description of the algorithm). The LDA model is calibrated on the full text of reviews for the full dataset, then topic distributions are assigned separately for the "Pros" and "Cons" sections, with "pro minus con" reflecting the weighted difference between the pro and con distributions, where the pro (con) distributions are weighted according to the share of the review's text written in the pro (con) section. Individual (un)helpful votes comprised of a subsample of voters for whom: (i) data on each up-vote and down-vote are available, and (ii) an employer review was also submitted on the website before submitting the up- or down-vote.

Table A-2: Predicting Helpful Vote Count of Volunteer's Review

	Log count helpful votes						
Star rating	-0.157***						-0.094***
	(0.002)						(0.001)
Would recommend employer to a friend		-0.453***					-0.198***
		(0.007)					(0.005)
Approves of the CEO			-0.312***				-0.024***
			(0.006)				(0.001)
Positive business outlook for the firm				-0.321***			-0.007***
				(0.005)			(0.001)
Pro share of review text					-0.860***		-0.321***
					(0.009)		(0.004)
Polarity of text						-0.270***	0.005***
						(0.005)	(0.002)
Log length of review (total characters)	0.190***	0.194***	0.236***	0.224***	0.189***	0.226***	0.187***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Flesch kincaid reading grade	-0.002***	-0.002***	-0.002***	-0.002***	-0.002***	-0.003***	-0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Subjectivity of text	0.003**	-0.014***	-0.035***	-0.027***	-0.020***	0.048***	0.003
	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)
Employer, year-month FE	✓	✓	✓	✓	✓	✓	✓
N	5958112	4849978	3642220	4373338	5958112	5958112	3191431
Adjusted R^2	0.47	0.47	0.42	0.44	0.44	0.39	0.50

Notes: The table above predicts the helpfulness of an employer review using the logarithm of review votes helpful plus one as the dependent variable. Sample is restricted to reviews for which the Flesch Kincaid reading grade is non-negative and no greater than 20. Polarity and subjectivity of each review are measured through natural langugage processing using the *TextBlob* library in Python. Standard errors clustered at the employer level. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. For a summary of the dataset, see Table A-1.

Table A-3: Job Satisfaction: Effects of Volunteer Attribute Evaluations on Their Overall Rating of Employer

	Overall rating		First difference overall rating
Pro minus con: attribute 1	3.143***	2.860***	3.007***
	(0.014)	(0.011)	(0.028)
Pro minus con: attribute 2	2.864***	2.658***	2.976***
	(0.009)	(0.008)	(0.023)
Pro minus con: attribute 3	3.498***	3.200***	3.280***
	(0.015)	(0.014)	(0.031)
Pro minus con: attribute 4	2.996***	2.745***	2.636***
	(0.011)	(0.009)	(0.023)
Pro minus con: attribute 5	2.497***	2.361***	2.647***
	(0.012)	(0.009)	(0.024)
Pro minus con: attribute 6	1.848***	1.939***	1.680***
	(0.012)	(0.010)	(0.025)
Pro minus con: attribute 7	2.490***	2.315***	2.536***
	(0.012)	(0.012)	(0.032)
Pro minus con: attribute 8	2.314***	2.131***	2.403***
	(0.010)	(0.011)	(0.023)
Pro minus con: attribute 9	2.217***	2.013***	2.312***
	(0.016)	(0.013)	(0.023)
Pro minus con: attribute 10	2.016***	1.806***	1.846***
	(0.011)	(0.010)	(0.023)
Pro minus con: attribute 11	1.711***	1.633***	1.707***
	(0.008)	(0.006)	(0.035)
Pro minus con: attribute 12	2.865***	2.610***	2.434***
	(0.011)	(0.009)	(0.026)
Pro minus con: attribute 13	2.589***	2.410***	2.426***
	(0.011)	(0.010)	(0.042)
Pro minus con: attribute 14	1.948***	1.825***	1.855***
	(0.014)	(0.011)	(0.036)
Pro minus con: attribute 15	1.871***	1.779***	1.799***
	(0.013)	(0.012)	(0.030)
Pro minus con: attribute 16	1.382***	1.264***	1.318***
	(0.014)	(0.007)	(0.021)
Pro minus con: attribute 17	1.610***	1.577***	1.725***
	(0.008)	(0.006)	(0.019)
Pro minus con: attribute 18	1.380***	1.312***	1.358***
	(0.007)	(0.006)	(0.017)
Pro minus con: attribute 19	1.193***	1.069***	1.136***
	(0.011)	(0.011)	(0.032)
Pro minus con: attribute 20	1.317***	1.142***	1.058***
	(0.011)	(0.007)	(0.055)
Year-month FE	✓	✓	
Employer FE		✓	
Former x latter industry, year-month FE			✓
N	6809127	6606642	439530
Adjusted R^2	0.42	0.47	0.43

Notes: The table above relates the overall star rating for an employer review to the "pro minus con" sentiment charge of the 20 attributes detailed in the review's text. Sample in first-difference specification restricted to volunteers who switch employers and do not submit the reviews in the same year-week. Standard errors clustered at the employer (former x latter industry) level for the level (first difference) specifications. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. For a summary of the dataset, see Table A-1.

Table A-4: Likelihood of Concealing Identifying Information by Employer Rating

	$\mathbb{1}$\{Conceal job title\}		$\mathbb{1}$\{Conceal location\}	
Star rating	-0.020***	-0.019***	-0.019***	-0.019***
	(0.001)	(0.001)	(0.001)	(0.001)
Current employee	0.028***	0.030***	0.026***	0.027***
	(0.001)	(0.001)	(0.001)	(0.001)
——— x $\mathbb{1}$\{1–2 star rating\}		0.010***		0.028***
		(0.003)		(0.003)
Log(employment)	-0.016***	-0.016***	-0.008***	-0.008***
	(0.000)	(0.000)	(0.000)	(0.000)
——— x $\mathbb{1}$\{1–2 star rating\}		-0.011***		-0.011***
		(0.000)		(0.000)
Sample mean	0.397	0.397	0.339	0.339
N	989996	989996	989996	989996
Adjusted R^2	0.27	0.27	0.23	0.24

Notes: When submitting a review, users are asked to provide their job title and location, both of which are optional. If these fields are not completed but the review is still submitted, then they are left "blank." Additionally, respondents can choose to leave a job title but anonymize their position. A job title is classified as anonymized if either the word "anonymous" is included in the job title or the job title ends with "employee." A concealed job title is one that is left blank or is anonymized. A concealed location is one that is left blank. Overall star rating, dummy for is current employee, and the logarithm of employment are demeaned by their sample averages. Each specification includes volunteer, industry and year-month fixed effects. Standard errors are clustered at the employer level. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. For a summary of the dataset, see Table A-1.

A.4 Additional figures

Figure A-1: Example Reviews of the University of Minnesota

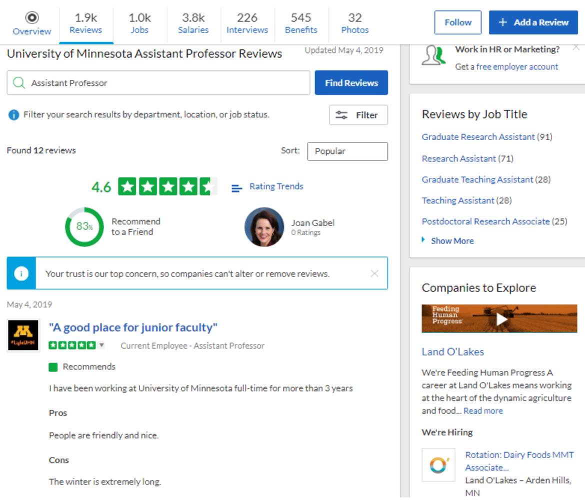

(a) Asst. prof. 5-star w/ header, search & filter

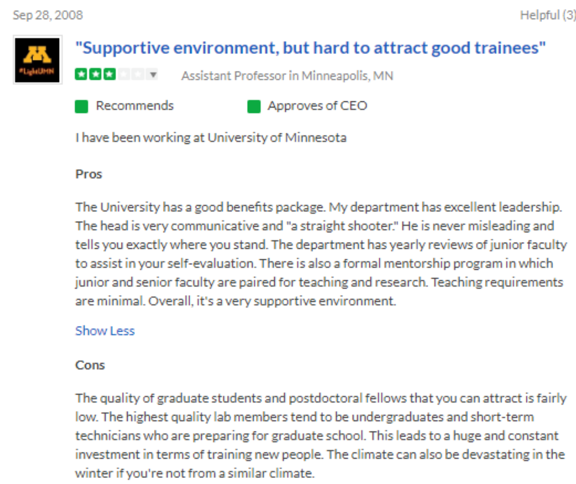

(b) Asst. prof. 3-star

(c) Admin. asst. 4-star

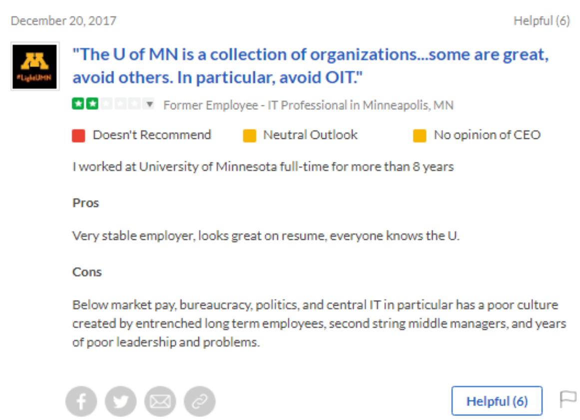

(d) IT 2-star

Notes: Figures (a)-(d) are screenshot images of example reviews that were submitted voluntarily and anonymously on the Glassdoor website for the University of Minnesota. In addition to displaying an employee's review, Figure (a) includes the top-line information and search capabilities available to users when reading through an employer's collection of reviews.

Figure A-2: Share of Total Votes Unhelpful over Time

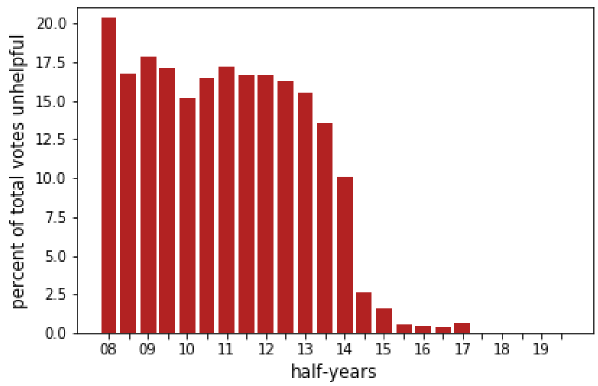

Notes: When reading a user-written review of an employer on the Glassdoor website, the reader can signal to others that the review supplied helpful information by submitting a “helpful” vote for the review. Conversely, if the review supplied information that was not informative to her, she can signal the review was not helpful by submitting an “unhelpful” vote. This figure calculates the fraction of submitted votes that were unhelpful votes, partitioning reviews into half-years bins based on the date the review was submitted.

Figure A-3: Distribution of Review Text across Topics

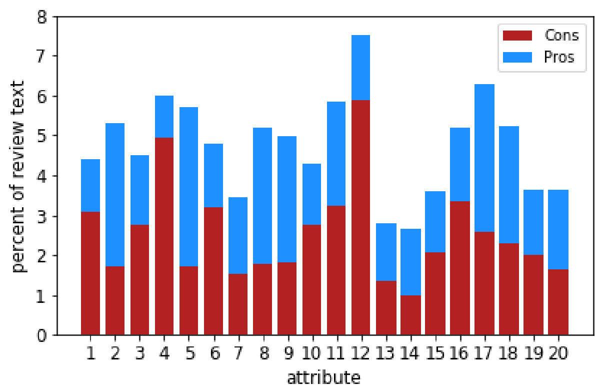

Notes: Attribute distributions are assigned separately to the “Pros” and “Cons” sections of each review, using a single LDA algorithm that captures all review text. To obtain the distribution across sections and attributes, the pro (con) attribute distribution is weighted by the pro (con) section’s character share of the review text. The bars displayed sum to 100 percent of the review text, with each attributes’ bar partitioned into the share attributable to the “Pros” section (upper-portion) and “Cons” section (lower-portion).

Figure A-4: Highest Incidence Words within Each Attribute

(a) Attribute 1

(b) Attribute 2

(c) Attribute 3

(d) Attribute 4

(e) Attribute 5

(f) Attribute 6

(g) Attribute 7

(h) Attribute 8

(i) Attribute 9

(j) Attribute 10

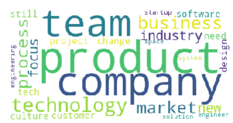

(k) Attribute 11

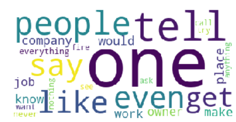

(l) Attribute 12

(m) Attribute 13

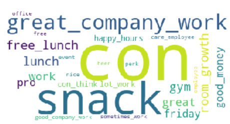

(n) Attribute 14

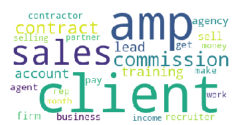

(o) Attribute 15

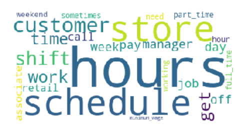

(p) Attribute 16

(q) Attribute 17

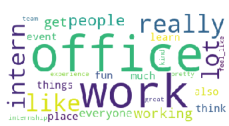

(r) Attribute 18

(s) Attribute 19

(t) Attribute 20

Notes: Distribution of review text into 20 attributes assigned using an unsupervised Latent Dirichlet Allocation (LDA) model. The LDA model is calibrated on the full text of reviews through August 2019, which reflects a sample of 6,809,217 reviews. Attributes are ordered in descending rank according to their relative helpfulness (i.e., the coefficients in Table 4: Column 2). Each panel displays the 25 highest-incidence words within an attribute.

Figure A-5: Relationships Between Probability of Identity-Aspect Concealment and Leaving a More-Negative Review by Measures of Retaliation Risk: Sample of Only Negative Reviews

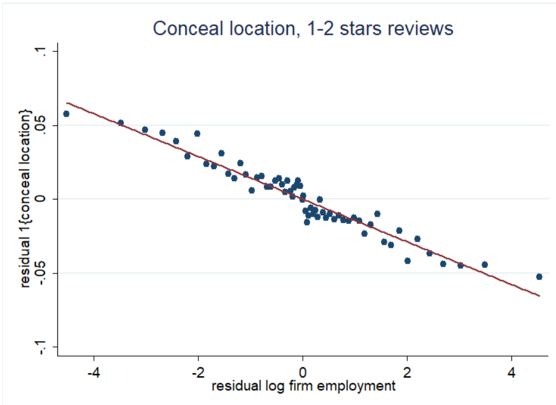

(a) Location by firm employment

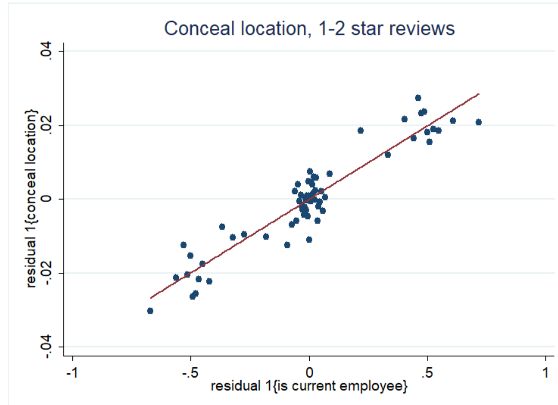

(b) Location by current employee

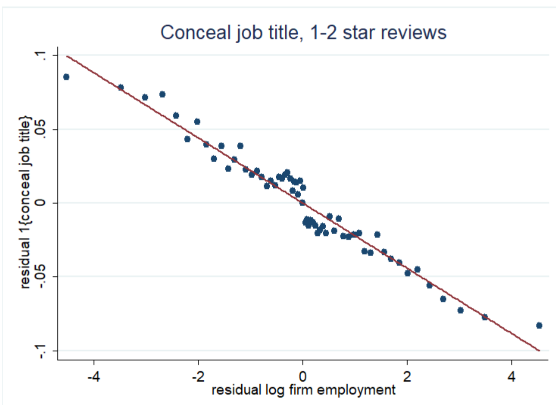

(c) Job title by firm employment

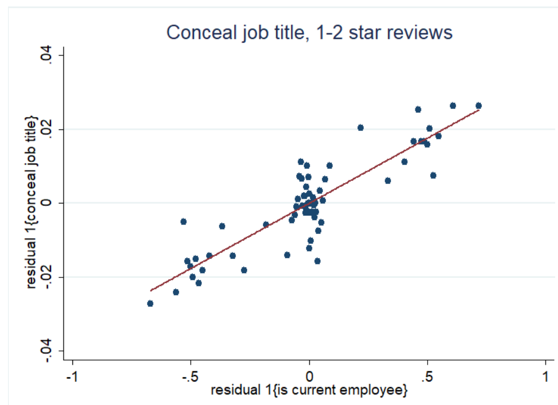

(d) Job title by current employee

Notes: The figures above detail the residual rate at which volunteers conceal potentially identifying information depending upon the size of the employer and whether the employee is still working for the firm. Sample of volunteers is restricted to those who leave multiple reviews on the website, and sample of reviews restricted to only one- and two-star reviews. Residuals come from a regression that includes the overall review rating; fixed effects for the volunteer, year-month and industry; and log firm employment (panels a & c) or a dummy for is current employee (panels b & d). Concealing job title reflects reviews that leave the job title blank or anonymize the job title. Concealing location reflects reviews that leave the location blank. Solid red lines reflects a linear line of best fit.