

MIT SLOAN SCHOOL OF MANAGEMENT
WORKING PAPER 6607-22

Workplace Happiness and Job Search Behavior: Evidence From A Field Experiment

George Ward
MIT Sloan, Institute for Work and Employment Research



This work is licensed under a Creative Commons Attribution-NonCommercial License (US/v4.0)
<http://creativecommons.org/licenses/bync/4.0/>Last Revised: October 2020

WORKPLACE HAPPINESS AND JOB SEARCH BEHAVIOR: EVIDENCE FROM A FIELD EXPERIMENT

George Ward*

Massachusetts Institute of Technology

January 26, 2022

Abstract

For over a century, organizational researchers and practitioners alike have been interested in the role played by employee happiness in shaping workplace performance. Whereas prior work has focused on the potential effects of happiness on productivity and retention, in this paper I study the effects of workplace happiness on a firm's ability to compete in the labor market and attract workers. I provide evidence from a field experiment on a large online jobs platform in the USA, in which treated job seekers were shown crowd-sourced information about the happiness of incumbent workers at the companies to which they were considering applying. I provide evidence that showing information about happiness has an impact on job search behavior, with treated job seekers increasing their application selectivity and redirecting applications from low-happiness companies. These labor supply effects are driven largely by job seekers "screening out" unhappy firms from their job search, a finding that is replicated in subsequent field experiments on the platform in Canada and the UK. The findings suggest that emerging large-scale sources of crowdsourced data affect behavior in the labor market and, ultimately, that employers face incentives to invest in organizational and management practices that are conducive to worker happiness.

*Correspondence: wardg@mit.edu. MIT Sloan, 100 Main Street, Cambridge MA 02139. I thank Moshe Barach, Clément Bellet, Alan Benson, Ambra Cantarutti, Andy Chiang, Jared Curhan, Dean Eckles, Roberto Fernandez, Dave Holtz, Jan De Neve, Apostolos Filippas, Andrey Fradkin, Laura Gee, John Horton, Simon Jäger, Will Johnson, Tom Kochan, Erin Kelly, Robert Metcalfe, Mike Norton, Paul Osterman, Andrew Oswald, Taylor Sansom, Aaron Sojourner, Janeane Tolomeo, and Ezra Zuckerman for help, feedback, and suggestions. For pre-registration details, see AEA RCT Registry AEARCTR-0007170 and AEARCTR-0007359. I am grateful to *Indeed* for their collaboration and generous data sharing.

1 Introduction

The role played by employee happiness in shaping workplace outcomes has long been a topic of intense interest to organizational researchers as well as practitioners. In a large body of literature that stretches back many decades (see, e.g., Fisher and Hanna, 1931; Hersey, 1932a,b; Houser, 1927; Kornhauser and Sharp, 1932), one of the most extensively studied aspects of workplace happiness has been its effects on firm financial performance. According to this so-called “holy grail” of organizational research (Landy, 1985; Wright and Cropanzano, 2004), happier workers will increase an organization’s economic performance since they are expected to (i) perform better in their jobs and (ii) be less likely to leave the firm. In this paper, I study a third possible reason: the increased ability of companies with a happier workforce to compete in the labor market and attract workers in the first place.

The past few years have seen a resurgence of interest in workplace happiness, and a growing number of companies are at least claiming to care about the happiness of their employees. For example, in a recent large-scale survey of 1,073 US executives, around 87% agreed that workplace happiness can provide their firm with a competitive advantage (HBR Analytical Services, 2020).¹ Nevertheless, work is still far from a happy experience for most people,² and despite the apparent widespread agreement among managers about the benefits of employee happiness, the large majority remain hesitant to do anything to actively improve it. Indeed, only a third of organizations see employee happiness as a strategic priority, and fewer than 20% actually have any sort of employee well-being strategy in place.³

However, for the firms who do exert effort to foster employee happiness, is there a business case for doing so? This depends, at least partly, on the extent to which prospective employees value happiness in the workplace. There is a long-running discussion on human motivation, centered around the key question: What is it that people want? A fundamental feature of many theories of motivation is that humans want to be happy (Lawler, 1973; Myers et al., 1993). But testing this empirically is difficult. Although it has long been recognized that workers want more out of a job than just a paycheck (see, e.g., Rosen, 1986) and would be willing to trade off wages for a wide range of specific job amenities such as flexibility, team work, autonomy, and purpose (e.g., Burbano, 2016; Jencks et al., 1988; Maestas et al., 2018; Mas and Pallais, 2017; Stern, 2004), the extent to which they value the more general concept of happiness in the workplace is an open question. In fact, happiness might be seen as a frivolous distraction from the things that really matter to people – like wages. Indeed, the fact that so few companies report having any sort of strategy to look after or improve employee happiness suggests that this is likely the actual view

¹Moreover, among the factors that these executives say they see as the most important to the success of their organization, happiness of the workforce comes second only to the happiness of customers—and above a range of other factors such as continuous innovation, investing in the best technologies, and having the right market strategy.

²Time-use data show that paid work is ranked lower in terms of happiness than almost any other activity that people engage in (the only thing worse is being sick in bed) and that the most unhappy moments of people’s day are when they are with their boss or supervisor (Bryson and MacKerron, 2016; Krueger et al., 2009).

³For more details of the survey of managers, which included responses from a wide range of sectors and organization sizes, but was not a randomly drawn sample, see HBR Analytical Services (2020).

of the majority of firms.

In this paper, I present evidence from a field experiment that provides job seekers with information about the aggregate happiness of incumbent workers at companies to which they are considering applying. The experiment took place on a major online jobs platform, *Indeed*, where job seekers can already see a range of crowdsourced information about companies—including extensive details of salaries, employee reviews, and star ratings left by workers. I explore the effects of providing additional information to job seekers about company-level happiness, in an experimental way. By assessing labor supply responses to information about the comparative happiness of firms, while keeping information about salaries constant, I am able to provide a test of the extent to which people are motivated to work at happier firms.

While much of the early literature on employee happiness operationalized the concept using measures of job satisfaction (Wright and Bonett, 2007), a more recent turn has seen a greater focus on hedonic (or affective) measures (Barsade and Gibson, 2007; Brief and Weiss, 2002). Here, I follow this line of work by studying the behavioral effects of showing job seekers aggregated company-level information on the extent to which employees agree with the statement “*I feel happy at work most of the time.*” *Indeed* has amassed a crowdsourced survey dataset of answers to this question from over 5.5 million workers across the USA, making it, to the best of my knowledge, the largest ever source of data on employee happiness. For the 20,000 or so companies that had individual-level survey responses from 20 or more workers during the period of the experiment, a “work happiness score” was calculated. To help motivate the field experiment, I first show that this measure of happiness varies significantly across companies, even within tightly-defined industries (see Figure 2). This begs the question of what effects they may be of this variation for firms. While long line of work has investigated the downstream effect of employee happiness on productivity and retention (see Walsh et al., 2018), the extent to which it may have implications on a firm’s ability to compete in the labor market is not yet clear.

I experimentally display this happiness information prominently to job seekers on the company profile pages of the platform, over a 10-month period. My main finding is that job seekers respond behaviorally to additional information on workplace happiness. Displaying happiness scores to job seekers serves to redirect applications away from low happiness companies to happier ones. I document an asymmetry in this effect for low and high worker happiness firms, whereby information revealing low happiness discourages applications more strongly than equivalent information revealing high happiness encourages them. This redirection effect is thus driven largely by jobseekers of the platform screening out low happiness organizations from their job search. For companies with scores below 60 on the 100-point happiness score scale ($mean = 63.8; SD = 8.8$), I find an estimated treatment effect of around -2.75% . For scores between 60 and 80, there is little discernible effect on applications, and for companies with scores over 80, I find an estimated treatment effect of 2% . In addition to studying effects on application behavior, I am also able to track actual job hires for a subset of the experimental sample. Although additional information on workplace happiness seems to make job seekers more cautious in their applications by screening

out low happiness companies, I show that this does not affect the probability that they ultimately get a job.

I distinguish between *information-provision effects* and *score-value effects*. The field experiment allows me to first estimate information-provision effects on labor supply by varying which job seekers see information about a company’s happiness. Here, I study what List (2007) refers to as a natural field experiment—where subjects behave naturally in the environment in which they are being studied, without knowing they are taking part in an experiment. This has the benefit of combining the advantages of randomization that come from laboratory or framed field experiments with the realism that comes from studying observational field data.⁴ One downside to this approach, however, is that although I am able to vary the provision of information, it is not feasible (without deception) to induce any experimental variation in the value of the score itself—which would allow me to estimate the effect of changes in the score on application behavior.⁵

To isolate score-value effects on labor supply, I turn to observational data from treated job seekers and rely on a variety of quasi-experimental empirical strategies. I make use of the fact that (i) the experiment takes place over a 10-month period and (ii) treated job seekers view multiple companies with different scores, to estimate application equations that include company, time, and job seeker fixed effects. First, in order to identify the effect of the score on applications, I leverage this within-company and within-jobseeker variation in the scores displayed. Going from being defined on the platform as having “Low” happiness (scores of 49 and below) to “Average” (scores of 60 to 69), the probability of a viewing job seeker applying goes up by 13.93%. Again, however, I find an asymmetry in this association, with there being little discernible relationship between the score and the probability of applying above scores of 65. Second, I build on this evidence by exploiting discrete jumps in the score caused by rounding rules (cf. Luca, 2016; Sockin and Sojourner, 2020), and, in doing so, present evidence to suggest that this is a causal relationship.

In two pre-registered replication experiments on the platform, I find a similar pattern of results in Canada and the UK. In these subsequent studies, treated job seekers responded to the provision of company-level happiness information and screened out low happiness firms from their job search. Finally, I also report follow-up analyses in the form of a pre-registered stated-preference vignette study that I embedded within nationally representative surveys in the USA, UK, and Canada, which allows me to investigate the generalizability of the paper’s main findings among the labor force in general. Here, I vary both happiness and wage levels at hypothetical companies and elicit respondents’ willingness to apply to work at them. Consistent with the field findings, I find evidence of a preference for working at happier workplaces—and, importantly, a willingness to trade off income to secure jobs at such companies. On average, respondents in the USA would be willing to take a 10.6% pay cut to take a job at a company with a happiness score

⁴In my setting, real job seekers view jobs at real companies in the US labor market and are provided with truthful information about the happiness of those companies. Moreover, the experiment involves a behavioral and potentially highly consequential outcome: the decision whether or not to apply to a company.

⁵In an ideal laboratory experiment, one would likely vary access to information about the happiness of incumbent workers, but also randomly vary the score itself (and likely also the pay level) across companies, such that treated and control job seekers are looking at the same company but with different levels of happiness.

one standard deviation above the mean, compared with working at a company at the mean level of happiness.

The remainder of the paper is structured as follows. Section 2 surveys some of the existing literature and discusses theoretical links between information about workplace happiness and job seeker behavior. In Section 3, I introduce the institutional setting by describing the survey instrument used to collect data on individual-level happiness and by explaining how company happiness scores are calculated. In Section 4, I report results from the natural field experiment that randomly allocates happiness information to job seekers, before focusing in Section 5 on the effect of the score itself on applications among treated job seekers. Section 6 provides additional estimates from a nationally representative survey. Section 7 discusses managerial implications, a number of key limitations, and some possible extensions. Section 8 concludes.

2 Background

2.1 Subjective wellbeing in the workplace

The philosophical study of human happiness is a long-running concern, but the empirical investigation of what is now more typically termed subjective well-being (SWB) has increased dramatically over the past three decades (Clark, 2018; Diener et al., 2017). SWB is often referred to loosely and collectively as happiness, but is actually typically thought of as having two separate components that measure (i) how people think about the state of their lives and jobs and (ii) how people feel moment to moment (Krueger and Stone, 2014). The former is usually thought of as evaluative SWB and is measured via concepts such as life or job satisfaction, while the latter is often referred to as hedonic or affective SWB. A third dimension is increasingly also recognized and refers to eudaimonic well-being or purpose (see, e.g., Cassar and Meier, 2018; Chadi et al., 2017; Gartenberg et al., 2019).

SWB varies significantly across people and across organizations, even within industries and locations – that is, across otherwise observationally similar firms facing the same business environment. On the one hand, a large literature has sought to understand the underlying causes of this variation, for example by studying the effects of different management practices as well as various aspects of organizational cultures (e.g., Bloom et al., 2014; Breza et al., 2018; Card et al., 2012; Clark, 2010; Gosnell et al., 2020; Jencks et al., 1988; Moen et al., 2016). On the other hand, work has sought to understand the downstream effects of this variation. At the level of the firm, happier workplaces have been shown to have higher financial performance (e.g. Edmans, 2011, 2012). Trying to understand the micro-foundations of this relationship, a large stream of research has focused on the so-called “happy-productive worker thesis” (Tenney et al., 2016), both in terms of job satisfaction (e.g., Iaffaldano and Muchinsky, 1985; Judge et al., 2001) as well as affect (e.g., Amabile et al., 2005; Bellet et al., 2021; Estrada et al., 1997; Oswald et al., 2015; Rothbard and Wilk, 2011; Staw and Barsade, 1993). A smaller body of research has also employed panel data to show a predictive link between happiness and subsequent employee quits (Clark, 2001; Green,

2010; Levy-Garboua et al., 2007). However, although the link between workplace happiness and recruitment has occasionally be discussed in theoretical terms, the extent to which happier workplaces are better able to compete in the labor market and recruit talent is an open empirical question.

2.2 Information problems in the labor market

Like almost all markets, the labor market is one in which participants face (often large) information problems (e.g. Carmichael, 1984; Jäger et al., 2021; Sockin and Sojourner, 2020). Although a large body of work has studied the employer’s information problem, such as the extent to which they can observe the ability of potential workers and deal with that problem by using performance-contingent contracts, referrals, and so on (see, e.g., Lazear and Shaw, 2007), a great deal less is generally known about the job seeker’s information problem. A recent series of papers has begun to address this by experimentally providing job seekers with information about the labor market in general and aspects of the search process.⁶ However, much less attention has been afforded to the job seekers information problem in terms of knowing what working is like at different companies.

Carmichael (1984) points toward employer reputation as a means through which job seekers may be able to deal with this information problem, at least partially. However, information on reputation was, for a long time, mostly limited to companies at the tails of the reputation distribution. Brown and Matsa (2016) show, for example, significant hiring effects for firms that are known publicly to be in severe financial distress. But while information on a wide cross-section of firms was long unavailable, recent digitization of labor market institutions means that detailed information on a broad distribution of organizations is now more freely and easily available to job seekers.

2.3 Digitization of labor market institutions

Growing digitization of the labor market has the potential to significantly alter the ways in which people look for and choose between employers and jobs. Platforms such as *Indeed*, *Glassdoor*, and *LinkedIn* have rapidly become the dominant way that people look for work, at least in high-income countries such as the USA (see, e.g., Horton and Tambe, 2015). A nascent body of work has begun to use data from digital platforms to understand a range of issues, such as how job seekers behave over the course of a period of unemployment (Faberman and Kudlyak, 2019; Marinescu, 2017) or respond to unemployment insurance (Baker and Fradkin, 2017; Marinescu and Skandalis, 2021).

⁶For example, Altmann et al. (2018) run a field experiment that gives unemployed job seekers a brochure including information on job search strategies (see also Belot et al., 2019). Gee (2019) reports on a natural field experiment on *LinkedIn* that informs jobseekers how many applications a job already has, and Bhole et al. (2021) show using a field experiment on *Jobs on Facebook* that providing information about competition for a vacancy on the platform serves to redirect applications to jobs with fewer prior applications. Coffman et al. (2017) also show in a field experiment that information on job take-up rates from prior years of the *Teach For America* program can have large effects on which jobs people choose.

But in addition hosting job adverts and helping to facilitate applications, such platforms also serve another—often overlooked—key function: to crowd-source and display information about companies.

In a very small literature beginning to try to better understand the information-aggregation function of digital labor market platforms, Sockin and Sojourner (2020) use data from *Glassdoor* to better understand the reasons why people give reviews of companies and, in particular, why they might conceal parts of their identity in situations where retaliation is more likely (see also Sockin et al., 2021). In addition, the authors exploit rounding rules in the overall star rating of companies on the platform, to show effects on labor supply, at least for small companies. In addition to this, Benson et al. (2020) also use a field experiment on *Amazon Mechanical Turk* and find effects of employer reputation on labor supply in the online market for gig work – and in doing so also develop of formal model to demonstrate the important point that online employer reputation systems have the potential to discipline firms who mistreat workers in different ways. That is, to the extent that attracting a larger application pool is beneficial to the firm,⁷ there is an incentive to ensure a good reputation as an employer.

2.4 Do workers care about happiness?

Although the question of what workers care about in a job has been studied for many decades, the literature has been re-enlivened in recent years with the creative use of surveys and experiments. For example, Maestas et al. (2018) pose hypothetical job choices to survey respondents while experimentally varying the wage and different bundles of job characteristics (see also Wiswall and Zafar, 2018). Mas and Pallais (2017) embed a survey experiment that elicits preferences for flexibility within a recruitment drive for a call center, and Flory et al. (2015) run a natural field experiment where interested job seekers on an internet jobs board were shown the same job with differing compensation regimes.⁸

In addition to this research on specific job amenities, other studies has also studied preferences for the characteristics of the organizations as a whole – such as a firm’s commitment to corporate social responsibility (CSR) (Carnahan et al., 2017; List and Momeni, 2021) or pro-social work (Ashraf et al., 2020). Burbano (2016) shows, for example, that providing positive information about company-level CSR on job adverts attracts applications from job seekers. When studying the mechanisms behind the link between CSR and application behavior more closely, Burbano (2016) finds evidence that people respond to such information because they interpret it as a signal that the company is likely to treat workers well, thus raising the expected utility of working there. In line with this, happiness is typically thought of as an overall measure that is a function

⁷A larger applicant pool is usually a positive outcome for an organization given that it typically allows firms to choose more a larger pool and in doing so be able to pick more talented and productive workers, as well as workers with a broader array of diversity along various demographic characteristics.

⁸In related work, Bandiera et al. (2010) find using a field experiment that workers are motivated to work together with friends (see also Nagaraj and Piezunka, 2020), while Benson et al. (2017) and Breza et al. (2018) show that workers have a strong preference for fair treatment.

of various aspects of organizational life include pay, flexibility, interpersonal relationships, and so on. In this paper, I thus take a more general approach than the literature studying specific amenities, and instead estimate the effects of providing information more directly on the experienced utility (cf Kahneman et al., 1997) of incumbent workers at different firms.

2.5 Social Information and Economic Behavior

An established stream of research has shown that word-of-mouth and referrals within social networks can affect job search behavior (e.g. Dustmann et al., 2016; Fernandez et al., 2000). However, the extent to which this translates to aggregate information about the experiences of strangers is far from clear. Indeed, even if job seekers are motivated to pursue happiness in their search, they may see crowdsourced information on digital platforms as only a very weak signal of their own expected happiness at a particular company. This is particularly the case to the extent that happiness in a job is the result of better matching between workers and employers, rather than a general quality of the employment relationship for all employees at a company.⁹ Moreover, platforms such as *Indeed* already show large amounts of information about salaries, reviews, star ratings, and so on. Job seekers may thus see little *additional* signal to happiness information.

A series of studies has shown the impact of social information on economic behavior in general. Much of this work has probed relatively low-stakes decisions such as contributing to a movie ratings website (Chen et al., 2010), the market for wedding services (Tucker and Zhang, 2011), and donation behavior (Frey and Meier, 2004; Gee and Schreck, 2018). A rapidly growing body of work is studying the effects of crowdsourced ratings on digital platforms (Dellarocas, 2003). While much of this research has up to now been focused almost entirely on product markets, using data from platforms such as *Yelp*, *TripAdvisor*, and *eBay* (for prominent examples, see, e.g., Chevalier and Mayzlin, 2006; Helmers et al., 2019; Luca, 2016; Reimers and Waldfogel, 2021), the extent to which this translates to the labor market is much less clear. Jobs are, after all, different to consumer goods. Decisions made in the labor market of where to work tend to have larger and much longer-lasting impacts than decisions about which products to buy or which restaurant to eat in—both because people spend a large percentage of their waking hours working at the company they choose and because the commitment is a longer-lasting one.

3 Empirical Setting

Indeed is a large online jobs platform, with over 250 million unique visitors worldwide each month.¹⁰ The website hosts job adverts, which job seekers can search and browse. In addition,

⁹If happiness were a general quality, then any representative worker’s experience would give useful information since everyone’s criteria for assessing the company would be the same. However, if everyone has different preferences for different types of organizational and managerial practices, then I would be much more likely to find the recommendation of someone in my social network to be information given that she is more likely to know my preferences (or have similar preferences to me, which is typically true within a social network).

¹⁰This figure is based on unique visitors calculated by Google Analytics, February 2020.

each company that is listed on the platform also has its own set of “company pages.” Here, the platform displays a wide range of information about each company, much of which is crowd-sourced from other jobseekers of the website. This includes job seeker-written text reviews, a headline star rating based on the question “Overall, how would you rate this company?”, questions about the interview process at the company, a large amount of salary information collected from current and former employees, and a more flexible question-and-answer section. In this paper, I study the effects of adding information on workplace happiness to this large existing corpus of information displayed to job seekers about companies.

3.1 Crowdsourced Happiness Survey

Beginning in October 2019, visitors to the website were invited as part of the platform’s general data collection efforts to answer survey questions on workplace happiness at the firms they were currently working at (or had previously been employed by). The data collection process is ongoing, and in this paper, I use data up to March 17, 2021, giving me a total of 5,338,631 individual survey responses in the USA.

The survey question was based on standard academic definitions of happiness. The headline question of the survey asks respondents the extent to which they agree, on a 1 to 5 (“strongly disagree” to “strongly agree”) scale, with the statement “*I feel happy at work most of the time.*” This is a hedonic well-being question in that it asks about feelings of happiness. However, it also a global judgment in that it asks not about a specific time frame but about workers’ experience at the company as a whole. The validity and reliability of measures such as this have been the subject of decades of academic research (for a full discussion of the issues surrounding the measurement of SWB, as well as a detailed overview of the ways in which the validity and reliability of such measures have been tested, see Krueger and Schkade (2008); Krueger and Stone (2014)).

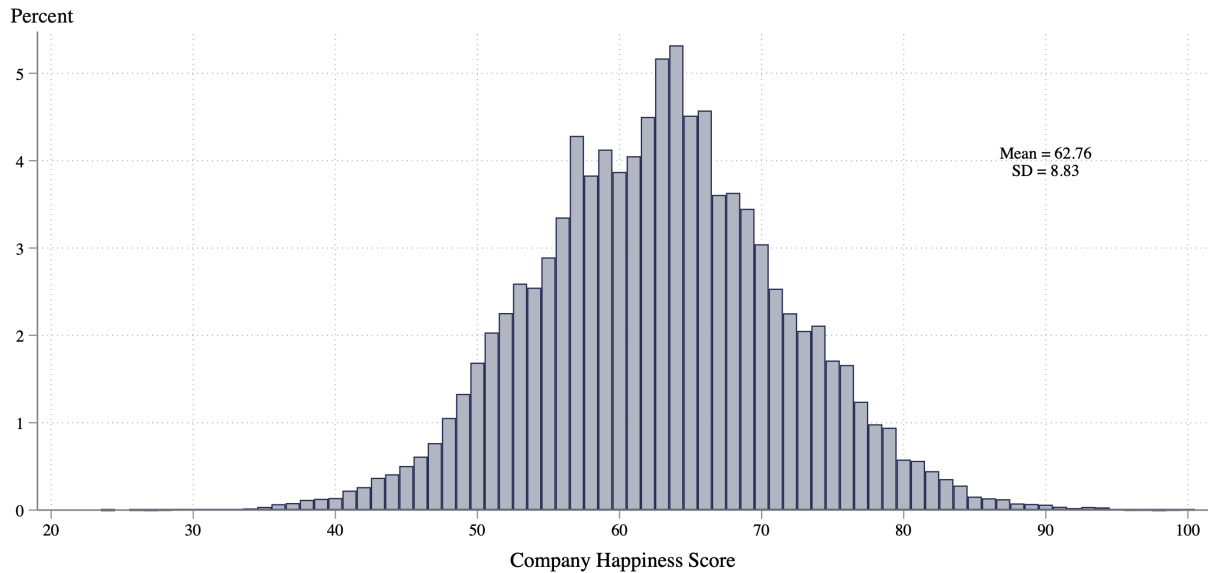
Respondents were assured their responses would be anonymous and told that their honest responses would help other job seekers. A further 12 questions follow this main question and ask the respondent about different sub-dimensions of workplace well-being. The exact question wording is provided in an Online Appendix, but the sub-dimensions include: achievement, appreciation, belonging, energy, flexibility, inclusivity, learning, management quality, fair pay, purpose, support, and trust (see Appendix B).¹¹

3.2 Company-Level Happiness Scores

Once a company has 20 or more completed surveys, it is eligible to have a workplace happiness score shown on its company pages. Company-level happiness scores are calculated to be on a scale out of 100. In practice, this entails taking the mean of happiness responses on the 1–5

¹¹More recent data collection has included questions on job satisfaction and stress. This gives a fuller picture of SWB that is in line with academic understanding of the concept, which, as noted above, includes evaluative, hedonic, and eudaimonic dimensions (Diener et al., 1999). Data on stress and satisfaction were not shown to any job seekers during the course of the experimental period I study in this paper.

Figure 1: Company-Level Work Happiness Scores



Note: Histogram shows the distribution of the work happiness scores that are displayed to job seekers during the experiment. Bin width=1.

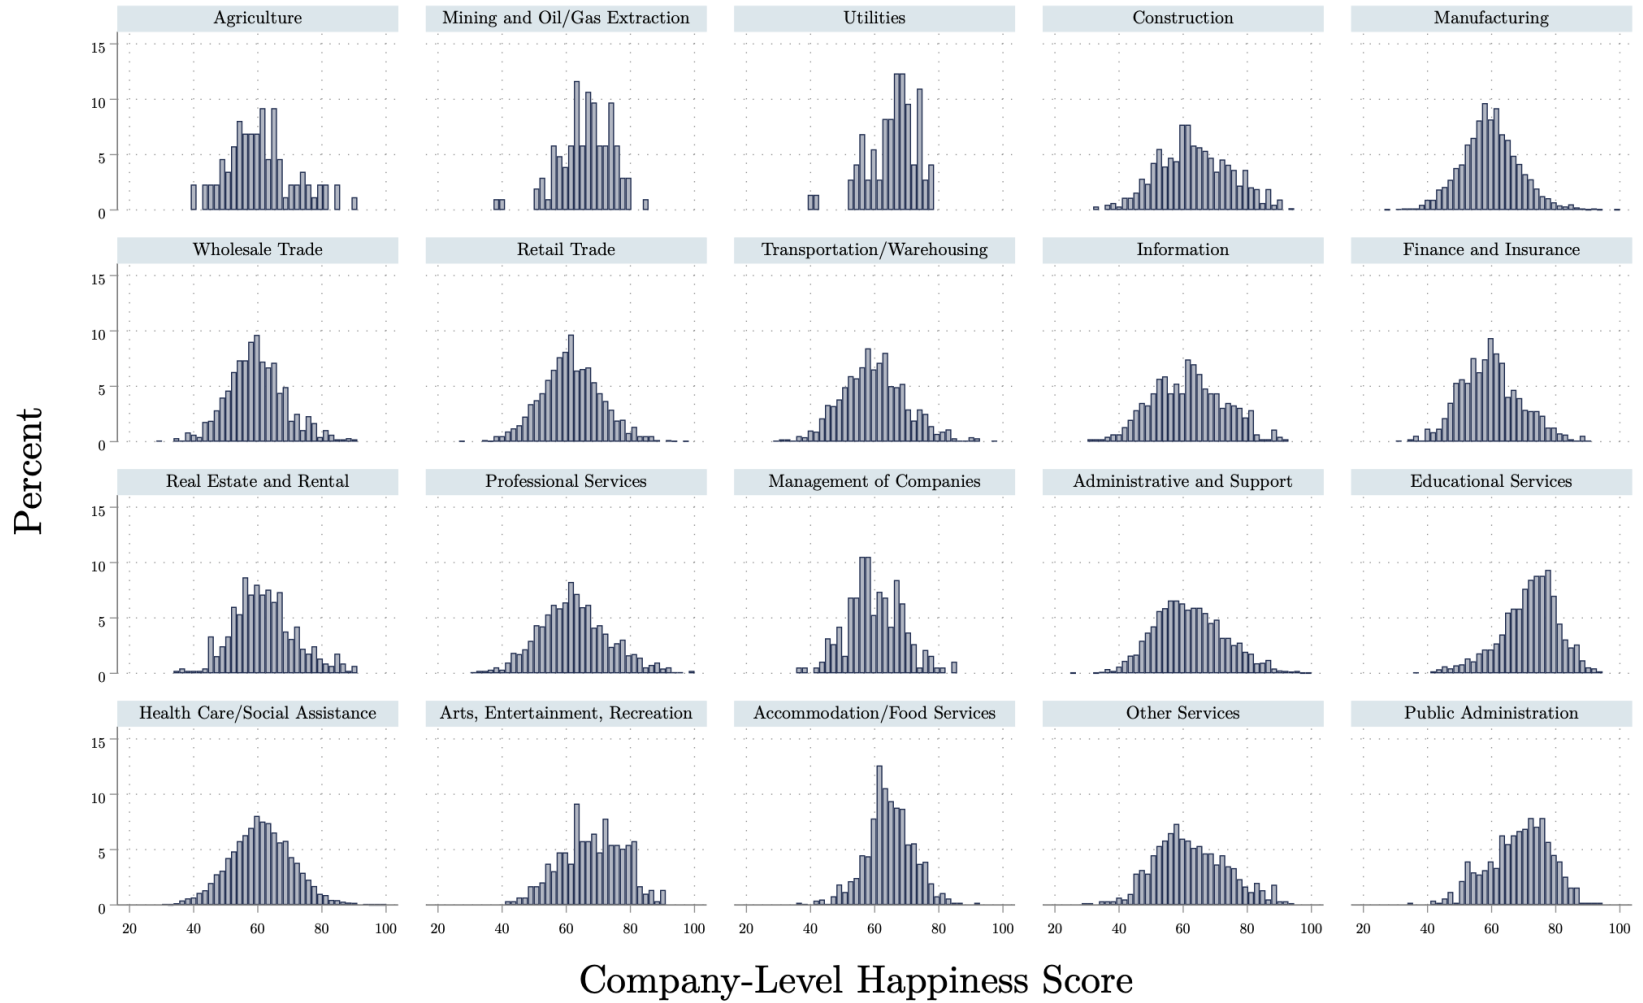
agreement scale and multiplying by 20—giving a score that is shown to job seekers as an integer between 20 and 100. Between October 2019 and April 2020, individual-level happiness surveys were conducted, but no company-level happiness scores were yet shown on the website. At the end of April, the company-level scores began to be experimentally shown to job seekers on the site.

Figure 1 plots the distribution of company-level happiness scores. Scores are centered around a score of 63. Very few companies have a score below 40 or above 90, and only a small number have a score below 50 or above 80. In Figure 2 I show the distribution of scores within 2-digit industry codes, and find that there is substantial variation across organizations. Further, in Figure S1 I show substantial variation in company-level happiness, even after partialling out fixed effects for 4-digit industry codes, suggesting that organizations vary in the happiness of their workforce, even with tightly-defined industries.

Clicking on a company’s page on the website, the job seeker arrives at the company’s landing page. There are a number of tabs, which can take the job seeker to other areas of the company’s pages (but always with a banner of tabs at the top such that the job seeker can navigate back to any of the company’s pages). The main landing page includes basic information about the company, such as industry, company size, year founded, and CEO name. Other pages the job seeker can navigate to include reviews written by job seekers, crowdsourced salary information, and job listings.

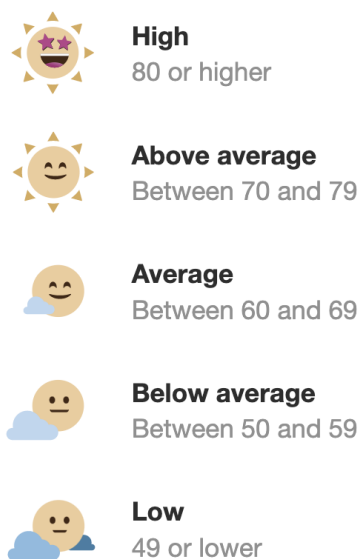
The happiness score is shown prominently as the first item on the company’s main landing

Figure 2: Company-Level Happiness by Industry in the USA



Note: Plotted are the distributions of the company-level happiness score on the final day of the experiment, within 2-digit NAICS industries.

Figure 3: Emojis Displayed to Job Seekers Alongside The Happiness Score



Note: Screenshot of the accompanying emojis and distributional “signposting” statements. When the happiness scores are displayed to job seekers on the site, they are accompanied by an emoji and a statement to help people situate the score they are seeing within the overall distribution.

page (for more details, see Section 4.1). Given that the score does not have any natural units, whenever the score is posted on the website, it is accompanied by an emoji and a piece of text explaining that the score is high, above average, average, below average, or low (see Figure 3). This information is based on cutoffs that were fixed throughout the course of the experiment and were the same for all companies.

Although happiness is the headline piece of information, a further 12 questions were also asked on sub-dimensions of workplace well-being. The happiness score and its emoji are always shown first and most prominently. Two sub-dimensions are shown alongside the happiness score. These are the two sub-dimensions with the highest two scores for that company. The job seekers can choose to expand to a “full report” to see the scores of the remaining sub-dimensions if they wish to.

4 Field Experiment: Information-Provision Effects

4.1 Experimental Design

Job seekers are randomized on the basis of their internet cookie to be in either the treatment or the control group. Treated job seekers will see the happiness score for all eligible companies (i.e., companies that have at least 20 responses to the happiness survey, so that a happiness score can be calculated) if they navigate to that company’s page. Control job seekers visiting the same company’s pages will see those pages as normal, but without the happiness information. Figure

4 shows what the two experimental conditions look like, for a treated and control jobseeker who visit the page of an eligible company (see Figure S2 for what the treatment looks like on a tablet computer or smartphone). This is a somewhat conservative test of the the importance of happiness to job seekers. First, I am estimating the experimental effect of showing happiness information, over and above all of the information about workers' subjective experiences at the firm that are already contained in star ratings, reviews, and so on. Second, the experiment is conditional on a job seeker having already navigated to the company's page, and thus does not include any effects of happiness attracting job seekers to view the page in the first place.

The experiment ran for 10 months beginning in May 2020 and was based on a roll-out design. Happiness scores were shown to the majority of job seekers when the score was launched as a product on the platform. A hold-out group comprising a randomly assigned 5% of job seekers, which I refer to as the control group, did not see any happiness scores. I am able to analyze the data at the level of jobseeker-company-day triads. The sample is restricted to jobseeker-company-days on which the company had a happiness score (since treatment and control will look the same if the company does not have the 20 or more happiness surveys required to calculate a score to display). Further, I restrict to jobseeker-company-day where the job seekers navigates to the company's page.¹² For the main analyses of the paper, I further limit the sample to include only the first observation per jobseeker-company pair. For robustness, I also limit the sample to include only the first jobseeker-company-day for each job seeker, given that there may be path dependence in job search strategy arising from the treatment assignment.¹³ The main outcome is whether or not the job seeker applied to the company, from anywhere on the site, during that day. This means that the job seeker clicks on the "Apply Now" button of at least one job listed by the company.¹⁴

4.2 Summary and Balance Statistics

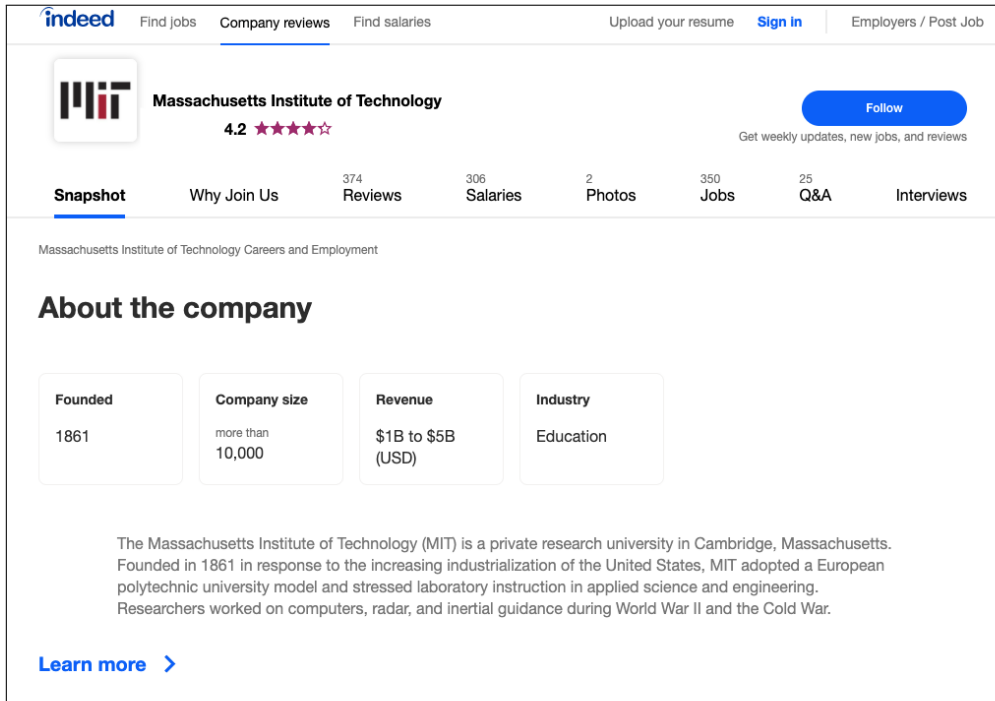
Overall, the sample includes 23,376,519 job seekers and 37,373,151 jobseeker-company-day observations. Only relatively basic information about the characteristics of job seekers is collected routinely by the website—I am not, for example, able to observe characteristics such as age, race, or gender. Nevertheless, I am able to observe the age of the job seeker's cookie, whether or not the job seekers is registered with the website, whether or not the job seekers is viewing on a desktop

¹²The happiness score appears at the top of the page, so that I assume that a job seeker who navigates to the page sees the score (or would have seen it if they were not in the control group).

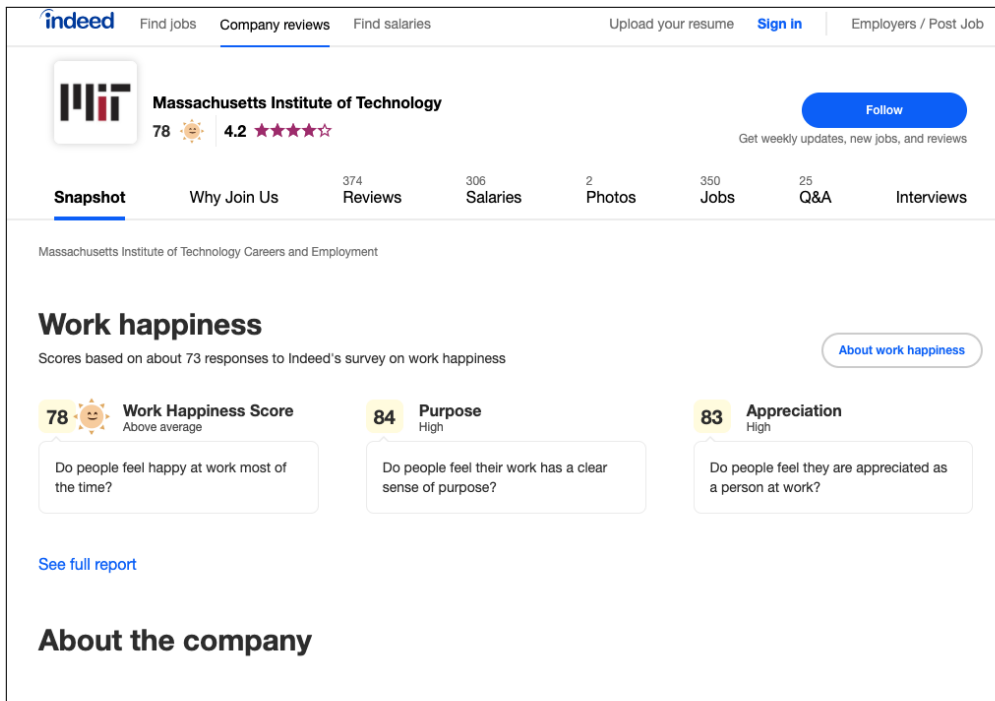
¹³That is, the initial treatment on the first company-day that a job seeker sees with an eligible score could have a knock-on effect on things such as search intensity or which companies and jobs the job seeker subsequently searches for.

¹⁴Note, however, that clicking on the apply button does not necessarily mean that the job seeker goes through with the whole application. In many cases, the apply button will direct the job seeker to a third-party website, usually the company's own website, to make the application. In these instances, I am not able to track anything beyond this initial click. Nevertheless, "application click" remains a good behavioral proxy for applying for a job, and is itself often used as a behavioral check in survey experiments that ask job seekers to state their hypothetical interest in different types of jobs and companies.

Figure 4: Experimental Conditions



(a) Control



(b) Treatment

Note: Screenshots from the desktop version of the platform, showing treatment and control conditions for an eligible organization. Job seekers who are randomized into the treatment condition visiting any company page that has 20+ individual-level happiness surveys see the full Work Happiness Score widget, as depicted in Panel (b). Job seekers in the control condition see all company pages as usual—with no happiness information—as depicted in Panel (a).

computer, and the job seeker’s county (see Table S1 for summary statistics).¹⁵ Randomization was carried out using the firm’s in-house A/B experimental software. The job seekers in the treatment and control groups are similar in terms of these observable characteristics, and the intended proportion of job seekers in both treatment and control is as intended (see Table S1). There is a marginally statistically significant difference in the proportion of desktop job seekers ($p = 0.058$); however, the magnitude of the difference is extremely small (62.6% versus 62.5%) and a joint test for all covariates is not close to being statistically significant.

At the company level, I am able to observe a number of characteristics, such as whether the company is in the Fortune 500, the number of employees it has, the number of jobs it has listed on the site, as well as other crowdsourced information routinely shown to job seekers, such as reviews and star ratings. Table S2 shows balance across treatment and control groups in these characteristics. In the experiment, job seekers view the happiness score of companies. This is on average 62.8 and is well balanced across treatment and control conditions.

4.3 Empirical Strategy

I begin by estimating a simple model whereby

$$A_{ijt} = \beta T_i + \varepsilon_{ijt}, \tag{1}$$

where A_{ijt} is an indicator variable equal to 1 if the job seeker i clicks to apply for at least one job at company j on calendar day t , and 0 otherwise. T_i is a treatment indicator equal to 1 if the job seeker is in the treatment group, 0 otherwise. ε_{ijt} is an error term that is adjusted for clustering on cookies (the level of randomization).¹⁶

In the main analyses, I estimate linear probability models (LPMs). I multiply the outcome variable by 100, such that it is equal to either 0 or 100. This does not change anything other than making the coefficients easier to interpret as percentage point changes (essentially, multiplying them by 100). The use of LPMs allows for more readily interpretable estimates, particularly once interactions are included in the equation. They also more easily allow for the inclusion of high-dimensional fixed effects. In an appendix, however, I also show that results are similar when estimating logistic regressions (see Table S3).

One concern is that treatment may interact with various characteristics of job seekers and companies—most pertinently, the happiness score that is displayed on the page. To this simple equation, I add date fixed effects as well as a series of observable job seeker and company characteristics. In a more restrictive specification, I omit the observable company characteristics and date fixed effects and instead introduce a company-by-day fixed effect, τ_{jt} , such that I am com-

¹⁵As part of the website’s normal business practices, around 25% of the job seekers I observe also upload their resume to the site, which entails filling out details such as previous education and employment.

¹⁶I also experiment with clustering on companies, company-days, and two-way clustering on job seekers and companies as well as job seekers and company-days. Results are robust to all of these specifications and generally have slightly smaller standard errors than those reported in the main analysis.

paring job seekers seeing the same company on the same day. It is not possible to introduce a jobseeker fixed effect, since job seekers are always in the treatment or control group, but I do introduce into the equation a vector X_i' that controls for the job seeker’s commuting zone, the age of their cookie, and whether or not they are a desktop jobseeker and a registered *Indeed* job seeker.

The sample consists of observations that correspond to jobseeker-company-days. For the main analyses, I limit the sample such that each jobseeker has only one observation per company, which is the first day they visit that company’s page.¹⁷ Equation (1) provides a causal estimate of showing the score, since T_i is randomly allocated by design across job seekers. However, the perhaps more interesting and important question is the extent to which this effect may vary according to the score itself. Here, I follow two broadly different strategies. One is to split the sample according to five bins of the happiness score that correspond to the emojis shown (as in Figure 3). The other is to interact the treatment dummy with a linear term for the happiness score of the company H_{jt} , such that,

$$A_{ijt} = \beta_1 T_i + \beta_2 (T_i \times H_{jt}) + X_i' + \tau_{jt} + \varepsilon_{ijt}. \quad (2)$$

4.4 Main Experimental Results

Showing happiness scores has, on average, a small negative effect on applications. Column (1) of Table 1 reports a pooled treatment effect estimate of -0.273 [95% confidence interval (CI): $-0.350, -0.197$]. In a fuller model of the pooled effect, reported in column (2), which includes controls for job seeker observables as well as company-by-date fixed effects, I find an estimate of -0.276 [95% CI: $-0.351, -0.202$]. The mean of the outcome variable, which is equal to 100 if the job seeker applied for the job and 0 otherwise, is 20.05 in the control group.

The average effect masks significant heterogeneity according to the company’s score. This is to be expected since the treatment is very different depending on whether positive or negative information is being shown to the job seeker. In Figure S3, I simply plot the application rate at different happiness levels, for treatment and control job seekers. Within the regression framework, I interact the treatment dummy with the happiness score, which is z-scored to have a mean of 0 and standard deviation of 1 across the sample. Here, I find in column (3) of Table 1 that the interaction term is oppositely signed and well-defined statistically. This suggests a negative effect that is larger at lower levels of happiness, and a less strong effect at higher levels.

To explore this heterogeneity across scores more fully, in column (4) of Table 1, in my preferred specification I allow the treatment effect to vary more flexibly by the happiness score shown. Here, I divide the score into five bins, according to the cutoffs for which emojis are shown. These are below 50, 50–60, 60–70, 70–80, and 80 or higher. Below 40 and above 90, there are very few observations, as can be seen in Figure 1, which is why the lowest and highest bins have a wider range. The main effect of treatment suggests that for companies with a score below 50 (the


¹⁷I do this because the initial treatment may itself determine whether the job seeker returns to the page or not. In robustness tests, I instead limit the sample to include only one observation per job seeker, in which case I include only the first company-day for each job seekers. Alternatively, I extend to the whole sample, such that, if a job seeker returns to that company’s page, I track them on multiple days.

Table 1: Effect of Showing Happiness Score on Application Behavior

	Applied = 100			
	(1)	(2)	(3)	(4)
Main Effect				
Treated	-0.273*** (0.039)	-0.269*** (0.037)	-0.271*** (0.037)	-0.509*** (0.125)
Interactions: Treated				
× Happiness (z-score)			0.203*** (0.031)	
× score is 50-59				-0.003 (0.135)
× score is 60-69				0.266** (0.132)
× score is 70-79				0.537*** (0.142)
× score is 80-100				0.890*** (0.213)
Observations	37,309,899	37,309,899	37,309,899	37,309,899
User Controls		✓	✓	✓
Company-by-Date FEs		✓	✓	✓

Notes: Robust standard errors are in parentheses, adjusted for clustering on individuals. Linear probability models are estimated, in which the outcome is an indicator variable equal to 100 if the job seeker applies to that company. jobseeker controls: logged-in, desktop job seeker, commuting zone fixed effects, cookie age. In column (4), the omitted happiness score category is 20–49. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 5: Percent Effects of Showing Happiness Score



Note: The graph illustrates the experimental effects of showing company-level happiness information to job seekers looking at a firm's Indeed page on the probability they will apply to that company. Percent effects are calculated using the control group mean of $Pr(\text{Apply})$ in each model. Results are derived from separate LPMs, with the sample split according to the happiness score of the company in question. Controls are included in all models for jobseeker characteristics as well as fixed effects for commuting zone and company-by-day. Standard errors are adjusted for clustering on job seekers. Full results are shown in Table S7.

omitted category), the treatment effect is -0.523 [95% CI: $-0.773, -0.273$]. The main effects of the happiness score bins are not estimated in this regression, since it includes company-by-date fixed effects. However, the first reported interaction term suggests that for companies with a score slightly higher—between 50 and 59—this treatment effect is not discernibly different to when the score is below 50. The interaction effects for the three score bins above 60 are oppositely signed to the main treatment effect and statistically different from zero. Once the score is over 70, the negative effect is canceled out, and once the score is over 80, the treatment effect becomes positive and significantly different from zero (see Figure S7 for a plot of the implied treatment effect at each level of happiness from this interaction).

To make the results more intuitive and interpretable, I instead split the sample according to these five happiness score bins. In this case, I report the percent effect, using the control group mean of the outcome variable in each regression. Figure 5 reports these estimates (for full details of these results, see Table S7). For companies with a happiness score below 50, showing this “bad news” to jobseekers reduced the likelihood of applying to that company by a little over 3%. For companies with scores between 50 and 59, the effect size is again similar. For average happiness companies, namely those with scores between 60 and 69, there is a negative effect of showing the score, albeit a smaller one of around 0.25 percentage points (column (3) of Table S7). For scores that are above average, between 70 and 79, there is no effect of treatment. For very high

happiness companies, with scores over 80, there is a positive effect of displaying this “good news” to job seekers, of around 2%.

The magnitude of these effects is likely to be conservative, in that I am experimentally estimating the effects of showing happiness information (i) over and above the existing plethora of information contained in star ratings, reviews, and so on and (ii) conditional on jobseekers navigating to the company’s page in the first place. The sharp asymmetry in effect accords with a long line of literature in psychology demonstrating that negative stimuli produce stronger responses than equivalently positive ones (Baumeister et al., 2001). Moreover, the asymmetry is perhaps to be expected, given the way in which jobseekers are likely to use the site. A job seeker who is interested in a job at a particular company, or is interested in the company in general, clicks through to the company’s page on *Indeed*—at which point, they can mostly only be discouraged. In fact, another way to interpret the point at which the information no longer has any impact is the level of the job seeker’s subjective beliefs about the happiness of the company in question. Consistent with recent work by Jäger et al. (2021), I find that workers are not perfectly informed about how they’d fare at other firms (see also Reynolds, 1951, for a classic study of the extent to which workers hold correct beliefs about the wage and non-wage conditions of employment at other firms).

4.5 Robustness and Replications

In Table S4, I show that the effects are little changed by including differing combinations of controls and fixed effects. Instead of introducing company-by-day fixed effects into the equation, in Table S4, I use a series of company and day fixed effects and find consistent results—regardless of whether controls for jobseeker observables are also included. Given that the outcome variable is binary, a nonlinear model may be more appropriate to the data. I show, however, that the use of LPMs gives similar results to when using a logistic regression to estimate the treatment effects (see Table S3 and Figure S4). In the main specification, I include in the sample only the first day that I am able to observe in the experiment for each jobseeker–company pair. I can instead use all observations, such that each jobseeker–company pair can appear twice in the data if the job seeker returns to the company’s page on subsequent days. In the other direction, I can restrict the sample such that I include only the first company-day that I observe for each job seeker, to reduce any concerns relating to path dependence of the treatment. In both cases, I find consistent results, as can be seen in Table S4 for the main effect and Table S7 for heterogeneous effects by score.

A further concern is that by randomizing across cookies, there is a possibility that I may observe the same jobseeker twice if she visits from another device. This is particularly a concern if the job seeker is then in both control and treatment group. It is worth noting that such a situation would likely bias any estimates toward zero by adding noise; nevertheless, I can restrict the analysis to the subset of jobseekers who are registered. In Table S9, I add into the main application equation a dummy for if the job seeker is registered and logged in to the platform, and interact it with the main treatment dummy. As before, for ease of interpretation, I split the sample into five,

according to the buckets of the score that determine the emoji shown, such that I am essentially looking at a three-way interaction between treatment, being logged in, and the score displayed as part of treatment. I find that at lower levels of happiness (below 60), the magnitude of the negative treatment effect is significantly stronger for registered job seekers.¹⁸

I also replicate the field experiment in Canada and the UK.¹⁹ The experiments were the same as in the USA, with scores experimentally displayed on eligible company pages. Randomization was carried out across jobseekers (1,542,329 job seekers in the UK and 1,270,398 in Canada). Some 25% of jobseekers were randomized into the treatment group, with the remaining 75% of jobseekers receiving business-as-usual. The studies took place between March and August 2021. Results are reported in Tables S5 and S6 and depicted for ease of interpretation in Figure S6. Similar to the USA, treated jobseekers make fewer job applications when faced with negative information. When the score is below 60, there is around a 3–4% reduction in applications in the UK, and a 4–5% reduction in Canada. In the UK, there is little effect of treatment for scores beyond this. In Canada, scores between 60 and 70 also see a reduction in applications (and no effect beyond this), suggesting potentially higher expectations about the happiness levels of companies in Canada than in the UK.

4.6 Other Outcomes

Although applying to the firm is the main outcome of interest, it is also possible to track various other behavioral outcomes. First, I am able to observe whether or not the job seeker clicks to “follow” the company on the platform, and in doing so be more informed in the future about job openings the company have. Second, I look whether or not the job seeker clicks onto the “reviews” tab of the company’s page, to test the hypothesis that negative news in particular will lead to further information searching on the part of the job seeker, who has been prompted to be more wary of applying to a firm with a poor reputation for workplace happiness. Third, rather than look at the number of jobs the job seeker applies to, I look at the number of job clicks, i.e. how many jobs the job seeker clicks to view the description of at that company.

Results are reported in the three panels of Figure S5. At the lowest levels of happiness (below 50), I find a small negative effect on the number of job clicks. There is no effect between 50 and 70, but showing higher levels of happiness leads jobseekers to click on more jobs that the company has to offer. The provision of negative information reduces follows, but this is not a significant effect. At higher levels of happiness, there is a small increase in follows, but again this is not statistically significant. At the highest and lowest levels of happiness, there is no effect of showing the happiness levels to jobseeker on whether they visit the reviews tab; however, at scores between

¹⁸However, it is worth noting that the base rate of applications (in the control group) is different: Among control group non-registered job seekers, the mean of the dependent variable is 13.49, compared with 25.05 among control group registered job seekers

¹⁹Micro-level data collection of happiness surveys began in August 2020, and by August 2021, there were 123,140 completed surveys in the UK and 86,255 in Canada. Given the generally smaller amount of data compared with the USA, companies are eligible to have a happiness score shown when they have responses from 10 employees. In the UK, this meant 1,196 companies with a score to be shown during the experiment, and 792 in Canada.

Table 2: Effect on Job Hires

	Treatment	Control	Difference
Number of Users	22,275,728	1,170,755	
Tracked Hires	267,934	13,942	
Users with Tracked Hire	1.203%	1.191%	0.012%

Notes: The table shows the number of jobseekers in the experiment who can be successfully tracked by the platform into actual job hires.

60 and 80, there is a small negative effect.

4.7 Effects on Job Hires

A final outcome of interest is job hires. While I am able to observe at a granular level the process of job searching in the applications phase, it is much more difficult to track what happens after that. Nevertheless, the website does track jobseekers through surveys (of both job seekers and employers) to see who actually gets hired. The data available to me here is much smaller and is likely to be a non-random subset of jobseekers who are traceable. Nevertheless, given the very large sample size and the fact that there is little reason to believe that attrition will vary because of treatment status, it is still possible to observe job hires as an outcome in the experiment.

Although the overall effect on applications is negative, there is no discernible difference in job hires across treatment and control groups. As can be seen in Table 2, I observe 13,942 hired among the control group of jobseekers in the sample, and 267,992 in the treatment group. This is 1.19% and 1.20% of job seekers, respectively. If anything, the treated jobseekers in the experiment have around a 1% higher number of hires, but a t-test suggests that the difference is not statistically different from zero ($p = 0.25$).

4.8 Heterogeneity

The experiment randomly varies, across job seekers, the provision of information on workplace happiness. The headline number is always happiness (see Figure 4). However, a further two sub-dimension scores are shown by default, with the full set of sub-dimensions available if the job seeker clicks “see full report.” In Table S8, I interact the treatment dummy with the value of the happiness score (as above in the main results) as well as the 12 sub-dimensions. The interaction effect with happiness is oppositely signed and statistically well defined. However, none of the other interactions is statistically different from zero. Overall, this suggests that jobseekers respond principally to the headline score in their application behavior. Further research may look to better understand the extent to which job seekers might use this information in different stages of job search, as well as the extent to which the provision of a range of well-being measures helps workers to match with employers that will maximize their happiness (i.e., obtain jobs at firms with high levels of the workplace characteristics they value most) rather than take jobs with companies that have the highest aggregate levels of happiness.

Guided by theoretical considerations, I also investigate differences in the treatment effect across subgroups of jobseekers and companies. I first look at differences between registered and non-registered job seekers, since we may expect that jobseekers who are more actively looking for a job—as opposed to just browsing—will be more responsive to happiness information. When browsing through jobs less seriously, the potential downsides of working at a low happiness company may not be as salient. In Table S9, discussed above, I find that the negative effect of showing low-happiness information is stronger among logged-in job seekers. For this subset of registered job seekers, I am able to observe various piece of information about their characteristics.²⁰, which in Table S10, I interact with the treatment effect. I first test for any differential effect by cookie age, since in the job search literature that focuses on wages, it is often found that people’s reservation wage drops the longer that they are unemployed—as the job seeker’s liquidity constraints become more binding (Krueger and Mueller, 2016; Mortensen, 1986). A similar dynamic may be at play here with people’s “reservation happiness”—that is, the longer someone has been looking, the more they are willing to compromise on the desire to work at a high happiness company, given that they are more in need of a job. However, I find no evidence of significant differences by cookie age, which is a proxy for length of time looking for a job (Table S10).

We may expect that someone who is currently employed will be less urgently in need of a job, and so has more ability to “shop around” for a higher happiness company.²¹ In this sense, the job seeker has more power—given that they are more able to pick and choose. However, I find no significant differences in the effect for employed versus unemployed job seekers. Job seekers who have higher education and/or more work experience may also have more power in the labor market. Equally, they may have different preferences, particularly if higher education increases expectations about workplace treatment. The direction of the interaction effect for both education (those with a bachelor’s degree or more) and work experience (cumulative employed months inputted to the job seeker’s resume) is negative at lower levels of displayed happiness, as expected, suggesting that such jobseekers may be more selective. However, the interaction terms are not statistically different from zero (Table S10). Labor market tightness may also play a role in relation to power: In areas where unemployment is high, job seekers are less likely to have the ability to shop around, given that they know competition is more fierce for jobs. At low levels of happiness, I estimate an interaction term that is positive, suggesting that in high unemployment areas, the negative effect of showing low happiness scores is less pronounced; however, this interaction effect is not discernibly different from zero (Table S10).

Testing for treatment effect heterogeneity across types of organization, I first look at firm size. A great deal of the theoretical link between happiness information and job seeker behavior out-

²⁰These are mostly gleaned from the resume section of the platform, where jobseekers input information about their past work experience and so on. Unfortunately this information is limited, and does not include characteristics such as race and gender.

²¹Although a different setting, Marinescu and Skandalis (2021) show that job seekers decrease the quality of jobs they target around the time that their unemployment benefits will run out, both in terms of wage as well as other characteristics such as educational and experience requirements and contract type. This highlights the link between the extent to which a job seeker is in need of a job and their selectivity.

lined above relies on the extent to which this new information helps to solve the worker’s information problem. Firm size is likely to be important here, since people tend to already possess information on larger firms. Thus, the hypothesis here is that if jobseekers are already familiar with very large companies, additional information is not new information—and is not likely to sway job seeker behavior, since they already know what they are getting. A similar dynamic is found for restaurants on *Yelp*, where the effects of star ratings are non-significant for chain restaurants (Luca, 2016). This is also in line with the idea that young organizations face the “liability of newness” (Stinchcombe, 1965), in that they have a credibility problem that makes it more difficult for them to trade with other organizations, compared with larger, more established firms (see also Benson et al., 2020). In column (2) of Table S10, I find that the effect of showing low-happiness information is less negative for firms with a larger number of jobs listed, which may be thought of roughly as a proxy for company size. I also find a positive interaction with a dummy variable equal to 1 if the company has 10,000 or more employees; however, in this case, it is not statistically different from zero. Finally, I am also able to observe the star rating that a company already has on the page, which is displayed to both control and treatment job seekers. This star rating already signals what working at a company may be like. It may be expected that only *additional* information will have a stronger effect on jobseeker behavior. At low levels of happiness (column (2) of Table S10), I find a positive interaction effect with the company’s displayed star rating. In situations where the experiment shows jobseekers a low happiness score, the effect is more negative when the star rating shown alongside is higher, such that the negative happiness information is more surprising. However, I do not find significant differences across star ratings for higher happiness companies.

5 Score-Value Effects

Thus far, we have seen that job seekers respond to information about the happiness of prospective workplaces. This suggests that employers have incentives to improve the well-being of their workforce if they want to attract workers as information about workplace well-being continues to become more widely available. There is, however, a subtle difference between the effect of showing jobseekers scores—even across companies with different levels of happiness—and the effect of a company’s score *per se*. An ideal laboratory experiment would most likely not only vary the provision of information, but also randomize the signal contained in that information as well—for example, by adding a noise parameter to the score each time it is shown. While this was not possible in the field experiment—for various legal, ethical, and logistical reasons—in this section, I nevertheless attempt to make progress in identifying the effect of the score’s value that is displayed to treated job seekers.

5.1 Fixed-Effect Estimates

5.1.1 Identification Strategy

Individual-level data collection from jobseekers on their workplace happiness was ongoing throughout the course of the 10-month experiment, and company-level scores change in real time. As a company receives surveys, the score is updated in real time, providing me with within-company variation in the score over the 10-month period I study. In addition to this within-company variation in the score over time, I also benefit from having micro data on job seekers. In particular, I am able to observe jobseekers as they browse various companies, such that I have within-jobseeker variation in the scores they see. Around a quarter of the treated jobseekers (5,441,500) that I observe during the experiment view two or more company pages with a happiness score. On average, these jobseekers view the pages of 3.5 companies, giving a total of 18,666,892 jobseeker-company-day observations that are usable when I rely on jobseeker fixed effects.

Restricting the sample to treated jobseekers who saw more than one happiness score, I estimate equations of the following form:

$$A_{ijt} = \beta H_{jt} + U_i + C_j + T_t + X'_{jt} + \varepsilon_{ijt} \quad (3)$$

where A_{ijt} is an indicator variable equal to 100 if jobseeker i applies to a job at company j on calendar day t , or 0 otherwise; H_{jt} is the happiness score of company j on day t ; U_i is a jobseeker fixed effect, C_j is a company fixed effect, and T_t is a date fixed effect; ε_{ijt} is an error term that is adjusted for two-way clustering on jobseeker and company.

Whereas in equation (1), the assumptions for identifying the causal effect of the main variable of interest—the treatment dummy—were relatively straightforward, given that the treatment was randomly assigned, here, the identifying assumption when estimating the causal effect β is that the happiness score H_{jt} is exogenous to application decisions, conditional on these job seeker, date, and company fixed effects. One major concern is that some time-varying third variable—which is presumably reputational in nature—may affect the company’s score as well as applications. Usefully in this context, the website shows not only happiness scores, but also reviews and a star rating of the company. I am able to include in the equation a time-varying control for the company’s star rating on the website, which should help to control for any general reputational shocks. I also include controls in the vector X'_{jt} for the number of jobs the company has on the site on that day and the number of happiness surveys that make up the firm’s happiness score.

5.1.2 Results

Column (1) of Table 3 suggests that a one-unit increase in the happiness score—which lies between 20 and 100—is associated with a 0.084 percentage point increase in the probability of applying for a job [95% CI: 0.053, 0.115], from a base of 20.167. This suggests that a one-point increase in the score has around a 0.41% positive effect on application behavior. A one standard deviation

Table 3: Effect of a Company's Score on Application Behavior

	Applied = 100				
	(1)	(2)	(3)	(4)	(5)
Linear					
Happiness	0.084*** (0.016)			0.077*** (0.016)	0.077*** (0.020)
Piecewise Linear					
Spline: below mean		0.144*** (0.018)			
Spline: above mean		0.012 (0.028)			
Piecewise Linear					
Spline: 20-39			0.390*** (0.128)		
Spline: 40-49			0.143** (0.056)		
Spline: 50-59			0.134*** (0.021)		
Spline: 60-69			0.070*** (0.026)		
Spline: 70-79			-0.061* (0.037)		
Spline: 80-89			0.093 (0.084)		
Spline: 90-100			-0.115 (0.343)		
Polynomial					
Happiness ²				-0.003*** (0.001)	-0.003*** (0.001)
Happiness ³					0.000 (0.000)
Company FEs	✓	✓	✓	✓	✓
Job Seeker FEs	✓	✓	✓	✓	✓
Date FEs	✓	✓	✓	✓	✓
Observations	18,660,412	18,660,412	18,660,412	18,660,412	18,660,412
R2	0.423	0.423	0.423	0.423	0.423

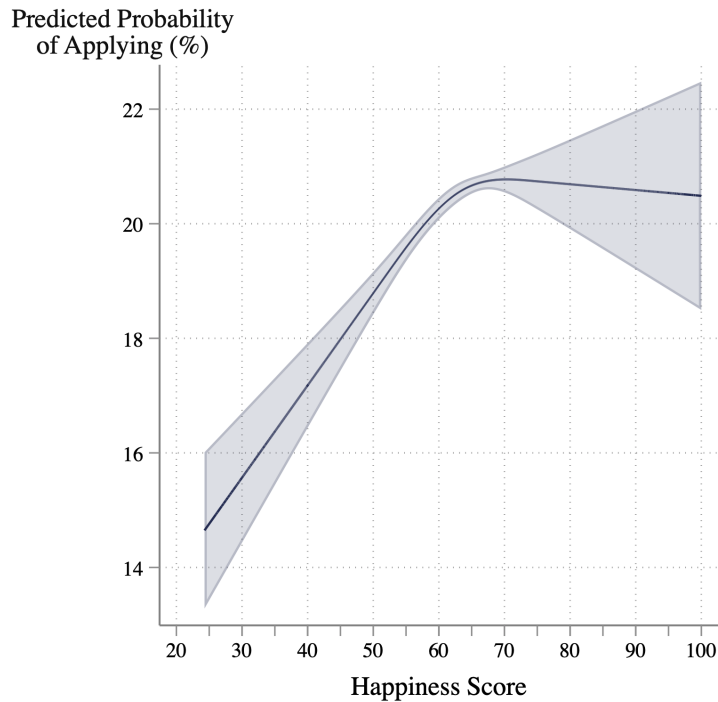
Notes: Robust standard errors are in parentheses, adjusted for clustering on jobseekers and company. Outcome in all models is equal to 100 if the job seeker applied to a job at that company, 0 otherwise. Time-varying company controls are included in all models for number of happiness surveys, company's star rating displayed on company pages, and number of jobs listed by the company. Linear probability models are reported. Happiness score is re-centered around 0 in models (4) and (5). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

increase in the score increases application probability by 3.65%, on average. However, given the findings of the prior field experiment, there is reason to suspect that any effects the score may have are likely to be nonlinear since jobseekers were more sensitive in their application behavior to being exposed to negative information than they were to positive information about the hap-

piness of prospective companies. In Figure S9, I introduce into the regression, indicator variables for a series of equally spaced happiness score bins (leaving out the 60–69 indicator, which is the bin that contains the mean value of the score) instead of the linear happiness term H_{jt} . Compared with scores between 60 and 69, lower scores are significantly associated with lower application probabilities. Improving the score beyond the 60–69 bin, however, does not have much of an effect. Although the point estimates for scores of 70 and above are positive, they are not statistically different from zero.

Building on this, I estimate piecewise linear regressions using splines of the happiness score variable. I first split happiness above and below the mean. In column (2) of Table 3, I find a significant positive effect of increasing the happiness score up to the mean, with a coefficient of 0.145 [95% CI: 0.107, 0.179], which is around 1.7 times greater than the simple linear coefficient shown in column (1). Below the mean of 63, a one standard deviation increase in the happiness score increases the application probability by around 6.2%. Above the mean, however, I find a positive coefficient of 0.012 [95% CI: -0.042, 0.0675], which is not statistically different from zero.

Figure 6: Non-Linear Effects of Happiness on Application Behavior



Note: The figure plots a linear probability model in which the outcome is an indicator variable equal to 100 if the job seeker applies to that company. The happiness score is a restricted cubic spline with 3 knots located at percentiles suggested by Harrell (2015). The regression includes company fixed effects, job seeker fixed effects, date fixed effects, and a set of time-varying company controls. 95% confidence intervals are reported, using standard errors that are adjusted for two-way clustering on companies and job seekers.

However, splitting above and below the mean is relatively crude. In column (3) of Table 3, I go further in separating the happiness score into seven linear splines, with equally spaced cut-points between 40 and 90. The slope is steepest at the lowest levels of happiness and becomes steadily less steep as the score increases. For the four splines where the score is below 70, the piecewise linear effect is in each case positive and statistically different from zero. As the score increases, the magnitude of this effect declines. Over 70, the slope becomes indistinguishable from zero.

In addition to this piecewise linear strategy, I use a restricted cubic spline (see Figure 6).²² This provides similar results to the piecewise linear estimates, with application probabilities increasing as the score goes up to around 65, whereupon the relationship flattens out. Using this restricted cubic spline model, I find that for observations with scores below 50 (“Low”) the average predicted probability of applying is 18.07 whereas this rises to 20.58 for observations with scores between 60 and 69 (“Average”). This represents a 13.9% increase in the probability of applying, going from Low to Average. Going above Average, however, there is little effect on the probability of applying.²³

5.2 Local Randomization Regression Discontinuity Estimates

In this section, I build on the fixed-effects estimates by making use of discrete jumps in the score that lead to a different emoji being displayed to the job seeker. This, with appropriate assumptions, provides an approximation of what we might call a framing field experiment (cf., List, 2007). For an alternative causal inference strategy, with a different set of identifying assumptions, see Appendix Appendix H, where I instrument the happiness score using plausibly exogenous variation in the source of individual-level happiness scores across companies.

5.2.1 Identification Strategy

The happiness score is shown to the nearest integer on a company’s page. Recall that accompanying the score is an emoji and piece of text describing where in the overall distribution the company’s score is (Figure 3). The emojis and distributional signposts (low, below average, average, above average, and high) are fixed during the whole period and are assigned according to five bins of the score—with integer cutoffs at 50, 60, 70, and 80.

For example, once a company goes from a mean happiness level of 69.4 to 69.5, three things happen. First, the score displayed goes up by one unit, from 69 to 70. Not only does it go up by one integer overall, but so too does the left-hand digit. This is important since there is a well-documented “left digit bias” in consumer behavior, which suggests that people’s judgments

²²In this model, I create a restricted cubic spline of the score with three knots, located at percentiles suggested by Harrell (2015), and introduce it into the application equation. In Figure 6, I report an alternative model where I specify the knots to be located at 50, 60, 70, and 80.

²³Introducing into the equation a quadratic of the happiness score also confirms these findings. This regression is reported in column (4) of Table 3. One benefit of the quadratic term is that, while being a reasonably good approximation of the functional form established in the spline and categorical analyses, its properties are well understood in two-stage least squares (2SLS) settings using instrumental variables.

are disproportionately influenced by the leftmost digit in a number (see, e.g., Bhattacharya et al., 2012; Lacetera et al., 2012). Second, the neutral-face emoji that appears next to the score is replaced by one with a happier face. Third, the text next to this emoji changes from “Average” to “Above Average.”

I restrict the sample to observations where the happiness score is within a window of 0.2 either side of the four cutoffs (49.5, 59.5, 69.5, and 79.5). On the original 1–5 scale on which happiness is measured, this corresponds to a mean score that is 0.01 either side. That is, if the score were shown on the 1–5 scale and the threshold were 4, this restriction rule would limit the sample to include only instances where the score is in the window between 3.99 and 4.01. I make the identifying assumption that within this tight window, the treatment of having the score rounded up can be thought of as being “as-good-as-randomly assigned.” This local randomization regression discontinuity (RD) design implies that within the window, we can think of the latent score as being the same—and see the rounding-up treatment in a similar way to it being an experiment (for a more formal discussion of this local randomization RD approach, see Cattaneo et al. (2015) and Cattaneo et al. (2017); see Sockin and Sojourner (2020) for a recent example of this approach using star ratings on *Glassdoor*).

5.2.2 Results

In Table 4, I find a positive effect of rounding up of the score. In the simplest model—reported in column (1)—I take all jobseeker-company-day observations that fall up to 0.2 either side of the four cutoffs, pool them, and regress the application outcome on a dummy equal to 1 if the score is equal to or above the cutoff (and a set of fixed effects for the cutoffs). The coefficient is positive and statistically significant [95% CI: 0.071, 0.332], suggesting that tipping over the threshold makes jobseekers more likely to apply to the company. In the remaining columns, I add into the equation a series of fixed effects and further control variables, which do little to affect the result—suggesting that within this tight window, being either side of the threshold is not likely to be influenced by the characteristics of companies or job seekers. In Figure S13, I vary the window size around each of the thresholds, up to a maximum of 0.5 either side. This does not significantly alter the pattern of results. Finally, as an additional sensitivity check, I follow the randomization inference approach of Cattaneo et al. (2015) and again find similar results (Table S12).

The issue of whether or not changes in a company’s score over time have causal effects on applications is an important one, since it directly addresses the incentives that companies face to invest in management and organizational practices that are conducive to happier workers. In this section, I have followed three strategies: fixed effects, regression discontinuity, and instrumental variables. Each approach has advantages and disadvantages, with varying degrees of plausibility to their identifying assumptions—and each should be interpreted with caution since there is no explicit randomization. Nevertheless, taken together, the three different approaches converge on the conclusion that improvements to a company’s score have a positive effect on job seekers’ probability of applying to that company.

Table 4: Discontinuity Estimates of the Effect of Happiness Score on Application Behavior

	Applied = 100			
	(1)	(2)	(3)	(4)
Rounded Up	0.202*** (0.067)	0.233*** (0.087)	0.203** (0.086)	0.211** (0.088)
Observations	1,453,747	1,453,747	1,453,747	1,453,747
Cutoff Point FEs	✓	✓	✓	✓
Date FEs		✓	✓	✓
Company FEs		✓	✓	✓
User Observables			✓	✓
Extra Controls				✓

Notes: Robust standard errors are in parentheses, adjusted for clustering on job seekers. Outcome in all models is equal to 100 if the job seeker applied to a job at that company, 0 otherwise. The sample is all jobseeker-company-days within 0.2 points of 49.5, 59.5, 69.5, and 79.5. Rounded up is equal to 1 if the score is above the threshold and thus rounded up to the nearest integer, which triggers the display of a different emoji. jobseeker observables: logged in, desktop job seeker, commuting zone, cookie age. Extra controls: number of happiness surveys, company star rating, number of jobs listed by company. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6 Nationally Representative Survey

I turn in this section to present evidence from a survey that is representative of the US labor force. The survey was in collaboration with *Indeed* in April 2021 and included 4,033 respondents who were interviewed online. The survey was fielded by Forrester Research. To ensure a broadly representative sample, quotas were set by age, education, gender, region, and income level (see Appendix Appendix J for more details of the survey).

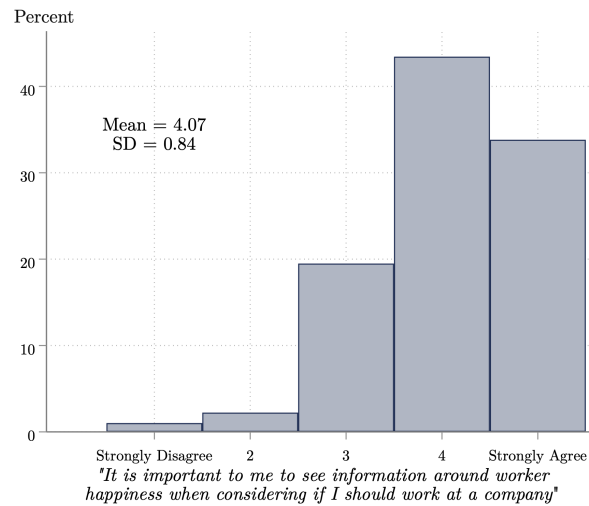
To better understand how people respond to information about others' happiness in the labor market, all employed and unemployed respondents were first asked to:

Imagine that a website has asked millions of workers across the country about their happiness at work (including various dimensions like whether people feel a sense of purpose, belonging, trust, appreciation, growth, etc.), in order to help job seekers make more informed decisions about where to work. After surveying 4.5 million people working at different companies, this company has published this new data for each company.

Following this introduction to the concept of the happiness score, respondents were asked a relatively simple, non-behavioral question about the usefulness of this information for making job choice decisions. Figure 7 reports the distribution of responses to this question, showing that 77% of respondents agree or strongly agree that it would be important to them to see information about worker happiness when considering whether to work at a company. To build on this initial piece of suggestive evidence, I attempted to elicit respondents' willingness to pay for workplace happiness as a job amenity. To do so, respondents were asked:

Imagine you are looking for a job on the website we have described. You are comparing two positions, both of which are in the same industry and location as your [current/last] job.

Figure 7: Preferences for Company Happiness Information



Note: The figure shows the distribution of responses on a 1–5 agreement scale to the statement shown. The sample is 4,033 survey responses from an online survey representative of the US labor force.

In this instance, assume the positions, job description and companies are the same outside of pay and happiness levels. Which position would you be more likely to choose?

1. *Position A: The workplace happiness score of the company is 65. This position pays you the same as your [current/last] job.*
2. *Position B: The workplace happiness score of the company is [45/55/75/85]. The position pays you [2, 5, 10, 20, 35]% [more/less] than your [current/last] job.*

In the cases where happiness is higher in Position B, respondents were asked about the extent to which they would be willing to be paid less; in cases where happiness is lower in Position B, they were asked how much more they would need to be paid to take it. Half the sample were randomly assigned to be in the group that was asked two questions about higher levels of happiness, and the other half was asked two questions about lower levels of happiness (giving a total of around 8,000 choice scenarios).

In all cases, the default job is 65. One possibility is that the choice of baseline job may bias any estimates. Encouragingly, the similar approach of Maestas et al. (2018) found little evidence of sensitivity of estimates according to the baseline job they used. Nevertheless, in order to calibrate the experiment, I also first asked respondents to estimate the happiness scores of their own company and industry. As can be seen in Figure S17, the mean response is 62.7 and 62.6, respectively (compared with the actual mean across companies in the *Indeed* data of 62.8). However, despite this, the spread of responses is greater: the standard deviation when asked to guess the happiness level of their company and industry is 20.1 and 17.8, respectively (compared with the actual standard deviation in the *Indeed* data of 8.8).

In the instances where a respondent was asked to contemplate switching to a job with lower happiness but higher pay, they were first asked whether they would accept a 2% pay raise for a given loss of happiness—and if they declined, they were then sequentially offered 5%, 10%, 20%, and 35% pay cuts until they accepted (or declined even at 35%). In the instances where the respondent was asked to consider switching to a job with higher happiness but lower pay, they were first asked whether they would accept a 35% pay raise, and if they declined, were sequentially offered 20%, 10%, 5%, and then 2% pay raises.²⁴

Results from this exercise are reported in Figure 8. In Panel (a), I plot the cumulative share of respondents choosing Position B at differing levels of salary, in the instances where Position B involves a move down from a company with a score of 65 to a company with a score of either 45 or 55. When comparing between staying at a company with a score of 65 to a much less happy company with a score of 45, when offered a 35% increase in pay, around 70% would switch—but 30% still say they would not go to the company with a much lower happiness score. Offered a 20% increase in pay, 60% would switch. Around 40% of those offered a 10% pay raise would take the job in the much less happy company. A similar pattern is found when offering respondents hypothetical positions at a company with a happiness level of 55, though willingness to switch to such a job is lower, as would be expected, given the smaller jump downward in happiness. Panel (b) of Figure 8 reports the extent to which respondents are willing to accept a lower level of pay to move to a company that has a generally happy workforce. The pattern of results is similar. Overall, the data suggest that workers value happiness as a workplace amenity and would be willing to trade off salary to get it.

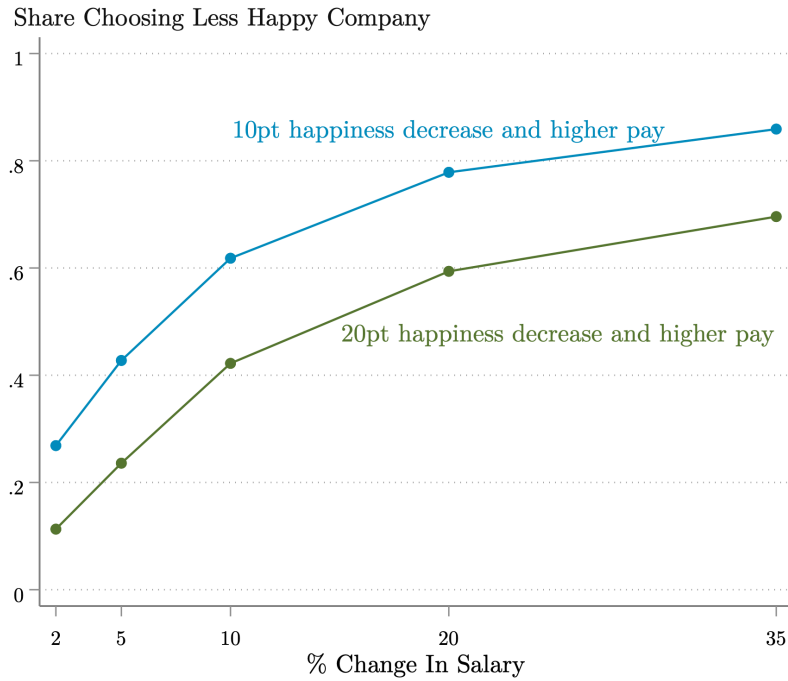
While Figure 8 provides the distribution of respondents' willingness to pay for higher levels of happiness, it is also informative to calculate the mean level in each of the four cases (a 10-point increase, 20-point increase, 10-point decrease, and 20-point decrease in happiness).²⁵ Table S13 reports mean pay changes across the sample at which respondents would be willing to choose Position B—i.e., switch away from the job that has a happiness score of 65. To take a job at a company with a score of 75 (compared with 65), respondents would be willing to take, on average, a 12.7% pay cut. Using the standard deviation of the happiness score across the 20,000 or so companies on *Indeed* of 8.83, to take a job at a company with a score of one standard deviation below the mean, this exercise suggests that people would be willing to take, on average, a 10.6% pay cut.

Nevertheless, this mean figure masks a large amount of heterogeneity. Indeed, even though just over 30% of people would take a 35% pay cut to work at a company with a high happiness score (85 compared to 65), nearly 20% would not even take a 2% pay cut to do so. However, in contrast to the field experience—where I was somewhat limited in the characteristics I could

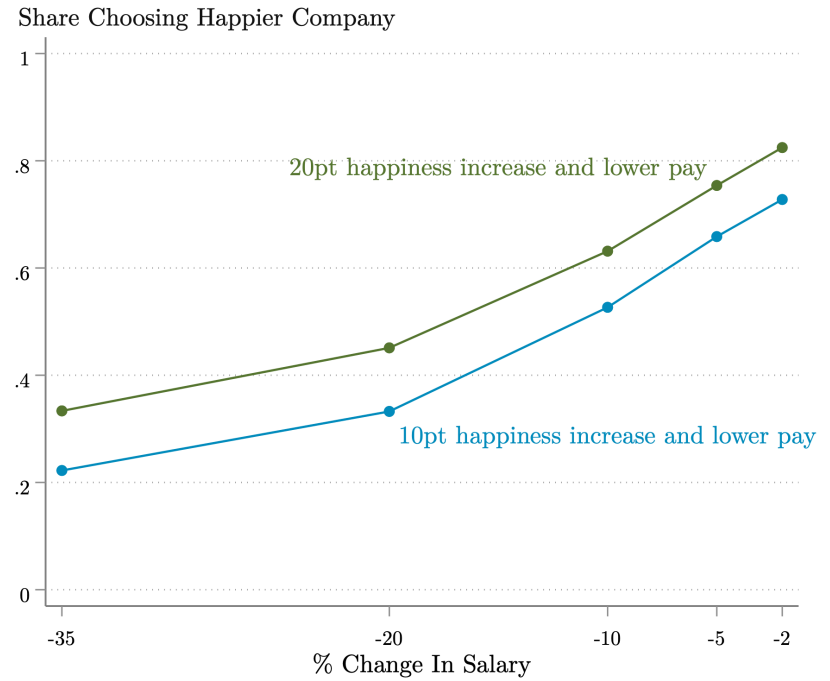
²⁴For more details of question wording and procedure, see Appendix Appendix J.

²⁵Given the design of the survey instrument, in each case, there are some respondents who chose Position A regardless of the wage offered for Position B. Here, I make the assumption that in the cases where the sequential wage offer was declining, those who refused to switch at 2% can be coded as 0%. In the other direction, I make the somewhat conservative assumption that those refusing even when offered a 35% pay increase can be coded as 35%.

Figure 8: Survey Evidence of Willingness To Pay for Workplace Happiness



(a) Higher Pay, Lower Happiness



(b) Lower Pay, Higher Happiness

Note: In all cases, the base job is at a company with a happiness score of 65. Panel (a) reports the share of respondents preferring the company with lower happiness, at differing levels of pay increase. Panel (b) reports the share of respondents preferring the company with higher happiness, at differing levels of pay decrease.

observe of job seekers—the survey allows for some exploratory analyses of who is more or less sensitive to happiness. In Table S14, I regress the wage at which the respondent switches to position B on a range of observables including race, gender, income, and education.

In terms of the wage cut people would be willing to take to work at happier companies, younger people say they would need less of a cut to do so. This suggests that on average younger people have higher tastes for happier workplaces. Equally, richer respondents would generally be willing to take larger pay cuts to work at happier firms. Looking instead at the wage raises that people would need to work at happier companies, highly education people (with at least a Bachelor's degree) would generally require a large raise to work at very low happiness firms. Those who are actively looking for a job are less sensitive to the happiness of prospective firms, however, and would generally need less of a raise to avoid unhappy workplaces. Happier individuals are generally less willing to work at low happiness firms – that is, they would require larger wage raises to do so. On the other hand, I find no systematic difference across race, gender, or parental status in the trade-offs people are willing to make, in either direction. Additional research is required in order to better understand these differences, particularly the extent to which they reflect differing tastes versus labor market constraints.

As with the field experiment, I also conducted the same analysis in Canada and the UK, in each case collecting a survey sample of around 1,500 respondents that is representative of the labor force. The distribution of responses is similar in all three countries, as can be seen in Figures S18 and S19. Mean levels of the willingness to trade off salary for working at a happier company are slightly lower than in the USA—to take a job with a happiness score of 75 (rather than 65), respondents are, on average, willing to be paid 10.5% and 11.3% less in the UK and Canada, respectively.

Naturally, the experiment is abstract and should be interpreted with the appropriate amount of caution. Nevertheless, people's preferences in the abstract world of hypothetical choice scenarios are consistent with behavior in the natural field experiment across all three countries, where job seekers view real companies and make real-stakes decisions on where to apply.

7 Discussion

When surveyed, around 87% of managers in the USA subscribe to the belief that workplace well-being can provide their firm with a competitive advantage (HBR Analytical Services, 2020). One explanation for this is a well-established link between happiness and productivity (see, e.g., Bellet et al., 2021; Oswald et al., 2015; Walsh et al., 2018). Nevertheless, this is not the only reason firms may routinely measure happiness, make it a strategic priority, and invest in organizational practices conducive to a happier workforce. 94% of managers believe that greater employee happiness would make it easier to attract workers (HBR Analytical Services, 2020). However, although this is seen as a key main mechanism linking workplace happiness to firm performance, empirical evidence is currently lacking. In this paper, I demonstrate that job seekers respond to information

about the happiness of workplaces to which they are considering applying.

7.1 Policy and Managerial Implications

My main finding is that workers place a positive value on workplace happiness over and above other key aspects of jobs such as salary and security, which are held constant. This suggests that companies face incentives to look after the happiness of their employees if they want to attract a larger pool of applicants. This is particularly the case as information about the well-being of workers at different companies becomes more widely accessible because of the growing digitization of the labor market.

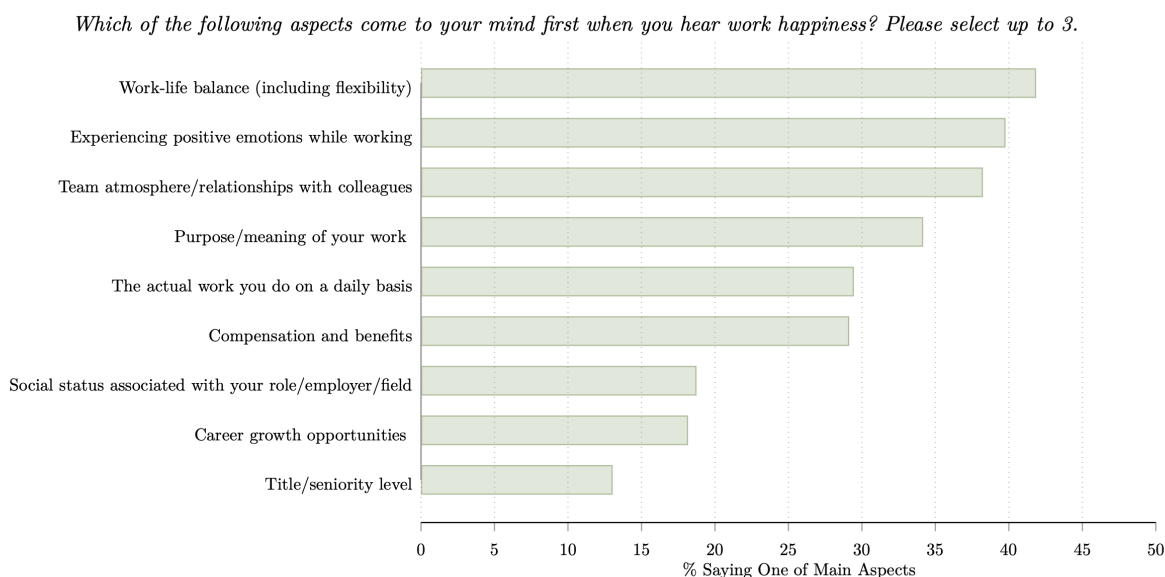
Figure S2 takes the average happiness score of the 20,000 or so companies in the study, and regresses it on a series of 4-digit industry codes. Plotting the residuals from this regression, it is clear that there is significant variation in happiness across companies even within tightly-defined industries – that is to say, companies that are observationally similar and facing the same business environment. This paper shows that this variation has downstream effects on behavior in the labor market (at least, once it is revealed). This begs a number of follow-up questions for firms and policymakers, such as i) what drives workplace happiness, ii) whether it can be influenced by managerial decisions, and iii) what job seekers are encoding when they see workplace happiness scores.

7.1.1 What do job seekers understand by workplace happiness?

The follow-up survey, which is representative of the US labor force at large, allows me to look more closely at what people understand by the term workplace happiness. That is, when a job seeker sees a high or low work happiness score on the platform, what is she likely to be inferring from that? I ask people a relatively straightforward question: “Which of the following aspects come to your mind first when you hear work happiness?” Respondents were asked to select up to 3. As can be seen in Figure 9, the most common response is work-life balance, followed by social relationships, and enjoyment while at work as well as purpose derived from it. At least part of workplace happiness is understood to be about compensations and benefits, though less so than the aforementioned aspects of work.

While more research is required in order to fully understand what job seekers are encoding when viewing happiness scores, potentially with job seekers on the platform itself, this relatively simple survey exercise suggests that people view it as an overall summary measure that is a function of various aspects of workplaces, including but far from dominated by wages and benefits. This is largely in line with a growing literature on the determinants of happiness in the workplace. Given that much of the existing information shown on *Indeed* is focused on things like salaries and career growth opportunities, adding a measure of workplace happiness appears to be providing a new layer of additional information that is useful to job seekers.

Figure 9: How do people interpret workplace happiness?



Note: N=4,033 from a representative sample of the US labor force. See text for further details.

7.1.2 Can firms influence the happiness of their workers?

One initial approach to answering the question of whether firms can reasonably be expected to do anything is to ask workers themselves. I included in the individual-level follow-up survey about who is responsible for employees' happiness. Here, I find that people say that individuals themselves take on at least some responsibility, but they also see a large role for management—from the CEOs down to line managers—to influence the well-being of workers (Figure S15). Similarly, from the other side, in a survey of US managers by HBR Analytical Services (2020), respondents were asked how much control they believed their organization has when it comes to influencing the happiness of its workforce. Some 38% said a high degree of control, and 57% said some control, with only a small number saying that organizational and management practices are irrelevant to people's happiness (HBR Analytical Services, 2020).

In line with this intuition, a growing body of work on happiness in the workplace suggests that employee subjective wellbeing—including measures of job satisfaction, happiness, and purpose—is at least determined by structural factors related to how work is organized as well as cultural factors and relationships within organizations. One strand of this literature regresses measures of subjective wellbeing against a range of workplace characteristics, in order to derive relative importance weights (see, e.g., Clark, 2010; Krekel et al., 2019). Results from this approach suggest that pay is an important predictor of workplace wellbeing, but that it is frequently not the variable that explains most of the variance in the happiness of workers.²⁶ Interpersonal relationships—both hor-

²⁶Indeed, even within the context of pay, it is typically found that happiness is more strongly predicted by relative pay rather than absolute pay, suggesting a key need for fairness in pay practices (see Breza et al., 2018; Card et al.,

izontally between co-workers and vertically between managers and staff—are highly important, as are things such as flexibility and work arrangements that are conducive to work–life balance.

This initial approach is potentially limited by only being able to make only cross-sectional comparisons between workers. Nevertheless, a more recent line of work is field experimental in nature – with researchers experimentally trying different managerial and organizational practices, and tracking measures of employee happiness. For example, Gosnell et al. (2020) ran a field experiment for airline pilots and found effects on workplace happiness of various management practices, including monitoring, performance information feedback, personal targets, and pro-social incentives. Equally, Breza et al. (2018) show that pay inequality can have large effects on happiness (see also Cullen and Perez-Truglia, 2018), while Bloom et al. (2014) find that working from home improves emotional experience as well as satisfaction of call center workers. Also in relation to autonomy, Moen et al. (2016) find that a program designed to increase supervisor support and employee independence improved various measures of worker subjective wellbeing.

This burgeoning literature on the causal determinants of subjective wellbeing demonstrates at least three things: i) organizational practices can be changed, ii) employee happiness is at least partly determined by these managerial practices, and iii) that there are specific areas that may be targeted based on this research. Nevertheless, more research is needed in order more fully understand the determinants of worker happiness, particularly how they may differ across different industries, occupations, and demographic characteristics of workers such as age, race, and gender.

7.2 Crowdsourced Happiness Data

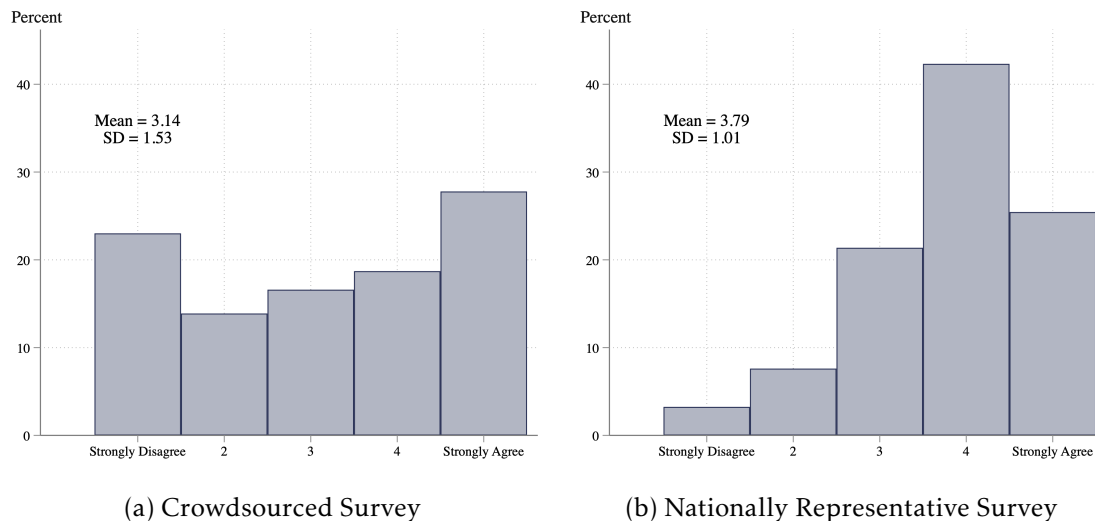
Subjective wellbeing in the workplace may be measured in a number of ways, such as surveys and, more recently, approaches like gleaning emotional states using natural language processing from social media posts or emails. One innovation of this paper is to use crowdsourced data. This has the key benefit of being able to provide data measured in a consistent way across a very large number of firms. However, the underlying sample is not random and not likely to be representative. Job seekers select into visiting the website, and then select into answering the happiness survey. The data is thus subject to a number of caveats that are common when working with such crowdsourced data. Indeed, one may expect that respondents selecting into filling in the survey may be likely to be either very happy or very unhappy, either of which may motivate them to respond.

In this paper, this is largely a moot point. I am interested in the ways in which job seekers *respond* to aggregated company-level summaries of crowdsourced information about employee happiness. The extent to which the underlying data may be noisy or biased in different ways falls outside the scope of this paper and should be subject to further research.

Nevertheless, in order to shed light on the issue, I included in the nationally representative survey the same questions that are asked on the platform. Panel (a) of Figure 10 shows the individual-level crowdsourced responses on *Indeed*. The two most common responses are 1 and 5—that is,

2012).

Figure 10: Individual-level responses across survey modes: “I feel happy at work most of the time”



Note: Panel (a) shows the distribution of 2,357,522 happiness responses from individuals in the USA who responded to the online survey between October 2019 and March 2021 in relation to a company that had at least 20 responses over the period. Panel (b) shows the distribution of 4,033 happiness responses collected in an online survey in March 2021 from a sample of US individuals that is representative of the US labor force.

strongly disagree and strongly agree. This sort of bimodal distribution is consistent with the idea that people answering these surveys are likely to be either very happy or very unhappy with their employer, leading them to answer a survey. In contrast, the representative sample shown in Panel (b) of Figure 10 is more normally distributed. While the disproportionately high 1 and 5 responses in the crowdsourced data to some extent cancel each other out when it comes to the overall mean, the very-unhappy-responder bias dominates, such that a given company’s score is likely to be lower than its “true” level of worker happiness. However, making comparisons between aggregate measures of company happiness (assuming a relatively high threshold for the number of surveys needed in order to calculate the score), this will only be problematic to the extent that bias varies systematically across companies. Ultimately, additional research is required in order to more fully understand this new source of large-scale data on workplace data; but the initial evidence suggests that it may provide an exciting “worker’s eye view” of thousands of organizations across the USA (and increasingly also other countries around the world) that could allow for the testing of a range of interesting hypotheses across the social and behavioral sciences.

7.3 Magnitude of Experimental Effects

During the experiment, jobseekers continued to see employee reviews, companies’ overall star ratings based on employee feedback, as well as information about salaries and other question-and-answer content. It is reasonable to expect that the magnitude of the effect of the provision of additional information about workplace happiness would be larger if it were provided by itself

versus nothing. The effects estimated in this experiment are the *additional* effects, conditional on all of this prior information. It is also the case that the effects presented here are conditional on a jobseeker navigating to a company’s page to see (or not) the score—that is, I study companies presumably already in job seekers’ choice sets. This is likely to underestimate the size of the happiness effect if companies’ scores also serve to attract job seekers to their pages in the first place.

The strength of the treatment effect is greater when negative information is shown, as compared with equivalently (in terms of distance from the mean score) positive information.²⁷ Job seekers appear to use the happiness score information to screen out miserable workplaces from their consideration. This accords with a long line of research in psychology that shows that “bad is stronger than good” (Baumeister et al., 2001). An alternative explanation is that workers are generally risk averse, such that they place a stronger value on negative information about companies they are considering working for (cf. Sockin and Sojourner, 2020). In addition, the asymmetric effect also makes sense insofar as jobseekers navigate the company pages of the platform when they are already interested in applying to that company. Having seen a job listed for the company elsewhere on the site and become interested in applying, jobseekers may then go to their page. At this point, they can really only be discouraged from applying.

One puzzle is that jobseekers have a negative (albeit small) reaction to happiness scores that are around the mean. The explanation for this should be subject to further research. One interpretation is that jobseekers are attentive and fully understand the scale, but are motivated to avoid “average” happiness companies as this is below what they want. Another interpretation is that a score of 65 is seen as low, particularly in the USA, where scores out of 100 are often intuitively understood on the basis of things such as educational grade scores in which 65 would be relatively low. This may be particularly the case if jobseekers are inattentive—that is, even though there is text next to the score saying that 65 is average, job seekers may not fully take in that piece of information and instead focus on the more pertinent piece of information—the number. Another interpretation may relate to so-called grade inflation. Filippas et al. (2020) suggest that ratings have become inflated over time on many platforms such as *Amazon*, *eBay*, and *Airbnb*, since raters feel pressure to give above-average scores. Given this, it might be that jobseekers have become so accustomed to skewed distributions when it comes to online reputation systems, where scores below 4 out of 5 are now frequently dismissed by jobseekers as very low, that their behavior is affected even in situations where scores are more normally distributed.

²⁷For companies with scores below 60, I find an estimated treatment effect of around -2.75% . Over the course of the experiment—taking all jobseeker-company-day observations—2,591,043 treated jobseekers applied to companies with scores below 60. In the absence of treatment, we would expect 71,254 more applications over the 10-month period. For companies at the other end of the spectrum—with scores over 80 (of which there are far fewer)—job seekers in the treatment group made 210,749 applications. Using the estimated treatment effect of 2% , the provision of the score led to around 4,215 more applications to these high happiness firms.

7.4 Generalizability

How generalizable are these findings? Here, I follow List (2020) in discussing the so-called SANS (selection–attrition–naturalness scaling) conditions. In terms of selection, the field experiment is not a random selection of the labor force, but it is nevertheless the population of online job seekers who navigate to company pages on *Indeed*. Moreover, results from a nationally representative survey provide evidence of similar patterns. Having access to administrative data on jobseeker behavior on the site, attrition is not a major concern. The field experiment is at-scale and is natural in that subjects do not know they are taking part. Altogether, the experiment has high generalizability. Nevertheless, the study does have a number of limitations.

7.5 Limitations and Extensions

The experiment was not able to randomly vary the score itself or the salary of the jobs offered. There is a trade-off here between being able to run a large-scale natural field experiment that is highly generalizable and the level of experimental control. To add to the experimental evidence of showing the score, versus not, I exploit discrete jumps in the score owing to rounding rules. This analysis suggests that the positive relationship between happiness scores and applications is a causal one. Moreover, I turn to a nationally representative survey in which I embed a survey experiment designed to elicit willingness to pay for workplace happiness. In this scenario, I have much more control over the salary and happiness score, but this equally comes at the expense of presenting respondents with highly stylized hypothetical scenarios. Taken together, however, the evidence from (i) the natural field experiment, (ii) the analysis of archival data on the website among treated job seekers, and (iii) the vignette study provide a broad picture that suggests that workers value happiness.

Although I have been able to quantify the extent to which improving workplace happiness raises the size of the applicant pool, which I interpret as meaning that there are incentives for companies to improve their score, two caveats are worth noting. First, one needs to know the marginal cost of improving workplace happiness, on the one hand, as well as the comparative benefits of a larger applicant pool on the other. In both cases, the strength of the incentives faced by firms to improve employee happiness will depend on a number of factors. Second, instead of providing incentives to improve workers' happiness, an alternative interpretation is that firms instead simply face an incentive to improve their score on crowdsourcing websites, which may well involve manipulation of these scores rather than any changes in employment or management practices. Sockin and Sojourner (2020) discuss the issue of retaliation by employers against employees who give negative reviews of companies on online jobs platforms such as *Glassdoor*. I cannot observe any manipulation (either overtly or through the threat of retaliation) directly in the data. However, it is worth noting that the website screens surveys for data quality and abnormal activity—for example by bots. Further research should investigate these dynamics further to understand more deeply the incentives faced by organizations.

Several further extensions of the experiment would add to the findings. For example, a very useful extension would be gather follow-up data that would allow not only to track the extent to which treated and control jobseekers get hired, but also the extent to which there are potential differences in terms of match quality. This may be in terms of how happy they are with the jobs they end up, or a more revealed preference approach of how long they stay in that job. Further, it would also be useful to display happiness information on job adverts themselves, rather than on company pages. The majority of job seekers on the website search for jobs using the search function, rather than navigating to company pages and deciding whether or not to apply. This would also allow for more fine-grained heterogeneity analysis, across occupations and types of job, rather than at the coarser level of company characteristics. A further extension would investigate more fully the sub-dimensions of happiness in the scores shown on the website. Experimentally manipulating which of these are displayed, as well as their scores, would shed light more fully on which workplace characteristics people value the most and the extent to which the provision of information on fine-grained aspects of work life might enable for better matching between workers and firms, assuming that workers likely have heterogeneous preferences for things such as flexibility, purpose, and appreciation.

Finally, the experiment took place in North America and the UK. The extent to which preferences for workplace happiness are culturally specific is not clear and cannot be estimated in these data. Further research is required to test the replicability of the findings in different contexts. Moreover, while it was not possible in this instance to observe demographic characteristics such as race and gender in the field experiment, further work is required in order to investigate the extent to which demographic and socioeconomic characteristics might moderate the main effects.

8 Conclusion

Happiness is seen as a key outcome of interest by a growing number of governments and public policymakers. For example, a number of national statistical offices are routinely collecting data on the SWB of citizens, and various governments are now using such data as a means of evaluating and informing their policy priorities and decisions (Graham et al., 2018; Krueger and Stone, 2014). A key question thus arises as to what sorts of labor market policies and institutions might serve to help improve well-being. The rapidly increasing digitization of the labor market, which has the ability to ease the flow of information to workers, is one potential—but unstudied—avenue through which to discipline opportunistic firms whose managerial and organizational practices typically induce low levels of happiness among their workforce.

Employee happiness is increasingly discussed by companies as a priority, a development that has only been accelerated by the COVID-19 pandemic in which large-scale workplace change was a key feature for most people. Moreover, measures of employee well-being are increasingly seen as a key data point in defining firms' environmental, social, and governance (ESG) impact.²⁸ But

²⁸See Allas and Schaninger (2020) for a detailed discussion of ESG goals and, in particular, the ways in which con-

do organizations have any incentives to invest in managerial and organizational practices that might improve the experience of work?

Workplace happiness is currently low in the USA, with many people employed in relatively miserable workplaces across the country. The evidence presented in this paper suggests that improving the happiness of these workers is not only in their own interests, but may also be in the broader interests of their employers.

References

- Allas, Tera and Bill Schaninger**, “The boss factor: Making the world a better place through workplace relationships,” *The McKinsey Quarterly*, 2020.
- Altmann, Steffen, Armin Falk, Simon Jäger, and Florian Zimmermann**, “Learning about job search: A field experiment with job seekers in Germany,” *Journal of Public Economics*, 2018, 164, 33–49.
- Amabile, Teresa M, Sigal G Barsade, Jennifer S Mueller, and Barry M Staw**, “Affect and creativity at work,” *Administrative Science Quarterly*, 2005, 50 (3), 367–403.
- Ashraf, Nava, Oriana Bandiera, Edward Davenport, and Scott S Lee**, “Losing prosociality in the quest for talent? Sorting, selection, and productivity in the delivery of public services,” *American Economic Review*, 2020, 110 (5), 1355–94.
- Baker, Scott R and Andrey Fradkin**, “The impact of unemployment insurance on job search: Evidence from Google search data,” *Review of Economics and Statistics*, 2017, 99 (5), 756–768.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul**, “Social incentives in the workplace,” *Review of Economic Studies*, 2010, 77 (2), 417–458.
- Barsade, Sigal G and Donald E Gibson**, “Why does affect matter in organizations?,” *Academy of Management Perspectives*, 2007, 21 (1), 36–59.
- Baumeister, Roy F, Ellen Bratslavsky, Catrin Finkenauer, and Kathleen D Vohs**, “Bad is stronger than good,” *Review of general psychology*, 2001, 5 (4), 323.
- Bellet, Clement, Jan-Emmanuel De Neve, and George Ward**, “Does Employee Happiness Have an Impact on Productivity?,” *Working Paper*, 2021.
- Belot, Michele, Philipp Kircher, and Paul Muller**, “Providing advice to jobseekers at low cost: An experimental study on online advice,” *Review of Economic Studies*, 2019, 86 (4), 1411–1447.
- Benson, A., A. Sojourner, and A. Umyarov**, “The value of employer reputation in the absence of contract enforcement: a randomized experiment. Work. Pap,” 2017.
- Benson, Alan, Aaron Sojourner, and Akhmed Umyarov**, “Can Reputation Discipline the Gig Economy? Experimental Evidence from an Online Labor Market,” *Management Science*, 2020, 66 (5), 1802–1825.

sistently measured company-level happiness data—such as those studied in the present paper—can help to fill a large and important data gap related to social impact.

- Bhattacharya, Utpal, Craig W Holden, and Stacey Jacobsen**, “Penny wise, dollar foolish: Buy–sell imbalances on and around round numbers,” *Management Science*, 2012, 58 (2), 413–431.
- Bhole, Monica, Andrey Fradkin, and John Horton**, “Information About Vacancy Competition Redirects Job Search,” *SocArXiv*, 2021, p82fk.
- Bloom, Nicholas, James Liang, John Roberts, and Zhichun Jenny Ying**, “Does working from home work? Evidence from a Chinese experiment,” *Quarterly Journal of Economics*, 2014, 130 (1), 165–218.
- Breza, Emily, Supreet Kaur, and Yogita Shamdasani**, “The Morale Effects of Pay Inequality,” *Quarterly Journal of Economics*, 2018, 133 (2), 611–663.
- Brief, Arthur P and Howard M Weiss**, “Organizational behavior: Affect in the workplace,” *Annual review of psychology*, 2002, 53 (1), 279–307.
- Brown, Jennifer and David A Matsa**, “Boarding a sinking ship? An investigation of job applications to distressed firms,” *The Journal of Finance*, 2016, 71 (2), 507–550.
- Bryson, Alex and George MacKerron**, “Are you happy while you work?,” *Economic Journal*, 2016, 127 (599), 106–125.
- Burbano, Vanessa C**, “Social responsibility messages and worker wage requirements: Field experimental evidence from online labor marketplaces,” *Organization Science*, 2016, 27 (4), 1010–1028.
- Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez**, “Inequality at Work: The Effect of Peer Salaries on Job Satisfaction,” *American Economic Review*, 2012, 102 (6), 2981–3003.
- Carmichael, H Lorne**, “Reputations in the labor market,” *The American Economic Review*, 1984, 74 (4), 713–725.
- Carnahan, Seth, David Kryscynski, and Daniel Olson**, “When does corporate social responsibility reduce employee turnover? Evidence from attorneys before and after 9/11,” *Academy of Management Journal*, 2017, 60 (5), 1932–1962.
- Cassar, Lea and Stephan Meier**, “Nonmonetary Incentives and the Implications of Work as a Source of Meaning,” *Journal of Economic Perspectives*, 2018, 32 (3), 215–38.
- Cattaneo, Matias D, Brigham R Frandsen, and Rocio Titiunik**, “Randomization inference in the regression discontinuity design: An application to party advantages in the US Senate,” *Journal of Causal Inference*, 2015, 3 (1), 1–24.
- , **Rocio Titiunik, and Gonzalo Vazquez-Bare**, “Comparing inference approaches for RD designs: A reexamination of the effect of head start on child mortality,” *Journal of Policy Analysis and Management*, 2017, 36 (3), 643–681.
- Chadi, Adrian, Sabrina Jeworrek, and Vanessa Mertins**, “When the meaning of work has disappeared: experimental evidence on employees’ performance and emotions,” *Management Science*, 2017, 63 (6), 1696–1707.

- Chen, Yan, F Maxwell Harper, Joseph Konstan, and Sherry Xin Li**, “Social comparisons and contributions to online communities: A field experiment on MovieLens,” *American Economic Review*, 2010, 100 (4), 1358–98.
- Chevalier, Judith A and Dina Mayzlin**, “The effect of word of mouth on sales: Online book reviews,” *Journal of marketing research*, 2006, 43 (3), 345–354.
- Clark, Andrew E**, “What really matters in a job? Hedonic measurement using quit data,” *Labour Economics*, 2001, 8 (2), 223–242.
- , “Work, jobs, and well-being across the millennium,” *International differences in well-being*, 2010, pp. 436–468.
- , “Four decades of the economics of happiness: Where next?,” *Review of Income and Wealth*, 2018, 64 (2), 245–269.
- Coffman, Lucas C, Clayton R Featherstone, and Judd B Kessler**, “Can social information affect what job you choose and keep?,” *American Economic Journal: Applied Economics*, 2017, 9 (1), 96–117.
- Cullen, Zoë and Ricardo Perez-Truglia**, “How much does your boss make? the effects of salary comparisons,” *NBER Working Paper 24841*, 2018.
- Dellarocas, Chrysanthos**, “The digitization of word of mouth: Promise and challenges of online feedback mechanisms,” *Management science*, 2003, 49 (10), 1407–1424.
- Diener, E., S.J. Heintzelman, K. Kushlev, L. Tay, D. Wirtz, L.D. Lutes, and S. Oishi**, “Findings All Psychologists Should Know From the New Science on Subjective Well-Being,” *Canadian Psychology/Psychologie canadienne*, 2017, 58, 87–104.
- Diener, Ed, Eunkook M Suh, Richard E Lucas, and Heidi L Smith**, “Subjective well-being: Three decades of progress,” *Psychological bulletin*, 1999, 125 (2), 276.
- Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Brücker**, “Referral-based job search networks,” *The Review of Economic Studies*, 2016, 83 (2), 514–546.
- Edmans, Alex**, “Does the stock market fully value intangibles? Employee satisfaction and equity prices,” *Journal of Financial economics*, 2011, 101 (3), 621–640.
- , “The link between job satisfaction and firm value, with implications for corporate social responsibility,” *Academy of Management Perspectives*, 2012, 26 (4), 1–19.
- Estrada, Carlos A, Alice M Isen, and Mark J Young**, “Positive affect facilitates integration of information and decreases anchoring in reasoning among physicians,” *Organizational behavior and human decision processes*, 1997, 72 (1), 117–135.
- Faberman, R Jason and Marianna Kudlyak**, “The intensity of job search and search duration,” *American Economic Journal: Macroeconomics*, 2019, 11 (3), 327–57.
- Fernandez, Roberto M, Emilio J Castilla, and Paul Moore**, “Social capital at work: Networks and employment at a phone center,” *American journal of sociology*, 2000, 105 (5), 1288–1356.
- Filippas, Apostolos, John J. Horton, and Joseph M. Golden§**, “Reputation Inflation,” *Working Paper*, 2020.

- Fisher, V.E. and J.V. Hanna**, *The dissatisfied worker*, New York: Macmillan, 1931.
- Flory, Jeffrey A, Andreas Leibbrandt, and John A List**, “Do competitive workplaces deter female workers? A large-scale natural field experiment on job entry decisions,” *Review of Economic Studies*, 2015, 82 (1), 122–155.
- Frey, Bruno S and Stephan Meier**, “Social comparisons and pro-social behavior: Testing” conditional cooperation” in a field experiment,” *American Economic Review*, 2004, 94 (5), 1717–1722.
- Gartenberg, Claudine, Andrea Prat, and George Serafeim**, “Corporate purpose and financial performance,” *Organization Science*, 2019, 30 (1), 1–18.
- Gee, Laura K**, “The more you know: information effects on job application rates in a large field experiment,” *Management Science*, 2019, 65 (5), 2077–2094.
- **and Michael J Schreck**, “Do beliefs about peers matter for donation matching? Experiments in the field and laboratory,” *Games and Economic Behavior*, 2018, 107, 282–297.
- Gosnell, Greer K, John A List, and Robert D Metcalfe**, “The impact of management practices on employee productivity: A field experiment with airline captains,” *Journal of Political Economy*, 2020, 128 (4), 1195–1233.
- Graham, Carol, Kate Laffan, and Sergio Pinto**, “Well-being in metrics and policy,” *Science*, 2018, 362 (6412), 287–288.
- Green, Francis**, “Well-being, job satisfaction and labour mobility,” *Labour Economics*, 2010, 17 (6), 897–903.
- Harrell, Frank E**, *Regression modeling strategies: with applications to linear models, logistic and ordinal regression, and survival analysis*, Springer, 2015.
- HBR Analytical Services**, “Cultivating Workforce Well-Being to Drive Business Value,” <https://bg.hbr.org/resources/pdfs/comm/workplacewellbeing.pdf>, 2020.
- Helmers, Christian, Pramila Krishnan, and Manasa Patnam**, “Attention and saliency on the internet: Evidence from an online recommendation system,” *Journal of Economic Behavior & Organization*, 2019, 161, 216–242.
- Hersey, R. B.**, “Rate of Production and Emotional State,” *Personnel Journal*, 1932, pp. 335–364.
- Hersey, R.B.**, *Workers’ emotions in shop and home: A study of industrial workers from the psychological and physiological standpoint*, Philadelphia, PA: University of Pennsylvania Press, 1932.
- Horton, John J and Prasanna Tambe**, “Labor economists get their microscope: Big data and labor market analysis,” *Big data*, 2015, 3 (3), 130–137.
- Houser, J.D.**, *What the employer thinks: Executives’ attitudes toward employees*, Cambridge, MA: Harvard University Press, 1927.
- Iaffaldano, Michelle T and Paul M Muchinsky**, “Job satisfaction and job performance: A meta-analysis,” *Psychological Bulletin*, 1985, 97 (2), 251.
- Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer**, “Worker Beliefs About Rents and Outside Options,” *Working Paper*, 2021.

- Jencks, Christopher, Lauri Perman, and Lee Rainwater**, “What is a good job? A new measure of labor-market success,” *American journal of sociology*, 1988, 93 (6), 1322–1357.
- Judge, Timothy A, Carl J Thoresen, Joyce E Bono, and Gregory K Patton**, “The job satisfaction–job performance relationship: A qualitative and quantitative review,” *Psychological bulletin*, 2001, 127 (3), 376.
- Kahneman, Daniel, Peter Wakker, and Rakesh Sarin**, “Back to Bentham? Explorations of Experienced Utility,” *Quarterly Journal of Economics*, 1997, 112 (2), 375–406.
- Kornhauser, A. and A. Sharp**, “Employee attitudes: Suggestions from a study in a factory,” *Personnel Journal*, 1932, 10, 3934, 31.
- Krekel, Christian, George Ward, and Jan-Emmanuel De Neve**, “What makes for a good job? Evidence using subjective wellbeing data,” in Mariano Rojas, ed., *The Economics of Happiness*, Springer, 2019, pp. 241–268.
- Krueger, A. B. and D. A. Schkade**, “The reliability of subjective well-being measures,” *Journal of Public Economics*, 2008, 92, 1833–1845.
- Krueger, Alan B and Andreas I Mueller**, “A contribution to the empirics of reservation wages,” *American Economic Journal: Economic Policy*, 2016, 8 (1), 142–79.
- Krueger, Alan B. and Arthur A. Stone**, “Progress in Measuring Subjective Well-Being: moving toward national indicators and policy evaluations,” *Science*, 2014, 346(6205), 42–43.
- Krueger, Alan B, Daniel Kahneman, David Schkade, Norbert Schwarz, and Arthur A Stone**, “National time accounting: The currency of life,” in “Measuring the subjective well-being of nations: National accounts of time use and well-being,” University of Chicago Press, 2009, pp. 9–86.
- Lacetera, Nicola, Devin G Pope, and Justin R Sydnor**, “Heuristic thinking and limited attention in the car market,” *American Economic Review*, 2012, 102 (5), 2206–36.
- Landy, F.W.**, *The psychology of work behavior*, 3rd ed., Homewood, IL: Dorsey Press, 1985.
- Lawler, Edward E**, *Motivation in work organizations.*, Monteray, CA: Brooks/Cole Publishing, 1973.
- Lazear, Edward P and Kathryn L Shaw**, “Personnel economics: The economist’s view of human resources,” *Journal of Economic Perspectives*, 2007, 21 (4), 91–114.
- Levy-Garboua, L., C. Montmarquette, and V. Simonnet**, “Job satisfaction and quits,” *Labour Economics*, 2007, 14, 2–251.
- List, John A**, “Field experiments: a bridge between lab and naturally occurring data,” *The BE Journal of Economic Analysis & Policy*, 2007, 5 (2).
- , “Non est Disputandum de Generalizability? A Glimpse into The External Validity Trial,” *NBER Working Paper 27535*, 2020.
- **and Fatemeh Momeni**, “When corporate social responsibility backfires: Evidence from a natural field experiment,” *Management Science*, 2021, 67 (1), 8–21.

- Luca, Michael**, “Reviews, reputation, and revenue: The case of Yelp.com,” *Harvard Business School Working Paper*, 2016.
- Maestas, Nicole, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger**, “The value of working conditions in the united states and implications for the structure of wages,” *NBER Working Paper 25204*, 2018.
- Marinescu, Ioana and Daphné Skandalis**, “Unemployment insurance and job search behavior,” *The Quarterly Journal of Economics*, 2021, 136 (2), 887–931.
- Marinescu, Ioana E**, “Job search monitoring and assistance for the unemployed,” *IZA World of Labor*, 2017.
- Mas, Alexandre and Amanda Pallais**, “Valuing Alternative Work Arrangements,” *American Economic Review*, 2017, 107 (12), 3722–3759.
- Moen, Phyllis, Erin L Kelly, Wen Fan, Shi-Rong Lee, David Almeida, Ellen Ernst Kossek, and Orfeu M Buxton**, “Does a flexibility/support organizational initiative improve high-tech employees’ well-being? Evidence from the work, family, and health network,” *American Sociological Review*, 2016, 81 (1), 134–164.
- Mortensen, Dale T**, “Job search and labor market analysis,” *Handbook of labor economics*, 1986, 2 (15), 849–919.
- Myers, D.G., D.G. Myers, and E. Diener**, “The pursuit of happiness,” *Psychological Science*, 1993, 6, 1 19.
- Nagaraj, Abhishek and Henning Piezunka**, “How Competition Affects Contributions to Open Source Platforms: Evidence from OpenStreetMap and Google Maps,” *Working Paper*, 2020.
- Oswald, Andrew J, Eugenio Proto, and Daniel Sgroi**, “Happiness and productivity,” *Journal of Labor Economics*, 2015, 33 (4), 789–822.
- Reimers, Imke and Joel Waldfogel**, “Digitization and pre-purchase information: the causal and welfare impacts of reviews and crowd ratings,” *American Economic Review*, 2021, 111 (6), 1944–71.
- Reynolds, Lloyd George**, *The structure of labor markets: Wages and labor mobility in theory and practice*, Westport, Conn: Greenwood Press, 1951.
- Rosen, Sherwin**, “The theory of equalizing differences,” *Handbook of labor economics*, 1986, 1, 641–692.
- Rothbard, Nancy P and Steffanie L Wilk**, “Waking up on the right or wrong side of the bed: Start-of-workday mood, work events, employee affect, and performance,” *Academy of Management Journal*, 2011, 54 (5), 959–980.
- Sockin, Jason, Aaron Sojourner, and Evan Starr**, “Non-Disclosure Agreements and Externalities from Silence,” *SSRN Working Paper 3900285*, 2021.
- and —, “What’s the Inside Scoop? Addressing Workers’ Information Problems with Firm Reputation,” *Addressing Workers’ Information Problems with Firm Reputation (May 4, 2020)*, 2020.

- Staw, Barry M and Sigal G Barsade**, "Affect and managerial performance: A test of the sadder-but-wiser vs. happier-and-smarter hypotheses," *Administrative Science Quarterly*, 1993, pp. 304–331.
- Stern, Scott**, "Do scientists pay to be scientists?," *Management Science*, 2004, 50 (6), 835–853.
- Stinchcombe, Arthur L.**, "Social structure and organizations," in J.G. March, ed., *Handbook of Organizations*, 1965, pp. 142–193.
- Tenney, Elizabeth R, Jared M Poole, and Ed Diener**, "Does positivity enhance work performance?: Why, when, and what we don't know," *Research in Organizational Behavior*, 2016, 36, 27–46.
- Tucker, Catherine and Juanjuan Zhang**, "How does popularity information affect choices? A field experiment," *Management Science*, 2011, 57 (5), 828–842.
- Walsh, Lisa C, Julia K Boehm, and Sonja Lyubomirsky**, "Does happiness promote career success? Revisiting the evidence," *Journal of Career Assessment*, 2018, 26 (2), 199–219.
- Wiswall, Matthew and Basit Zafar**, "Preference for the workplace, investment in human capital, and gender," *The Quarterly Journal of Economics*, 2018, 133 (1), 457–507.
- Wright, T.A. and D.G. Bonett**, "Job satisfaction and psychological well-being as nonadditive predictors of workplace turnover," *Journal of Management*, 2007, 33, 141–160.
- Wright, Thomas A and Russell Cropanzano**, "The role of psychological well-being in job performance:: a fresh look at an age-old quest," *Organizational dynamics*, 2004, 33 (4), 338–351.

Supplementary Materials

For Online Publication Only

Appendix A Balance and Summary Statistics

Table S1: Balance Between Control and Treatment job seekers

Variable	(1) Control		(2) Treated		(3) Total		T-test P-value (1)-(2)
	N	Mean/SD	N	Mean/SD	N	Mean/SD	
Cookie Age	1167000	125.229 (278.498)	22204762	125.512 (278.638)	23371762	125.498 (278.631)	0.285
Registered User	1167000	0.472 (0.497)	22204762	0.473 (0.497)	23371762	0.473 (0.497)	0.113
Desktop User	1167000	0.626 (0.484)	22204762	0.625 (0.484)	23371762	0.625 (0.484)	0.057
CZ: Unemployment Rate	1129172	0.083 (0.033)	21487432	0.083 (0.033)	22616604	0.083 (0.033)	0.344
User Has Resume	1167000	0.255 (0.436)	22204762	0.256 (0.436)	23371762	0.256 (0.436)	0.239
Number of Jobs on Resume	298067	3.534 (3.041)	5682210	3.533 (3.037)	5980277	3.533 (3.037)	0.810
Total Experience (months)	298067	95.405 (101.438)	5682210	95.123 (101.302)	5980277	95.137 (101.309)	0.137
Year Entered Labor Force	224743	2009.480 (8.668)	4282258	2009.508 (8.643)	4507001	2009.507 (8.645)	0.122
User is Employed	298067	0.473 (0.499)	5682210	0.473 (0.499)	5980277	0.473 (0.499)	0.967
User Wants Full-Time Job	301891	0.518 (0.500)	5755899	0.518 (0.500)	6057790	0.518 (0.500)	0.448
User has BA or higher	469519	0.402 (0.490)	8951705	0.401 (0.490)	9421224	0.401 (0.490)	0.412

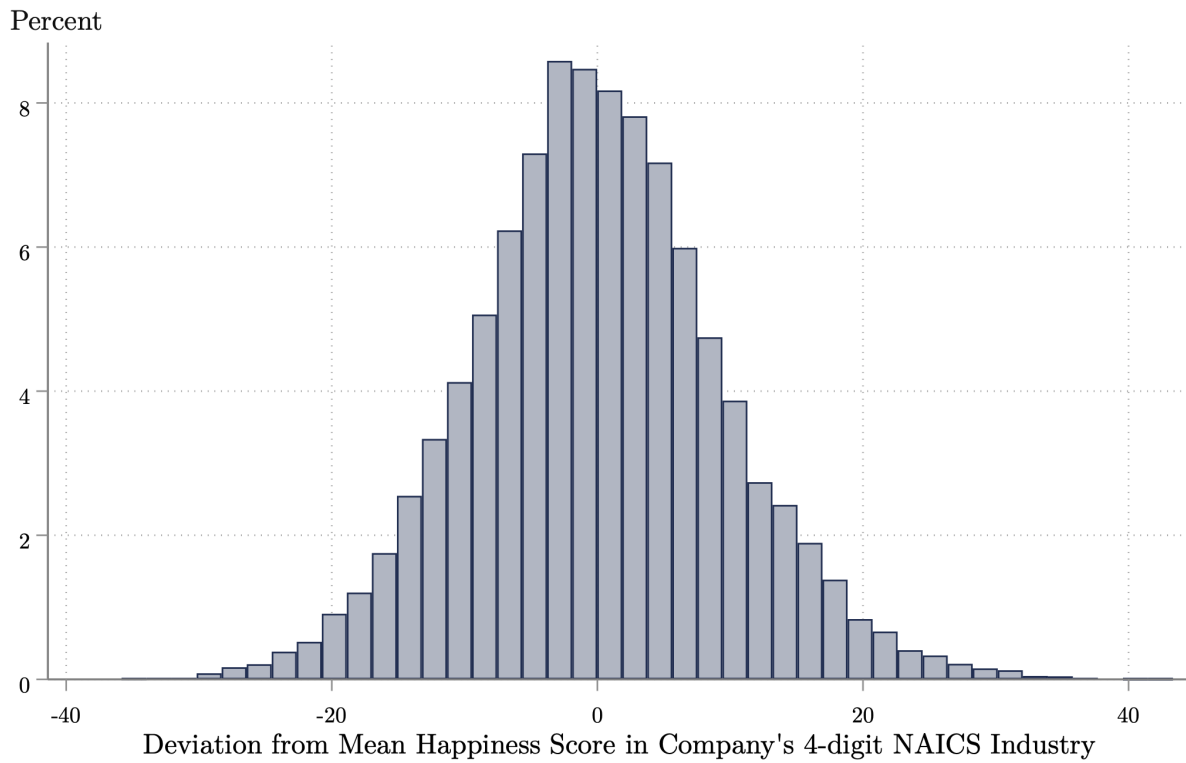
Notes: Each observation is a job seeker.

Table S2: Balance Between Control and Treatment Jobseeker–Company Observations

Variable	(1) Control		(2) Treated		(3) Total		T-test P-value (1)-(2)
	N	Mean/SD	N	Mean/SD	N	Mean/SD	
Applied = 100	1865835	20.049 (40.037)	35507316	19.775 (39.830)	37373151	19.788 (39.840)	0.000
Happiness Score	1865835	62.769 (8.836)	35507316	62.764 (8.833)	37373151	62.764 (8.833)	0.491
Number of Happiness Surveys	1865835	1278.794 (3723.817)	35507316	1281.649 (3724.995)	37373151	1281.506 (3724.936)	0.308
Company is Fortune 500	1865104	0.157 (0.364)	35493558	0.157 (0.364)	37358662	0.157 (0.364)	0.776
Company is Staffing Agency	1865104	0.014 (0.115)	35493558	0.013 (0.115)	37358662	0.013 (0.115)	0.925
Number of Jobs Listed	1865098	6068.453 (32154.709)	35493464	6084.784 (32099.241)	37358562	6083.969 (32102.012)	0.498
Company Star Rating	1865026	3.521 (0.492)	35492273	3.520 (0.492)	37357299	3.520 (0.492)	0.177
Number of Reviews	1865026	6300.275 (19105.253)	35492273	6317.131 (19134.965)	37357299	6316.289 (19133.483)	0.241
Number of Employees	1865835	5847.607 (9663.622)	35507316	5845.556 (9659.707)	37373151	5845.659 (9659.902)	0.777

Notes: Each observation is a jobseeker–company pair on the first day that the job seeker visits that company’s page.

Figure S1: Company Happiness Score Variation Within Industries



Note: Plotted are the residuals from a regression of company-level happiness score on the final day of the experiment on a set of industry fixed effects.

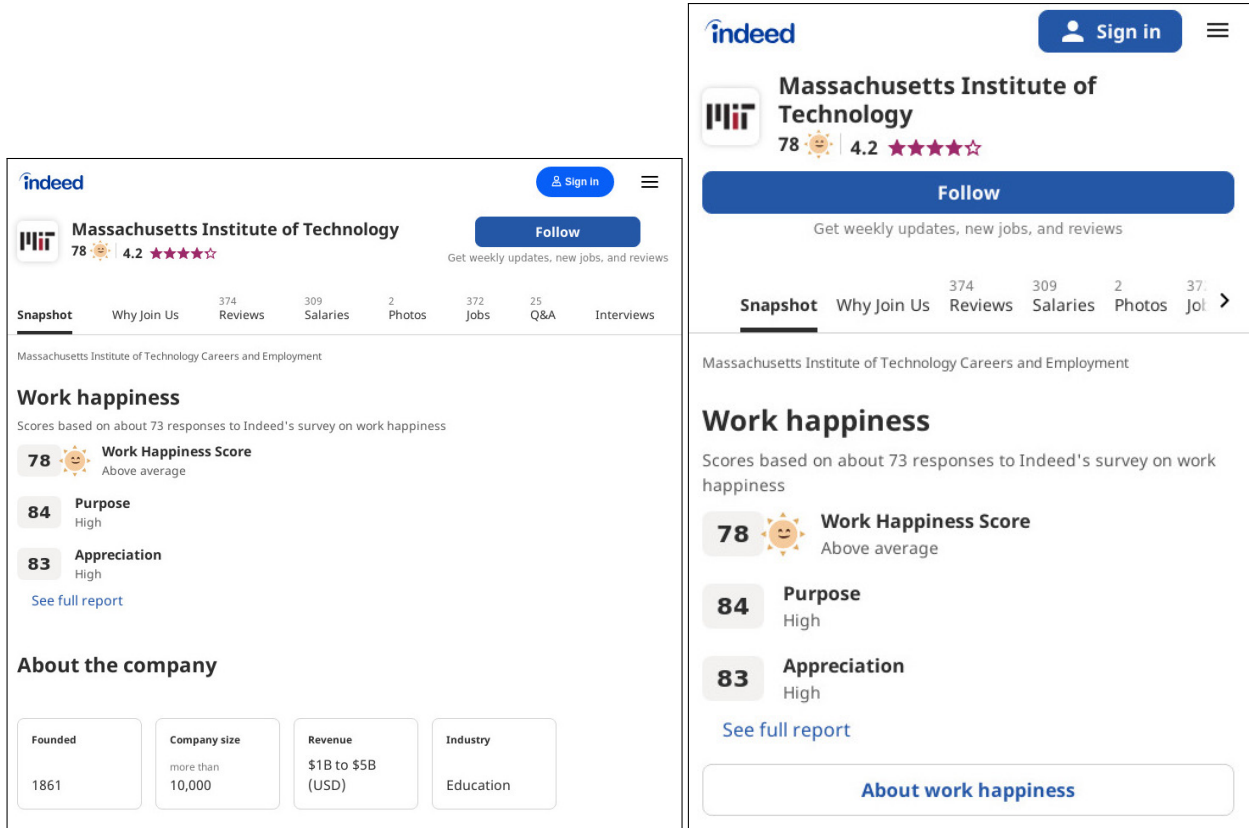
Appendix B Happiness Score: Survey Question Wordings

Users on the site are invited to answer questions on the workplace well-being of their workplace. The wording of these questions is as follows. See <https://www.indeed.com/about/happiness> for more details. All questions ask jobseekers the extent to which they agree with the statement from strongly disagree to strongly agree.

- *I feel happy at work most of the time.*
- *My work has a clear sense of purpose.*
- *I am paid fairly for my work.*
- *There are people at work who give me support and encouragement.*
- *There are people at work who appreciate me as a person.*
- *I can trust people in my company.*
- *I feel a sense of belonging in my company.*
- *My manager helps me succeed.*
- *My work environment feels inclusive and respectful of all people.*
- *My work has the time and location flexibility I need.*
- *In most of my work tasks, I feel energized.*
- *I am achieving most of my goals at work.*
- *I often learn something at work.*

Appendix C Treatment Conditions on Different Device Types

Figure S2: Treatment Screenshots on Other Devices



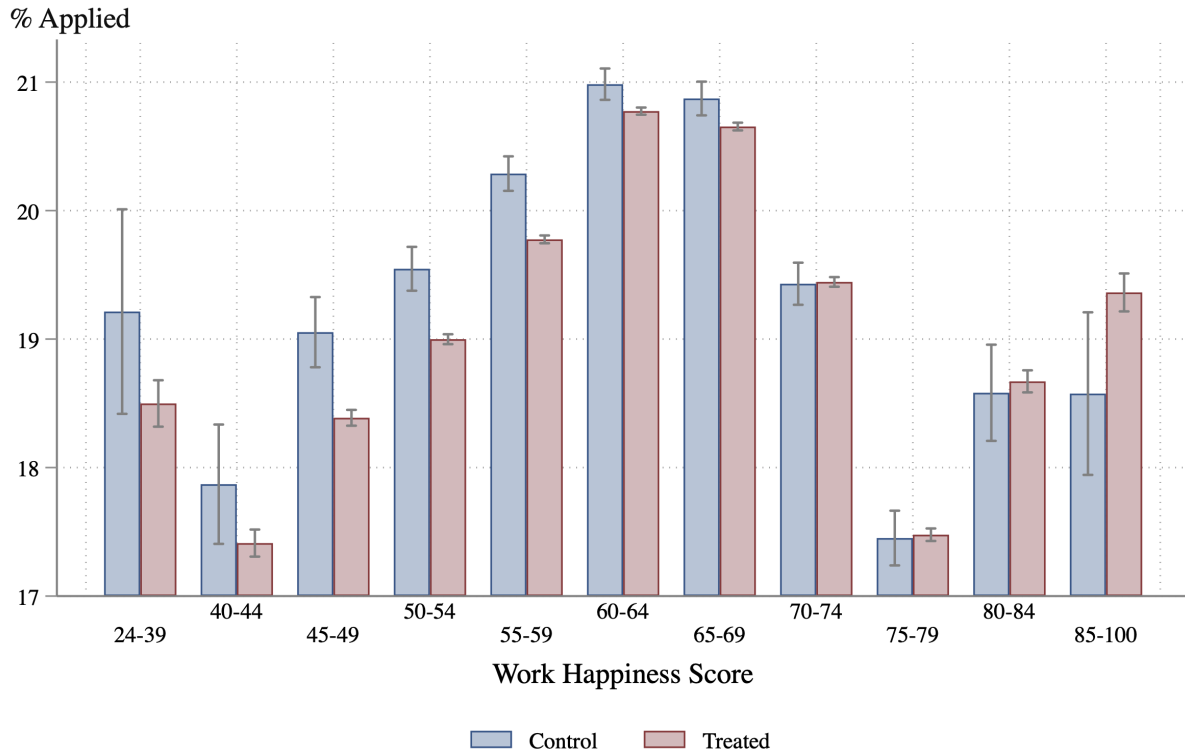
(a) Tablet

(b) Phone

Note: Screenshots in main text are for desktop computers. Shown here is what the treatment looks like on a mobile phone and on a tablet.

Appendix D Sensitivity to Alternative Specifications

Figure S3: Non-Regression Estimates



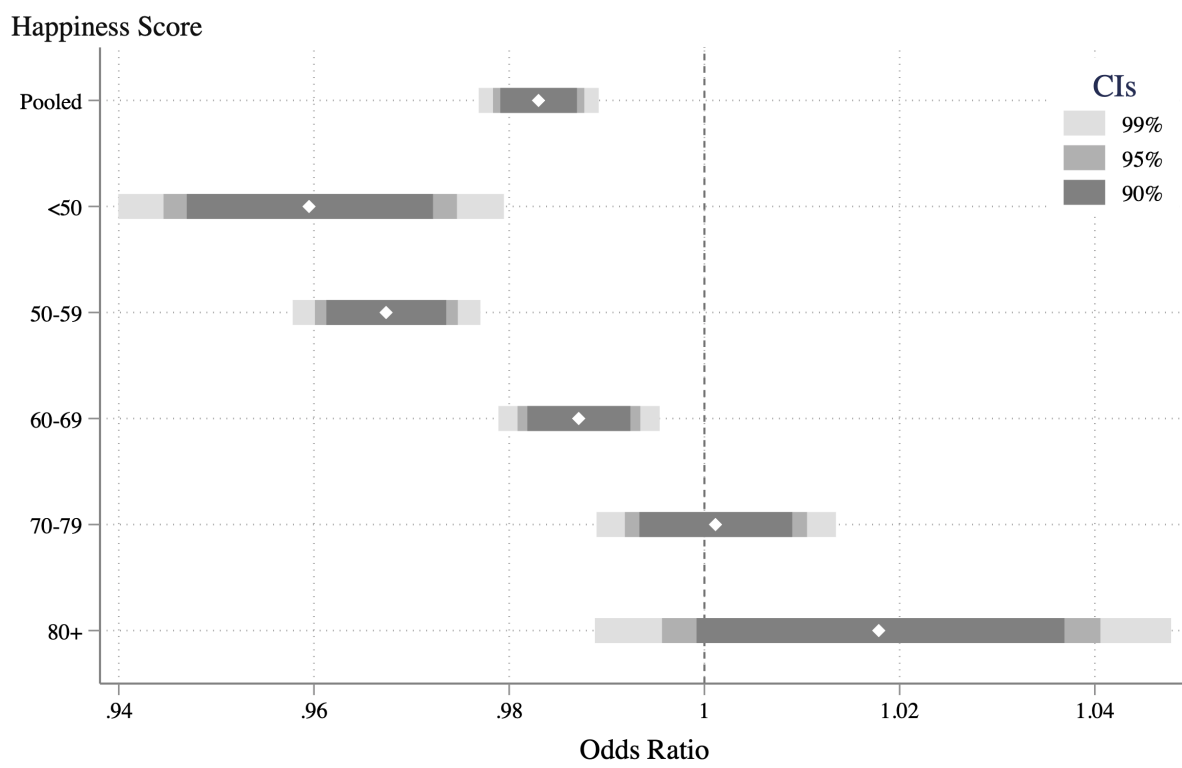
Note: Plotted are the simple means of $\text{applied} \times 100$ across treatment and control, at different levels of work happiness score displayed.

Table S3: Logistic Regression Estimates

	(1) All	(2) 20-49	(3) 50-59	(4) 60-69	(5) 70-79	(6) 80-100
Treated	-0.017*** (0.002)	-0.041*** (0.008)	-0.033*** (0.004)	-0.013*** (0.003)	0.001 (0.005)	0.018 (0.011)
Observations	37363079	2306602	11059770	15951318	6924676	1120713
Log-Likelihood	-18586738.3	-1094469.1	-5458399.7	-8140686.9	-3341673.1	-542246.8

Notes: Robust standard errors are in parentheses, adjusted for clustering on job seekers. Outcome in all models is equal to 1 if the job seeker applied to a job at that company, 0 otherwise. Logistic regression coefficients are reported. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure S4: Logistic Regression Estimates: Odds Ratios



Note: Plotted are the exponentiated coefficients reported in Table S3.

Table S4: Main Effect of Treatment: Sensitivity to Different Controls and Samples

(a) Sample restricted to first day per jobseeker–company pair

	Applied = 100							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treated	-0.273*** (0.039)	-0.273*** (0.039)	-0.277*** (0.038)	-0.260*** (0.038)	-0.278*** (0.038)	-0.280*** (0.038)	-0.283*** (0.037)	-0.269*** (0.037)
N	37,309,899	37,309,899	37,309,899	37,309,899	37,309,899	37,309,899	37,309,899	37,309,899
Date FEs		✓	✓			✓	✓	
Company FEs			✓				✓	
Company-Date FEs				✓				✓
User Controls					✓	✓	✓	✓

(b) Whole Sample: all jobseeker-company-days observed


	Applied = 100							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treated	-0.280*** (0.041)	-0.281*** (0.041)	-0.285*** (0.040)	-0.267*** (0.040)	-0.285*** (0.040)	-0.287*** (0.039)	-0.292*** (0.039)	-0.276*** (0.039)
Observations	39,712,785	39,712,785	39,712,785	39,712,785	39,712,785	39,712,785	39,712,785	39,712,785
Date FEs		✓	✓			✓	✓	
Company FEs			✓				✓	
Company-Date FEs				✓				✓
User Controls					✓	✓	✓	✓

(c) Sample restricted to first company-day pair per job seeker


	Applied = 100							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treated	-0.266*** (0.038)	-0.269*** (0.038)	-0.261*** (0.038)	-0.259*** (0.040)	-0.270*** (0.038)	-0.272*** (0.038)	-0.265*** (0.037)	-0.266*** (0.040)
N	22,997,526	22,997,526	22,997,526	22,997,526	22,997,525	22,997,525	22,997,525	22,997,525
Date FEs		✓	✓			✓	✓	
Company FEs			✓				✓	
Company-Date FEs				✓				✓
User Controls					✓	✓	✓	✓

Notes: Robust standard errors are in parentheses in Panel (c); robust standard errors adjusted for clustering on jobseekers are in parentheses in Panels (a) and (b). Outcome in all models is equal to 100 if the job seeker applied to a job at that company, 0 otherwise. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


Figure S5: Alternative Outcomes



(a) Outcome: Job Clicks



(b) Outcome: Followed Company = 1



(c) Outcome: Looked at Reviews Tab = 1

Note: panel (a) is based on a series of Poisson regressions, with the sample split according to the happiness score of the company in question; panels (b) and (c) are based on a series of LPMs. Percent effects are calculated using the control group mean of $\Pr(\text{Apply})$ in panels (b) and (c), and in panel (a) are based on (the exponent of) the coefficients from the Poisson models. Company-by-day fixed effects included in all models. Standard errors are adjusted for clustering on job seekers.

Appendix E Replication Studies in Canada and the UK

Table S5: Effect of Showing Happiness Score on Application Behavior in the United Kingdom

	Applied = 100			
	(1)	(2)	(3)	(4)
Main Effect				
Treated	-0.603*** (0.069)	-0.570*** (0.068)	-0.570*** (0.068)	-0.599*** (0.101)
Interactions: Treated				
× Happiness (z-score)			0.160*** (0.059)	
× score is 50-59				-0.206 (0.137)
× score is 60-69				0.406** (0.168)
× score is 70-79				0.197 (0.238)
× score is 80-100				1.197** (0.498)
Observations	2,496,398	2,493,551	2,493,551	2,493,551
User Controls		✓	✓	✓
Company-by-Date FEs		✓	✓	✓


Notes: Robust standard errors are in parentheses, adjusted for clustering on individuals. A linear probability model is estimated in which the outcome is an indicator variable equal to 100 if the job seeker applies to that company. jobseeker controls: logged in job seeker, desktop job seeker, cookie age. In column (4), the omitted happiness score category is 20-49. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table S6: Effect of Showing Happiness Score on Application Behavior in Canada


	Applied = 100			
	(1)	(2)	(3)	(4)
Main Effect				
Treated	-0.851*** (0.072)	-0.701*** (0.070)	-0.701*** (0.070)	-0.741*** (0.158)
Interactions: Treated				
× Happiness (z-score)			0.204*** (0.061)	
× score is 50-59				-0.286 (0.189)
× score is 60-69				0.050 (0.192)
× score is 70-79				0.746*** (0.222)
× score is 80-100				0.677* (0.376)
Observations	2,037,505	2,036,455	2,036,455	2,036,455
User Controls		✓	✓	✓
Company-by-Date FEs		✓	✓	✓

Notes: Robust standard errors are in parentheses, adjusted for clustering on individuals. A linear probability model is estimated in which the outcome is an indicator variable equal to 100 if the job seeker applies to that company. jobseeker controls: logged in job seeker, desktop job seeker, cookie age. In column (4), the omitted happiness score category is 20-49. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure S6: Replication Studies: Split-Sample Percent Effects of Showing Happiness Score



(a) United Kingdom

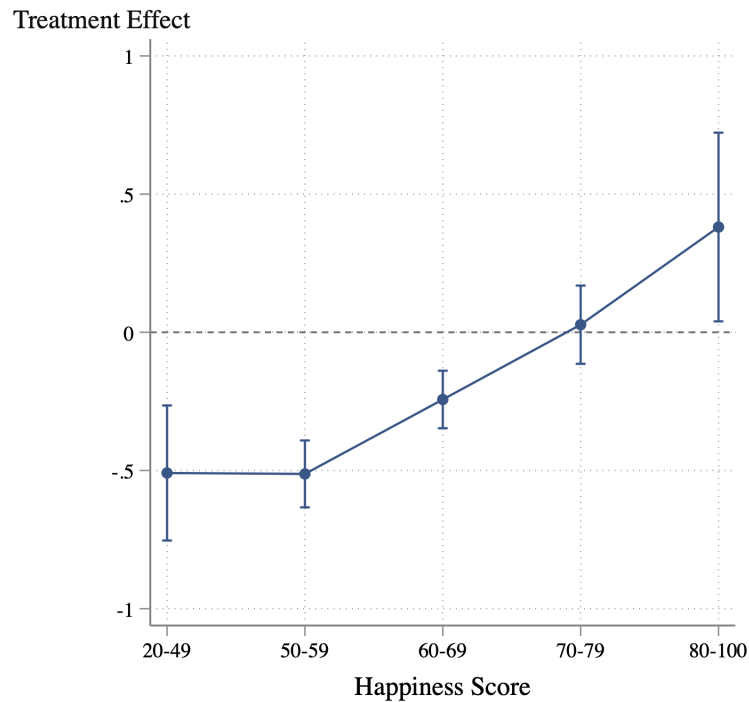


(b) Canada

Note: Percent effects are calculated using the control group mean of $Pr(\text{Apply})$ in each model. Results are derived from separate LPMs, with the sample split according to the happiness score of the company in question. Controls are included in all models for jobseeker characteristics as well as fixed effects for commuting zone and company-by-day. Standard errors are adjusted for clustering on job seekers.

Appendix F Treatment Effect Heterogeneity

Figure S7: Implied Treatment Effects from Interaction Model with Bins of Happiness Score Displayed



Note: Figure plots the implied treatment effects from the interaction models reported in column (4) of Table 1. In that model, the outcome is equal to 100 if the job seeker applies, and zero otherwise. The model includes controls for jobseeker characteristics, and a set of company-by-date fixed effects. The model includes a dummy for treatment, which is interacted with a series of dummies of the binned happiness score shown to job seekers. This figure plots the linear combination of treated and the interaction effect for each level of happiness, along with 95% confidence intervals.

Table S7: Treatment Effect By Happiness Score of the Company Visited

(a) Sample restricted to first day per jobseeker–company pair

	(1)	(2)	(3)	(4)	(5)
	20-49	50-59	60-69	70-79	80-100
Treated	-0.506*** (0.127)	-0.515*** (0.062)	-0.242*** (0.053)	0.026 (0.072)	0.385** (0.176)
Observations	2,299,865	11,042,034	15,933,238	6,916,461	1,118,283
Control Mean	18.81	20.01	20.94	18.74	18.59

(b) Whole Sample: all jobseeker-company-days observed

	(1)	(2)	(3)	(4)	(5)
	20-49	50-59	60-69	70-79	80-100
Treated	-0.536*** (0.126)	-0.489*** (0.063)	-0.256*** (0.054)	0.004 (0.072)	0.337* (0.174)
Observations	2,456,781	11,776,805	16,962,276	7,336,604	1,180,299
Control Mean	18.60	19.91	20.92	18.83	18.56

(c) Sample restricted to first company-day pair per job seeker

	(1)	(2)	(3)	(4)	(5)
	20-49	50-59	60-69	70-79	80-100
Treated	-0.598*** (0.163)	-0.472*** (0.075)	-0.243*** (0.061)	-0.015 (0.088)	0.413* (0.218)
Observations	1,396,195	6,700,754	9,848,111	4,321,047	731,390
Control Mean	18.81	20.11	20.97	18.51	18.35

Notes: Robust standard errors are in parentheses in Panel (c); robust standard errors adjusted for clustering on jobseekers are in parentheses in Panels (a) and (b). Outcome in all models is equal to 100 if the job seeker applied to a job at that company, 0 otherwise. Linear probability models reported. All models include company-day and commuting zone fixed effects, as well as controls for account age, whether logged in or not, and whether on desktop computer or not. Each column represents a separate model, with the sample determined by the Work Happiness Score of the company the job seeker was looking at. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table S8: Treatment Effect by Well-Being Scores

	Applied = 100	
	(1)	(2)
Main Effect		
Treated	-0.262*** (0.038)	-0.260*** (0.038)
× Happiness	0.022*** (0.004)	0.048*** (0.017)
× Achievement		-0.017 (0.013)
× Appreciation		-0.020 (0.020)
× Belonging		-0.006 (0.021)
× Energized		0.006 (0.016)
× Flexibility		0.013 (0.008)
× Inclusivity		-0.009 (0.016)
× Learning		0.022* (0.012)
× Management		-0.003 (0.014)
× Fair Pay		-0.006 (0.006)
× Purpose		-0.005 (0.012)
× Support		-0.000 (0.021)
× Trust		-0.002 (0.020)
Observations	37,309,899	37,219,457

Notes: Robust standard errors are in parentheses, adjusted for clustering on job seekers. Outcome in all models is equal to 100 if the job seeker applied to a job at that company, 0 otherwise. Linear probability models reported. All models include company-by-day fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table S9: Treatment Effect For Registered and Non-Registered job seekers

	(1)	(2)	(3)	(4)	(5)
	< 50	50-59	60-69	70-79	80+
Treated	-0.219 (0.168)	-0.375** (0.081)	-0.288** (0.069)	0.125 (0.091)	0.374* (0.221)
Registered User	11.620*** (0.241)	11.332*** (0.118)	10.715*** (0.101)	10.208*** (0.140)	9.581*** (0.347)
× Registered	-0.485** (0.247)	-0.286** (0.121)	0.081 (0.104)	-0.193 (0.144)	0.031 (0.355)
Observations	2,283,640	10,961,636	15,807,291	6,870,993	1,110,181
Control Mean (Registered)	11.53	13.24	14.63	13.06	13.20
Control Mean (Non-Registered)	23.90	25.22	26.34	24.18	24.15

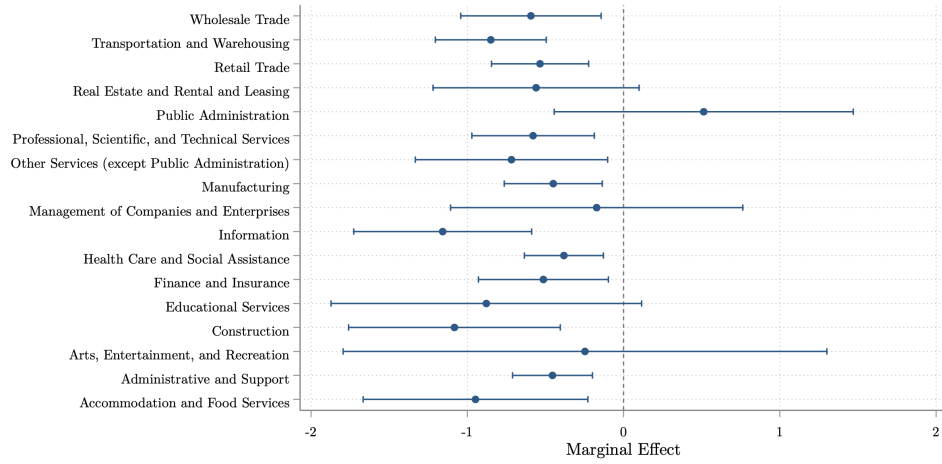
*Notes: Robust standard errors are in parentheses, adjusted for clustering on job seekers. Outcome in all models is equal to 100 if the job seeker applied to a job at that company, 0 otherwise. Linear probability models reported. All models include company-by-day fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

Table S10: Treatment Effect Heterogeneity: jobseeker and Company Characteristics


	< 50		50-59		60-69		70-79		80+	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Main Treatment Effect										
Treated	-0.880*** (0.251)	-2.382*** (0.915)	-0.658*** (0.128)	-0.500* (0.261)	-0.245** (0.115)	-0.232 (0.215)	-0.084 (0.159)	-0.440 (0.355)	0.512 (0.386)	0.857 (1.629)
Covariate Effects										
Cookie Age (z-score)	-0.507*** (0.051)	-0.452* (0.237)	-0.490*** (0.026)	-0.330*** (0.120)	-0.595*** (0.024)	-0.536*** (0.106)	-0.495*** (0.033)	-0.733*** (0.143)	-0.484*** (0.081)	-0.644* (0.351)
Employed User	0.346*** (0.112)	-0.006 (0.512)	0.220*** (0.057)	0.307 (0.266)	0.080 (0.052)	-0.243 (0.236)	0.529*** (0.073)	0.111 (0.326)	0.627*** (0.179)	0.749 (0.793)
User high Education (BA or more)	-2.933*** (0.114)	-2.305*** (0.497)	-2.754*** (0.058)	-2.531*** (0.255)	-2.209*** (0.054)	-2.164*** (0.228)	-0.019 (0.075)	-0.086 (0.315)	0.022 (0.184)	0.843 (0.767)
User Total Work Experience (z-score)	-0.942*** (0.055)	-0.712*** (0.244)	-1.063*** (0.028)	-1.074*** (0.127)	-1.346*** (0.026)	-1.239*** (0.116)	-0.908*** (0.037)	-0.861*** (0.165)	-0.954*** (0.095)	-1.052** (0.409)
Local Unemployment Rate (z-score)	0.246*** (0.084)	-0.010 (0.275)	0.431*** (0.040)	0.439*** (0.132)	0.550*** (0.035)	0.629*** (0.116)	0.653*** (0.049)	0.521*** (0.156)	0.553*** (0.118)	0.432 (0.370)
Num Happiness Surveys (z-score)	19.727*** (3.751)	22.835*** (4.370)	-1.029** (0.438)	-1.257*** (0.481)	0.200 (0.144)	0.193 (0.178)	4.865*** (0.800)	4.931*** (0.916)	0.523 (6.332)	-0.104 (7.306)
Large Company (10,000+ employees)	2.297** (0.942)	1.361 (1.206)	0.770 (0.477)	1.108** (0.540)	2.309*** (0.537)	2.593*** (0.576)	1.587** (0.716)	1.338* (0.781)	8.865*** (1.783)	8.774*** (2.091)
Company Star Rating (z-score)	0.133 (0.194)	0.712** (0.322)	0.091 (0.079)	0.145 (0.162)	-0.154*** (0.053)	0.046 (0.133)	0.073 (0.159)	0.122 (0.289)	0.088 (0.300)	-0.175 (0.683)
Company Num of Jobs Listed (z-score)	1.080** (0.535)	-0.656 (0.882)	1.093** (0.476)	1.757** (0.821)	3.395*** (0.174)	3.370*** (0.487)	0.311** (0.121)	0.222 (0.226)	10.312*** (0.425)	9.981*** (1.198)
Interactions: Treated										
× Cookie Age		-0.058 (0.242)		-0.168 (0.123)		-0.063 (0.109)		0.251* (0.147)		0.169 (0.361)
× Employed		0.370 (0.524)		-0.091 (0.272)		0.340 (0.242)		0.440 (0.334)		-0.129 (0.814)
× Education		-0.660 (0.508)		-0.234 (0.260)		-0.047 (0.233)		0.071 (0.322)		-0.865 (0.784)
× User Work Experience		-0.243 (0.250)		0.012 (0.130)		-0.112 (0.119)		-0.049 (0.169)		0.103 (0.420)
× Unemployment Rate		0.270 (0.276)		-0.008 (0.133)		-0.084 (0.117)		0.139 (0.156)		0.128 (0.369)
× Num Happiness Surveys		-3.267 (2.354)		0.242 (0.211)		0.007 (0.111)		-0.066 (0.473)		0.640 (3.758)
× Large Company		0.990 (0.798)		-0.354 (0.266)		-0.300 (0.219)		0.261 (0.329)		0.100 (1.141)
× Star Rating		-0.612** (0.274)		-0.056 (0.149)		-0.210 (0.129)		-0.052 (0.254)		0.275 (0.631)
× Num Jobs Listed		1.809** (0.755)		-0.696 (0.701)		0.026 (0.478)		0.093 (0.200)		0.352 (1.179)
Observations	638,671	638,671	2,971,360	2,971,360	3,900,874	3,900,874	1,661,821	1,661,821	260,162	260,162

Notes: Robust standard errors are in parentheses, adjusted for clustering on job seekers. Outcome in all models is equal to 100 if the job seeker applied to a job at that company, 0 otherwise. Linear probability models reported. All models include company fixed effects and day fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


Figure S8: Treatment Effect by Industry



(a) Low: Happiness Score < 60



(b) Average: Happiness Score 60 – 69

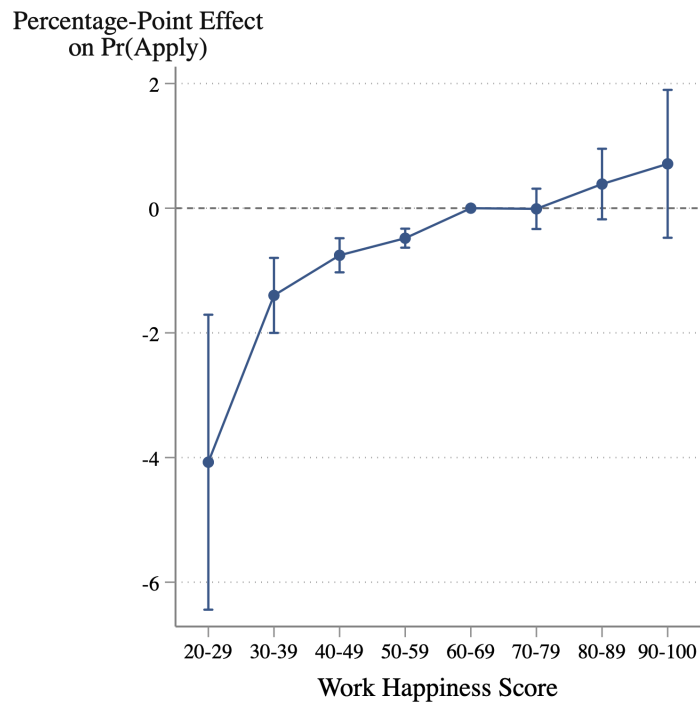


(c) High: Happiness Score 70+

Note: Each panel reports a regression with the sample split according to the happiness score shown. In each of the 3 regressions, the treatment indicator is interacted with industry fixed effects. Marginal effects are then calculated and reported, along with 95% confidence intervals.

Appendix G Effects of Happiness Score on Applications: Functional Form

Figure S9: Happiness Split into Categorical Variables



Note: The figure plots the coefficients and 95% confidence intervals from a regression of the probability of applying to a job on a series of indicators for the work happiness score displayed. 60–70 is the omitted category. The sample includes only treated job seekers. A linear probability model is estimated in which the outcome is an indicator variable equal to 100 if the job seeker applies to that company. The regression includes a full set of jobseeker and date fixed effects, as well as controls for company industry, company size, and indicators for whether the firm is a staffing agency and in the Fortune 500. Confidence intervals are derived from standard errors that are adjusted for two-way clustering on companies and job seekers.

Appendix H IV Estimates of Score Effects

In this section, I build on the fixed effects and RD estimates by making use of the fact that I can observe the data generating process of these scores. In doing so I am able to pursue an alternative identification strategy that leverages plausibly exogenous variation in the happiness score.

Appendix H.1 Identification Strategy

Happiness scores are calculated and shown when a company has 20 or more individual-level surveys. These surveys are filled in by job seekers, who arrive at the survey page from a variety of different sources. Some arrive directly, by going to their company's review page and clicking through to the survey. Others arrive after being directed to the survey from an email or mobile alert. Still others arrive there from the resume pages of the website. Jobseekers may upload their resume details to the website, and in doing so fill in their employment history. From here, the website directs these jobseekers to the happiness survey for each of the companies they have been employed by or are currently working for. As can be seen in Figure S12, those arriving from the resume pages report on average much higher happiness, on the 1-to-5 agreement scale. While those arriving to the survey from company pages answer on average just over 2.5, for those arriving from resume pages this is over 1 point higher.

Across companies, there is variation in the “make-up” of their happiness scores – in terms of where its individual-level surveys originate from. Among their respondents, some have a larger proportion of resume respondents than others, as can be seen in Figure S12. On average, around a quarter of responses come from the resume section, but this varies from zero up to around 60%. In order to identify the causal effect of the happiness score on applications, I make the (untestable) assumption that the percentage of a firm's responses coming from the resume section is as-good-as-random, controlling for observables like a company's industry and size. If this is the case, then the proportion of a score's responses arising from clicks from the resume section of the site will be a valid instrument for the score itself.

I estimate the equation

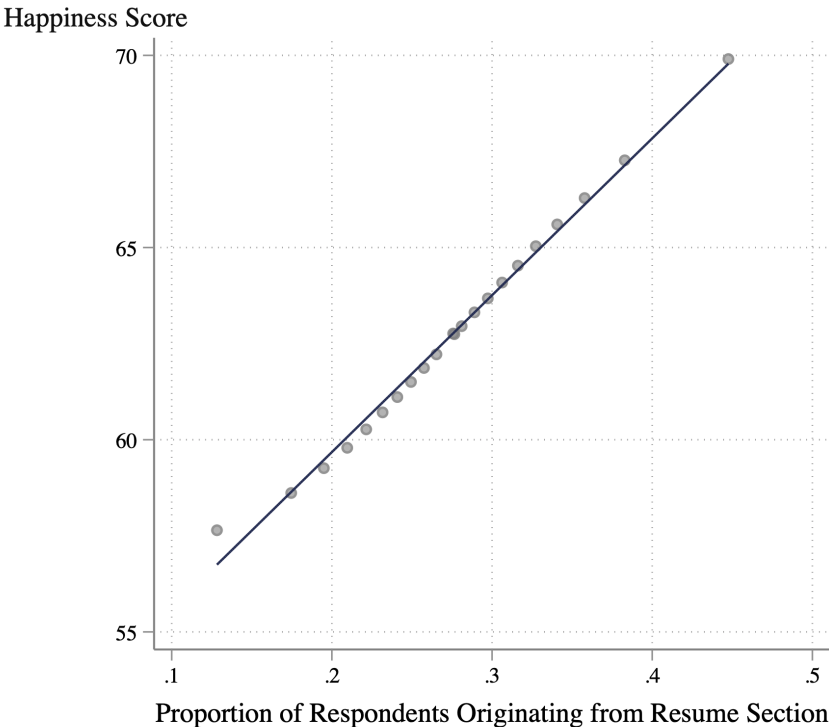
$$A_{ijt} = \beta H_{jt} + U_i + T_t + X'_j + \varepsilon_{ijt} \quad (4)$$

where U_i and T_t are jobseeker and date fixed effects, and X'_j includes 26 fixed effects for the company's industry, 8 fixed effects for the company's number of employees, and indicator variables for if the company is a staffing agency and is in the Fortune 500. I instrument for H_{jt} using the proportion of individual-level surveys that make up H_{jt} that come from click-throughs from the resume section of the site, Z_{jt} .

One key assumption in this set-up is that the share of resume respondents a company has is sufficiently correlated with the happiness score. Figure S10 shows a binned scatterplot of the “first-stage” relationship between survey source and happiness score. Included in the regression are the full set of jobseeker and date fixed effects, along with a set of company observables. The proportion of surveys coming from the resume section has a strong positive impact on the aggregate score. The instrument has a first stage cluster-robust F-statistic of over 990, suggesting that the instrument is sufficiently strong to be valid.

A further key (untestable) assumption is that the effect that the share of resume responses a company has on application behavior runs through its influence on the happiness score only, condition on the covariates and fixed effects in the model.

Figure S10: IV First Stage Effect



Note: Figure shows the relationship between the composition of a company-day's aggregate happiness score in terms of the source of individual-level surveys, a using binned scatter plot. Both the happiness score and the proportion of surveys coming from resumes are first regressed on the full set of controls and fixed effects. The residuals from these regressions are binned across 40 quantiles and plotted as grey dots. The blue line shows the linear fit from a regression using all of the data.

Table S11: IV Estimates of the Effect of Happiness Score on Application Behavior

	DV: Happiness		DV: Applied = 100		
	(1) 1st Stage	(2) Red.-Form	(3) 2SLS	(4) Red.-Form	(5) 2SLS
Proportion Surveys from Resume	4.064*** (0.128)	0.434*** (0.088)		0.611*** (0.085)	
Happiness Score			0.107*** (0.022)		0.104*** (0.023)
Resume-Source Surveys ²				-0.256*** (0.037)	
Happiness Score ²					-0.013*** (0.002)
Observations	18,639,593	18,639,593	18,639,593	18,639,593	18,639,593
Kleibergen-Paap rk Wald F-Stat			1007.4		159.7

Notes: Robust standard errors are in parentheses, adjusted for clustering on companies. Outcome in all models is equal to 100 if the job seeker applied to a job at that company, 0 otherwise. All models include jobseeker and date fixed effects, as well as controls for company industry, company size and dummies for being a staffing agency and in the Fortune 500. Happiness score is re-centered around 0. Proportion surveys from resume is the IV in model (3); proportion surveys from resume and its square are the instruments in model (5). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix H.2 Results

Column (1) of Table S11 reports the key coefficient and standard error from this first-stage regression. Column (2) of Table S11 reports the reduced form relationship between survey source and the probability of applying. The proportion of surveys coming from the resume section has a strong positive effect on application decisions. A coefficient of 4 suggests that displaying a score made up entirely of resume clickers increases the probability of applying by 4 percentage points, compared with a score made up with no jobseekers from the resume source. A one standard deviation increase in this proportion (which is around 0.11) increases the probability of applying by around half a percentage point.

Turning to the 2SLS estimate reported in column (3), a statistically significant coefficient of 0.104 [95% CI: 0.064, 0.149] suggests that a one-point increase in the score, on the 100-point scale, increases the application probability by around 0.1 percentage points.²⁹ In the remaining columns of Table S11, I again estimate the 2SLS model, but this time introduce a quadratic term for the happiness score, as well as a quadratic term of the instrument. The use of two instrumental variables weakens the strength of the first stage, though it remains relatively strong with a cluster-robust F-statistic of over 160.³⁰ In column (5) of Table S11, both linear and quadratic happiness terms enter significantly into the equation, with the expected signs. The happiness score has a positive causal effect on applications, which then flattens out at higher levels of the score.

²⁹ Although the point estimate is slightly larger than the fixed-effect estimate reported above, the fixed-effect estimate falls well within the confidence interval of the 2SLS estimate.

³⁰ As above, when using polynomials, I first re-centre the happiness score to have a mean of zero.

Figure S11: Source of Happiness Surveys Over Time

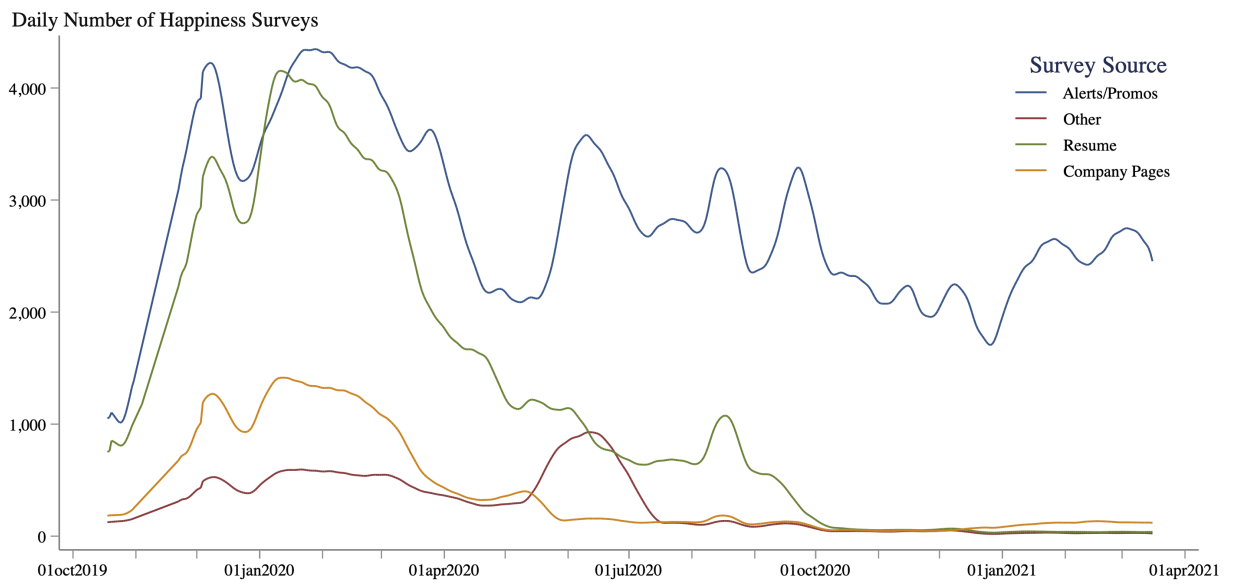
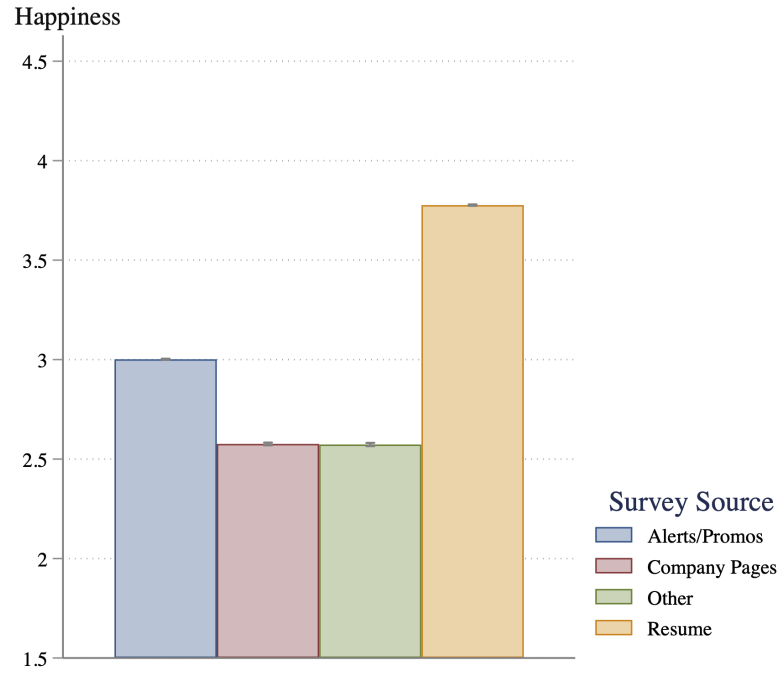
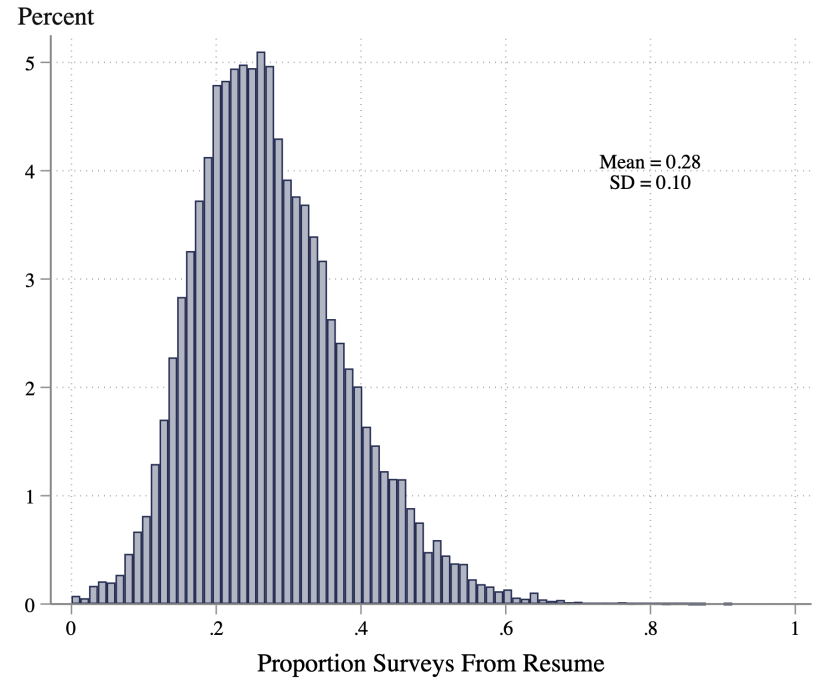


Figure S12: Happiness by Survey Source



(a) Happiness by Survey Source

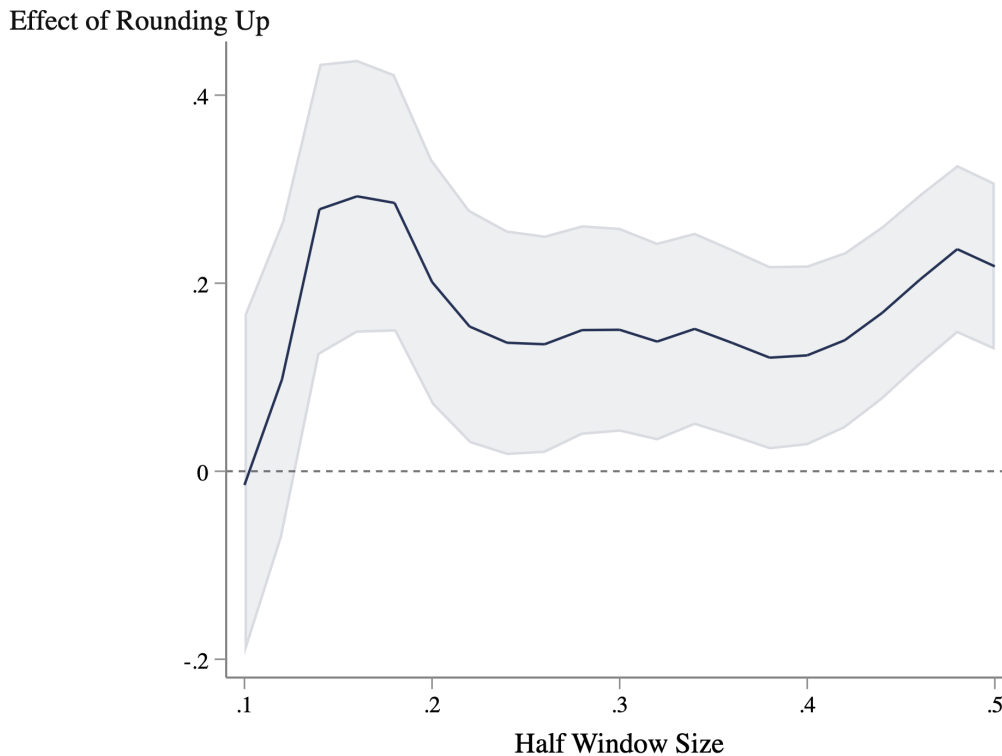


(b) Survey Source Across Companies

Note: Panel (a) plots the raw differences in happiness responses by survey source. Panel (b) shows the source of happiness surveys across company-days.

Appendix I Local Randomization RD Analysis: Robustness

Figure S13: Sensitivity of Window Size Selection



Note: Point estimates and 95% confidence intervals reported, from a series of regressions with differing window sizes around the cutoff points for inclusion (in increments of 0.02 from 0.1 to 0.5). All regressions are linear probability models, with the outcome equal to 100 if the job seeker applies and zero otherwise. All models include a dummy equal to 1 if the score is rounded up to the nearest integer, as well as cutoff fixed effects.

Table S12: Local Randomization RD Analysis: Extra Results

RD Effect	Fisher p-value	Large-sample p-value	Window	N_{below}	N_{above}	Kernel
0.313	< 0.0001	< 0.0001	[-0.2, 0.2]	728,047	725,941	Uniform
0.290	< 0.0001	< 0.0001	[-0.2, 0.2]	728,047	725,941	Triangular

Notes: Table reports local randomization RD estimates of the effect of rounding up the score, using data pooled over the 5 cutoffs. P-values based on 1,000 replications. See Cattaneo et al. (2015, 2017) for more details of the approach.

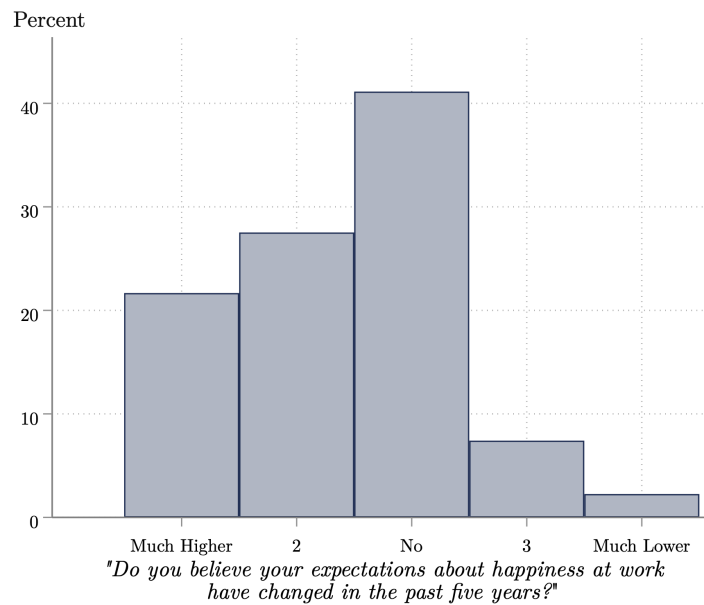
Appendix J Nationally Representative Survey

The survey was conducted online by Forrester Consulting, and funded by *Indeed*. The survey was fielded in March 2021, and sampled adults ages 18+ who reported working either full-time or part-time, or currently not working but were working prior to the COVID-19 pandemic and open to new opportunities. 4,033 US adults were surveyed. To ensure a representative sample, quotas were set by age, education, gender, region, and income.

The survey asked the same questions as are asked on *Indeed* about workplace happiness and the sub-dimensions of workplace happiness. The embedded survey experiment was conducted, with wording and logic outlined below.

Appendix J.1 Further Descriptive Statistics and Description

Figure S14: Worker Expectations



N=4,033.

Survey Wording and Flow

G1. Imagine that a website has asked millions of workers across the country about their happiness at work (including various dimensions like whether people feel a sense of purpose, belonging, trust, appreciation, growth, etc.), in order to help jobseekers make more informed decisions about where to work.

After surveying 4.5 million people working at different companies, this company has published this new data for each company.

How much do you agree or disagree with the following statement:

It is important to me to see information around worker happiness when considering if I should work at a company

01. 1 – Strongly disagree [1]
02. 2 [2]
03. 3 – Neutral [3]
04. 4 [4]
05. 5 – Strongly Agree [5]

G2. Sticking with this example, imagine that employers are all given a score on happiness out of 100 points based on employees' answers to the question "I feel happy at work most of the time." Across all employers on the website, the average happiness score is typically in the range between 60 and 70.

What do you estimate would be the average score for the industry you [WORK/WORKED] in?

G3. What do you estimate would be the average score for your [CURRENT/LAST] company?

RANDOMIZE: A HALF OF SAMPLE WILL ANSWER G4 AND ANOTHER HALF G5

G4_1. Imagine you are looking for a job on the website we have described. You are comparing two positions, both of which are in the same industry and location as your [CURRENT/LAST] job. In this instance, assume the positions, job description and companies are the same outside of pay and happiness levels.

Which position would you be more likely to choose?

1. **Position A:** The workplace happiness score of the company is 65. This position pays you the same as your [CURRENT/LAST] job.
2. **Position B:** The workplace happiness score of the company is 45. The position pays you [SHOW 2, 5, 10, 20, 35 – START FROM 2 AND SHOW EACH CONSECUTIVE LEVEL UNTIL RESPONDENT SELECTS POSITION B; OR "NONE" IF POSITION A WAS SELECTED AT ALL LEVELS]% more than your [CURRENT/LAST] job.

[IF G4_1 SELECTED POSITION B AT 2% SALARY INCREASE AUTOPOPULATE G4_2 WITH "2%" OPTION AND SKIP TO NEXT QUESTION]

G4_2. Imagine another scenario, similar to the previous example.

Which position would you be more likely to choose?

1. **Position A:** The workplace happiness score of the company is 65. This position pays you the same as your [CURRENT/LAST] job.
2. **Position B:** The workplace happiness score of the company is 55. The position pays you [SHOW 2, 5, 10, 20, 35 – START FROM 2 AND SHOW EACH CONSECUTIVE LEVEL UNTIL RESPONDENT SELECTS POSITION B AND UP TO THE LEVEL THAT PRECEDES SELECTION

Survey Wording and Flow

IN G4_1 OR UP TO 35 IF "NONE" IS SELECTED IN G4_1; IF POSITION A WAS SELECTED AT ALL SHOWN LEVELS, AUTO-POPULATE WITH RESPONSE FROM G4_1]% more than your [CURRENT/LAST] job.

G5_1. Imagine you are looking for a job on the website we have described. You are comparing two positions, both of which are in the same industry and location as your [CURRENT/LAST] job. In this instance, assume the positions, job description and companies are the same outside of pay and happiness levels.

Which position would you be more likely to choose?

1. **Position A:** The workplace happiness score of the company is 65. This position pays you the same as your [CURRENT/LAST] job.
2. **Position B:** The workplace happiness score of the company is 75. The position pays you [SHOW 35, 20, 10, 5, 2 – START FROM 35 AND SHOW EACH CONSECUTIVE LEVEL UNTIL RESPONDENT SELECTS POSITION B; OR "NONE" IF POSITION A WAS SELECTED AT ALL LEVELS]% less than your [CURRENT/LAST] job.

[IF G5_1 SELECTED POSITION B AT 35% SALARY DECREASE, AUTOPOPULATE G5_2 WITH 35% AND SKIP TO NEXT QUESTION]


G5_2. Imagine another scenario, similar to the previous example.

Which position would you be more likely to choose?

1. **Position A:** The workplace happiness score of the company 65. This position pays you the same as your [IF S6=1-2: current IF S6=8: last] job.
3. **Position B:** The workplace happiness score of the company is 85. The position pays you [SHOW 35, 20, 10, 5, 2 – START FROM 35 AND SHOW EACH CONSECUTIVE LEVEL UNTIL RESPONDENT SELECTS POSITION B AND UP TO THE LEVEL THAT PRECEDES SELECTION IN G5_1 OR UP TO 5 IF "NONE" IS SELECTED IN G5_1; IF POSITION A WAS SELECTED AT ALL SHOWN LEVELS, AUTO-POPULATE WITH RESPONSE FROM G5_1]% less than your [CURRENT/LAST] job.

Figure S15: Who is responsible for worker happiness?

"Allocate 100 points across the options below according to their impact on employees' happiness at work"



N=4,033.

Figure S16: Reasons for Job Search

Could you please tell us more about the reasons why you would consider new opportunities? Please select the top-3 reasons.

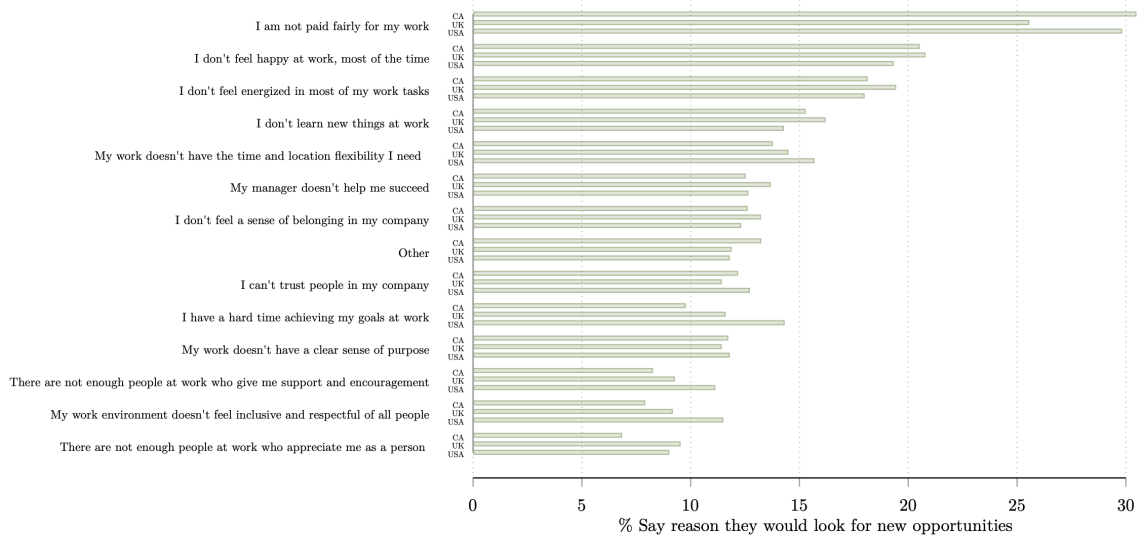

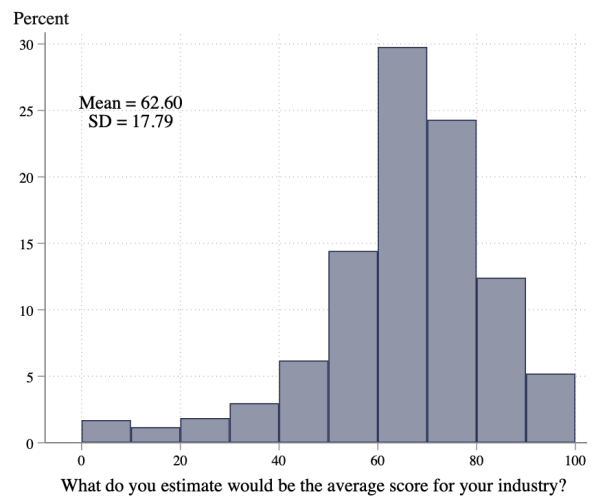


Figure S17: Estimates from Nationally Representative Survey



(a) Own Company



(b) Own Industry


Note: Figure plots the distribution of responses from 4,033 survey respondents in the USA. Survey was conducted in March 2021 – see text for more details.

Table S13: Value of Happiness: Survey Evidence


	N	Mean	Std Dev	Std Err	95% CIs	
<u>United States</u>						
20pt lower happiness and higher pay (45 vs 65)	2,009	20.35	13.20	0.29	19.78	20.93
10pt lower happiness and higher pay (55 vs 65)	2,009	14.20	12.57	0.28	13.65	14.75
10pt higher happiness and lower pay (75 vs 65)	2,024	-12.72	-13.33	-0.30	-12.14	-13.31
20pt higher happiness and lower pay (85 vs 65)	2,024	-16.58	-14.22	-0.32	-15.96	-17.20
<u>United Kingdom</u>						
20pt lower happiness and higher pay (45 vs 65)	771	18.09	12.80	0.46	17.19	19.00
10pt lower happiness and higher pay (55 vs 65)	771	11.91	11.69	0.42	11.08	12.73
10pt higher happiness and lower pay (75 vs 65)	763	-10.50	-11.78	-0.43	-9.66	-11.33
20pt higher happiness and lower pay (85 vs 65)	763	-14.33	-13.51	-0.49	-13.37	-15.29
<u>Canada</u>						
20pt lower happiness and higher pay (45 vs 65)	769	19.42	13.14	0.47	18.49	20.35
10pt lower happiness and higher pay (55 vs 65)	769	13.69	12.69	0.46	12.79	14.59
10pt higher happiness and lower pay (75 vs 65)	763	-11.33	-12.53	-0.45	-10.44	-12.22
20pt higher happiness and lower pay (85 vs 65)	763	-14.57	-13.62	-0.45	-10.44	-12.22

Notes: In the cases where the sequential wage offer was declining, those who refused to switch at 2% are coded as 0%. In the cases where the sequential wage offer was increasing, those who refused to switch at 35% are coded as 35%.

Figure S18: Distribution of Survey Responses in Canada




(a) Higher Pay, Lower Happiness




(b) Lower Pay, Higher Happiness

Note: In all cases, the base job is at a company with a happiness score of 65. Panel (a) reports the share of respondents preferring the company with lower happiness, at differing levels of pay increase. Panel (b) reports the share of respondents preferring the company with higher happiness, at differing levels of pay decrease.

Figure S19: Distribution of Survey Responses in the United Kingdom



(a) Higher Pay, Lower Happiness



(b) Lower Pay, Higher Happiness


Note: In all cases, the base job is at a company with a happiness score of 65. Panel (a) reports the share of respondents preferring the company with lower happiness, at differing levels of pay increase. Panel (b) reports the share of respondents preferring the company with higher happiness, at differing levels of pay decrease.

Table S14: Survey Experiment: Heterogeneity

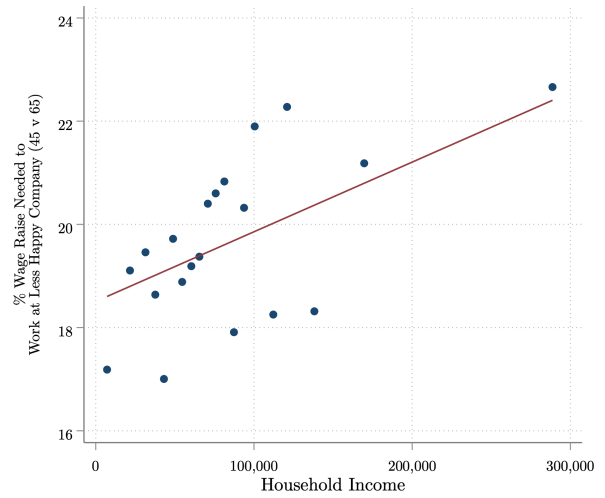
	ln(% Wage Raise Needed to Work at Less Happy Company, With Score):				ln(% Wage Cut Willing to Take to Work at Happier Company, With Score):			
	(1) 45	(2) 45	(3) 55	(4) 55	(5) 75	(6) 75	(7) 85	(8) 85
Age (z-score)	0.021 (0.031)	0.022 (0.029)	0.682 (0.399)	0.059 (0.029)	-0.207*** (0.019)	-0.206*** (0.018)	-0.268*** (0.025)	-0.267*** (0.025)
Men	-0.024 (0.048)	-0.037 (0.040)	0.571 (0.592)	0.054 (0.026)	0.119 (0.074)	0.115 (0.073)	0.085 (0.047)	0.079 (0.048)
Has Children	0.084 (0.039)	0.049 (0.037)	0.176 (0.797)	-0.032 (0.059)	0.175 (0.077)	0.168 (0.074)	0.212 (0.088)	0.202 (0.085)
Has BA	0.099** (0.022)	0.087* (0.025)	0.442 (0.322)	0.037 (0.021)	0.118 (0.055)	0.116 (0.054)	0.144 (0.061)	0.141 (0.060)
Income (log)	0.076 (0.032)	0.068 (0.030)	0.553 (0.238)	0.085** (0.010)	0.041** (0.008)	0.038* (0.010)	0.050*** (0.005)	0.046** (0.007)
Black (v. white)	-0.119 (0.046)	-0.117 (0.044)	-0.853 (0.772)	-0.109 (0.054)	0.153 (0.101)	0.152 (0.100)	0.151 (0.102)	0.150 (0.101)
Hispanic (v. white)	-0.058 (0.090)	-0.067 (0.088)	0.220 (1.468)	-0.002 (0.118)	0.290 (0.192)	0.291 (0.192)	0.299 (0.179)	0.300 (0.180)
Asian (v. white)	-0.040 (0.055)	-0.036 (0.063)	-0.061 (0.655)	-0.004 (0.064)	-0.064 (0.057)	-0.063 (0.054)	0.015 (0.095)	0.017 (0.091)
Other race (v. white)	-0.050 (0.066)	-0.046 (0.058)	-0.654 (0.894)	-0.103 (0.044)	0.009 (0.178)	0.010 (0.178)	0.068 (0.156)	0.070 (0.157)
Looking for job (v. not)	-0.206** (0.030)	-0.176** (0.024)	-2.339*** (0.109)	-0.191*** (0.010)	0.167 (0.067)	0.174 (0.065)	0.212* (0.070)	0.223* (0.069)
UK (v. USA)	-0.101** (0.015)	-0.082** (0.017)	-2.029** (0.247)	-0.140** (0.017)	-0.116*** (0.010)	-0.112** (0.014)	-0.162*** (0.005)	-0.155*** (0.009)
Canada (v. USA)	-0.048*** (0.002)	-0.034*** (0.001)	-0.458* (0.120)	-0.039** (0.005)	-0.005 (0.022)	-0.002 (0.026)	-0.091** (0.020)	-0.085* (0.023)
Own Happiness (z-score)		0.107*** (0.010)		0.086** (0.011)		0.023 (0.021)		0.035 (0.022)
Observations	3494	3494	3494	3494	3492	3492	3492	3492

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


Figure S20: By Income: Wage Willing to Give up to Work at Happier Companies & Wage Raise Needed to Work at Lower Happiness Companies




(a) Happiness Score of 55 (v. 65)



(b) Happiness Score of 45 (v. 65)




(c) Happiness Score of 75 (v. 65)




(d) Happiness Score of 85 (v. 65)

Note: Pooled data from UK, USA, and Canada. Binned scatterplots reported, having adjusted for country fixed effects and controls for gender and age.


Figure S21: By Occupation: Wage Willing to Give up to Work at Happier Companies & Wage Raise Needed to Work at Lower Happiness Companies




(a) Happiness Score of 55 (v. 65)



(b) Happiness Score of 45 (v. 65)



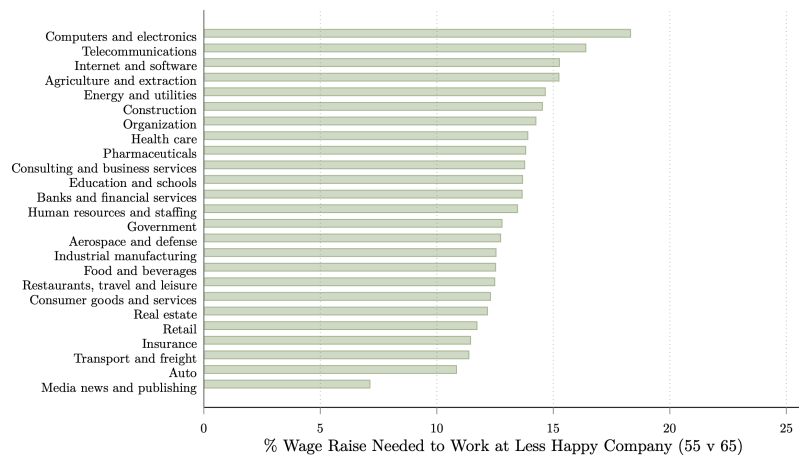
(c) Happiness Score of 75 (v. 65)



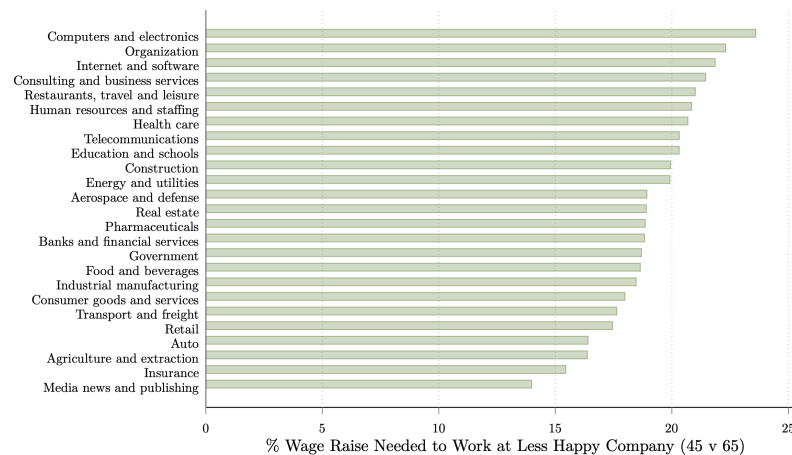
(d) Happiness Score of 85 (v. 65)

Note: Pooled data from UK, USA, and Canada.

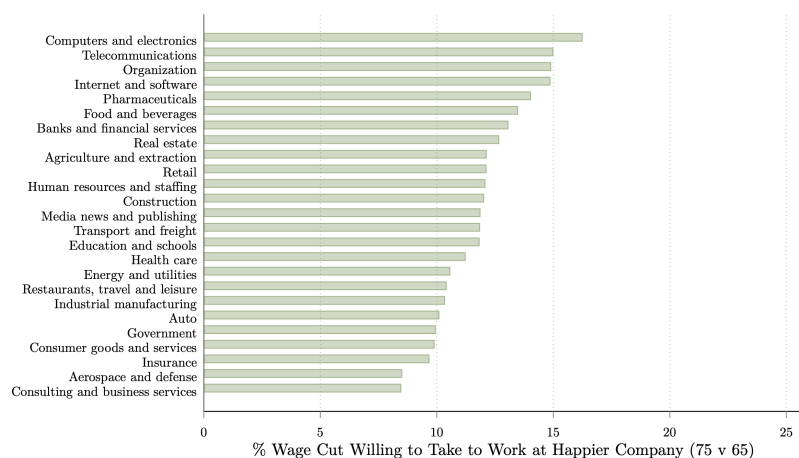
Figure S22: By Industry: Wage Willing to Give up to Work at Happier Companies & Wage Raise Needed to Work at Lower Happiness Companies



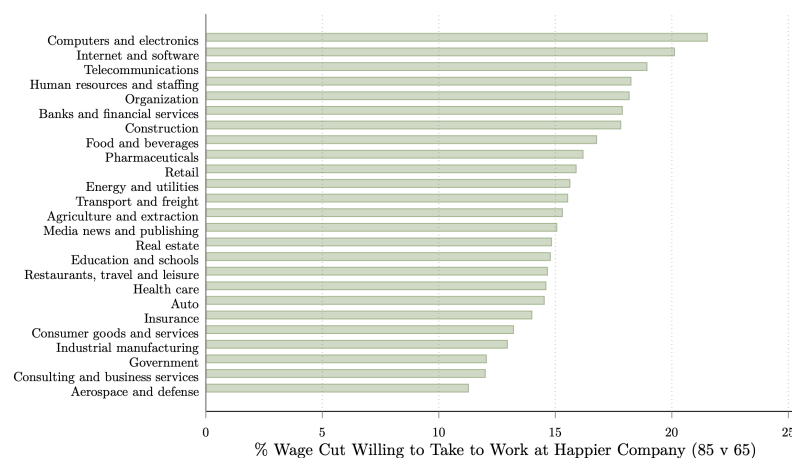
(a) Happiness Score of 55 (v. 65)



(b) Happiness Score of 45 (v. 65)




(c) Happiness Score of 75 (v. 65)




(d) Happiness Score of 85 (v. 65)

Note: Pooled data from UK, USA, and Canada.


Figure S23: By Own Current Happiness: Wage Willing to Give up to Work at Happier Companies & Wage Raise Needed to Work at Lower Happiness Companies




(a) Happiness Score of 55 (v. 65)



(b) Happiness Score of 45 (v. 65)




(c) Happiness Score of 75 (v. 65)




(d) Happiness Score of 85 (v. 65)

Note: Pooled data from UK, USA, and Canada.


Figure S24: By Age Group: Wage Willing to Give up to Work at Happier Companies & Wage Raise Needed to Work at Lower Happiness Companies




(a) Happiness Score of 55 (v. 65)



(b) Happiness Score of 45 (v. 65)



(c) Happiness Score of 75 (v. 65)



(d) Happiness Score of 85 (v. 65)

Note: Pooled data from UK, USA, and Canada.