What's My Employee Worth?

The Effects of Salary Benchmarking

Zoë Cullen
Harvard University

Shengwu Li
Harvard University

Ricardo Perez-Truglia
University of California, Berkeley

Abstract

Firms are allowed to use aggregate data on market salaries to set pay, a practice known as salary benchmarking. Using national payroll data, we study firms that gain access to a tool that reveals market benchmarks for each job title. Using a difference-in-differences design, we find that the benchmark information reduces salary dispersion by 25%. Thus, salary dispersion must stem partly from aggregate uncertainty about the salaries offered by other firms. Our model formalizes how salary dispersion can arise even in competitive labor markets for identical workers when such uncertainty exists, and we discuss implications for an ongoing policy debate.

JEL Classification: D83, J31, J38, M52.
Keywords: compensation, salary benchmarking, inequality, information frictions.

*This draft: March 24, 2024. First Draft: July 25, 2022. Cullen: zcullen@hbs.edu, Rock Center 211, Boston, MA 02163. Li: shengwu_li@fas.harvard.edu, 1805 Cambridge St., Cambridge, MA 02138. Perez-Truglia: ricardotruglia@berkeley.edu, 545 Student Services Building #1900, Berkeley, CA 94720. Special thanks to Brent Weiss and Ben Hanowell for all of their help and feedback. We are also thankful for comments by Sydnee Caldwell, Laura Giuliano, Claudia Goldin, Matthew Grennan, Simon Jäger, Larry Katz, Asim Khwaja, Pat Kline, Ray Kluender, Felix Koenig, Claudio Labanca, Alex MacKay, Alex Mas, Filip Matejka, Enrico Moretti, Bobby Pakzad-Hurson, Simon Quinn, Ben Roth, Benjamin Schoefer, Jesse Shapiro, Isaac Sorkin, Shoshana Vasserman and other colleagues and seminar discussants at NBER Summer Institute (Labor Studies), Harvard University, Columbia University (Econ), Columbia University (GSB), University of Southern California, University of North Carolina (Kenan-Flagler), UC Santa Barbara, INSEAD, University of Chicago (Booth), U.S. Census, CEPR Labor Studies, Essex University, Universita della Svizzera Italiana, Norwegian School of Economics, UC-Berkeley (PF), UC-Berkeley (Labor), Amazon, University of Delaware, University of Copenhagen, University of Cologne, Goethe Univeristy, Firms and Labor Workshop, and the Texas A&M Labor and Public Economics Workshop. This project was reviewed and approved in advance by the Institutional Review Board at Harvard Business School (IRB #20-1779). We thank the collaborating institution for granting access to their data and for all of their help. The collaborating institution did not provide any financial support for the research being conducted. Julia Gilman, Dylan Balla-Elliott, Romina Quagliotti and Xinmei Yang provided excellent research assistance.

1 Introduction

Employee compensation is the largest source of expenditure for companies. Setting the right salaries is of first-order importance. How do companies find out what their employees are worth? While U.S. legislation, in an effort to hinder collusive practices, prohibits employers from sharing compensation information with each other, employers are still allowed to acquire and use more aggregated data provided by third parties. This practice of using market pay data to identify typical market salaries for an internal position is known as *salary benchmarking*.

Survey data suggests that a strong majority of employers use salary benchmarking. For example, in our survey of Human Resources (HR) managers, 87.6% report using salary benchmarks to set pay. Interviews with executives (Adler, 2020) and historical accounts (Adler, 2022) suggest that salary benchmarking plays a prominent role in pay-setting practices. Even HR textbooks dedicate entire chapters to how to use salary benchmarking tools (e.g., Berger and Berger, 2008; Zeuch, 2016). Despite their ubiquity, benchmarking tools rarely make their way into public view, and their effects have not been studied by economists, perhaps partly due to a long-standing assumption that employers have complete information about market pay. Understanding how these tools affect pay setting can shed light on how labor markets operate in practice. Furthermore, the effects of these tools are of direct interest to policy makers, who are seeking to determine whether there is a "procompetitive justification" for these tools (White House, 2021).

We collaborated with the largest U.S. payroll processing company, which serves 20 million American workers and approximately 650,000 organizations. In addition to providing payroll services, the company aggregates salary data from its payroll records in the form of salary benchmarks. Clients can access these benchmarks through a website. This online tool allows users to look up a specific position title (e.g., accountant) and then observe salary statistics for that position, such as the median salary. This is arguably the most advanced benchmark tool and is used by some prominent firms.

Our analysis combines three sources of administrative data. The first dataset corresponds to the payroll records, which include detailed information such as the hire date, position, and compensation. The second dataset contains information about the usage of the benchmark tool, allowing us to reconstruct which firms looked up which positions and when. The third dataset contains historical data on the salary benchmarks, allowing us to impute the salary benchmarks that a firm saw (or would have seen) in the online tool when searching for a specific position at a specific time.

Our data cover the roll-out of the benchmark tool when it was introduced to the market.

Our sample has 586 "treatment" firms that gained access to the tool and 1,419 "control" firms that did not gain access but were selected to match treatment firms according to observable characteristics. We focus on new hires that took place between January 2017 and March 2020, and during a narrow window of 10 quarters around the firm's onboarding date.

We use a difference-in-differences design that makes use of three sources of variation. First, some firms gain access to the tool, and other firms do not. Among the firms that gain access to the tool, some gain access earlier than others. And even within firms with access to the tool, some positions are searched and others are not. According to the provider of the benchmark tool, which firms end up gaining access to the tools, and when they gain access, is arbitrary. For example, when the benchmark tool was introduced to the market, its adoption relied heavily on direct contact from the sales representative of the payroll firm with its clients. As a result, some firms adopted earlier than others due to the arbitrary order in which the sales team approached them. In any case, rather than assuming this variation to be exogenous, we conducted a series of empirical tests (e.g., event-study analysis) to check exogeneity.

We assign each new hire to one of three categories. Searched positions correspond to the 5,266 unique hires in the positions that are (eventually) searched in treatment firms. Non-Searched positions correspond to 39,686 hires in positions that are not searched by treatment firms. Non-Searchable positions correspond to 156,865 hires in control firms, which, by construction, could not search in the benchmark tool. For treatment firms, we analyze how the salaries in Searched and Non-Searched positions evolved around the date when the firm gained access to the benchmark tool. For control firms, we analyze how the salaries in Non-Searchable positions evolved around the date when the firm could have gained access to the benchmark tool: for each control firm, we assign a "hypothetical" onboarding date, equal to the actual onboarding date of the treatment firm that is most similar in observables.

We start by measuring the effects of the benchmarking tool on the distribution of salaries for new hires. We find that after a firm is exposed to the benchmark information in a position, it sets salaries that are closer to the median salary benchmark. On the one hand, firms that would otherwise have paid more than the median benchmark reduce salaries toward the median (for the sake of brevity, we refer to this as "compression from above"). On the other hand, firms that would otherwise have paid less than the median benchmark increase salaries toward the median ("compression from below").

To quantify the effects on salary dispersion around the median benchmark, we construct a dependent variable equal to the absolute %-difference between the employee's starting salary and the corresponding market benchmark.[1] Among Searched positions, the dispersion to the benchmark was on average 19.8 percentage points (pp) before the firms gained access to the tool. After gaining access to the tool, the dispersion dropped from 19.8 to 14.9 pp. This drop is statistically significant (p-value<0.001) and economically significant too, corresponding to a 25% decline.

We present evidence suggesting that the reduction in salary dispersion corresponds to a causal effect. First, an event-study analysis indicates that the reduction in salary dispersion coincides precisely with the timing of access to the benchmark: dispersion was stable in the quarters before the firm gained access to the tool, dropped sharply in the quarter after the firm gained access, and remained stable at the lower level afterwards.

Next, we look at the evolution of pay dispersion for Non-Searched positions. Contrary to the case of Searched positions, we should not expect compression toward the benchmark for Non-Searched positions, because the firms do not see the relevant benchmark. We show that, indeed, for Non-Searched positions the salary dispersion is stable before the firm gains access to the tool and remains stable at the same level after the firm gains access to the tool. For Non-Searchable positions, we should not expect compression toward the benchmark either, because firms did not see the relevant benchmark information. Consistent with that expectation, we find that salary dispersion is stable before the (hypothetical) onboarding date, and remains stable at that same level afterward.

We estimate the effects of salary benchmarking in a difference-in-differences fashion, using two alternative control groups (Non-Searched and Non-Searchable positions, respectively). The results are similar between the two identification strategies: the dispersion around the median benchmark drops by 5.0 pp (p-value<0.001) when Non-Searched positions are used as control, and by 6.2 pp (p-value<0.001) when Non-Searchable positions are used as control. The results are robust to a host of additional validation checks. Moreover, we show that these findings are consistent with evidence from two additional identification strategies: an alternative quasi-experimental design and a survey experiment.

According to anecdotal accounts from interviews with compensation managers, salary benchmarking may play a more prominent role for low-skill positions. Intuitively, employers see candidates for a low-skill position as "interchangeable" (Adler, 2020), so they want to identify the market rate and offer that amount to all candidates. In contrast, in high-skill positions, employers may focus more on tailoring offers to each specific candidate. We categorize low-skill versus high-skill positions using data on education, age, and income. Approximately 42% of the positions in the sample are classified as low-skill (e.g. bank teller,

[1]This formula is related to a common measure of dispersion in statistics and economics: the Mean Absolute Percentage Error. More precisely, the relevant "error" in our context is the difference between the employee's starting salary and the corresponding benchmark (i.e., the median salary for that position).

receptionist), and the remaining as high-skill (e.g. opthalmic technician, software developer). Consistent with the anecdotal accounts, the drop in salary dispersion for Searched positions is larger in magnitude for low-skill than for high-skill positions. The dispersion around the benchmark drops from 14.5 pp to 8.7 pp (p-value<0.001) in low-skill positions, equivalent to a 40% decline. For comparison, in high-skill positions the salary dispersion drops from 21.9 pp to 18.9 pp (p-value=0.021), a 14.6% decline.

Leveraging the same difference-in-differences design, we measure the effects of salary benchmarking on other outcomes such as the average salary and average retention. The estimated average effect on salary is small and statistically insignificant; on average, salaries change by -0.2% (p-value=0.756) when Non-Searched positions are used as control and by $+1.7\%$ (p-value=0.308) when Non-Searchable positions are used as control. For low-skill positions, we find a modest increase in the average salary: $+5.0\%$ (p-value=0.014) when using Non-Searched as control, and $+6.7\%$ (p-value=0.001) when using Non-Searchable as control.

If, after looking up the benchmark information, firms are choosing to increase salaries in low-skill positions, it must be that they expect to get something in return, such as higher retention rates. To explore this conjecture, we measure the effects of salary benchmarking on the retention rates of low-skill employees: more precisely, the probability that the employee is still working at the firm 12 months after the hiring date. Indeed, we find suggestive evidence that salary benchmarking increases retention rates in low-skill positions, by $+6.6$ pp (p-value=0.101) when using the Non-Searched control and by $+6.8$ pp (p-value=0.029) when using the Non-Searchable as control. Moreover, the ratio between the effects on average salary and retention imply a retention elasticity that is consistent with average estimates in the literature (e.g. Sokolova and Sorensen, 2021).

Motivated by the evidence, we propose a simple model with firms that are uncertain about the salary distribution. The model shows that this uncertainty can cause pay dispersion in equilibrium, and its predictions are broadly consistent with our data. In our model, each firm faces aggregate uncertainty about the wages offered by other firms. We assume away standard explanations for wage dispersion, so that we can clearly see the workings of the new mechanism. Workers are identical, firms have identical amenities, firms have no monopsony power, and efficiency wages play no role.

A unit mass of firms simultaneously makes offers to a mass $Q < 1$ of workers. The highest Q offers are accepted. Thus, each firm faces a trade-off: offering a high wage means paying more, but offering a low wage risks leaving the position unfilled. Firms differ only in their marginal revenue from filling the position (their 'value'), which is private information. When one firm has a high value, other firms are also more likely to have high values, and hence to offer high wages. In our model, this relationship is implied by affiliation across firm values, a standard technical condition from auction theory (Milgrom and Weber, 1982). Using this machinery, we characterize a labor market equilibrium in which firms with higher values make higher offers, because they (rationally) have different beliefs about the wage distribution. Thus, the model exhibits wage dispersion in equilibrium; the law of one price does not hold, even though there are many firms and many identical workers.

Suppose that one of the firms covertly gains access to a salary benchmark, learning the population distribution of wages, and hence learning the threshold wage needed to hire a worker. Suppose that all other firms continue to behave as before. If the informed firm's value is above the threshold wage, then the informed firm should raise its offer when it would otherwise be below the threshold. Similarly, the informed firm should lower its offer when it would otherwise be above the threshold. The model predicts that firms will compress their offers in response to the benchmark, raising low offers and reducing high offers. Thus, our empirical results are broadly consistent with a model of profit-maximizing firms and competitive labor markets, once we account for incomplete information.

Furthermore, our model allows us to explore the equilibrium effects of salary benchmarking. Suppose that the benchmark is common knowledge so that all firms make offers with full knowledge of the wage distribution. In the equilibrium with the benchmark, the firms with the highest values hire workers, at a uniform wage that makes the marginal firm indifferent between hiring and not hiring. Compared to the no-benchmark equilibrium, some firms will make higher offers and others will make lower offers. However, we prove that the mean salary is higher (in expectation) under the benchmark equilibrium. The intuition for this result is that without a benchmark, the marginal firm underestimates the strength of the labor market and offers less than it would under full-information competitive equilibrium. When the marginal firm makes a low offer, that makes it easier for other firms to hire, so each firm makes profits that exceed (in expectation) its contribution to social surplus. Salary benchmarks remove those extraordinary profits by resolving the aggregate uncertainty.

Our paper has implications for the study of labor markets. One key observation from labor economics is that seemingly similar workers are often paid different wages (Abowd et al., 1999). There are various putative explanations for wage dispersion; for instance, that workers are different in unobserved ways (Murphy and Topel, 1990), that firms offer different non-wage amenities (Rosen, 1986), that firms have monopsony power (Burdett and Mortensen, 1998), that firms pay efficiency wages (Krueger and Summers, 1988) or share rents for various reasons, such as equity concerns (Card et al., 2018; Kline et al., 2019). We study a novel source of wage dispersion: firms are unsure about the wages paid by other firms, so they offer different wages because they hold different beliefs. Standard labor market models rule out this explanation, because they assume that firms behave as if they know everything about the prevailing distribution of wages (Diamond, 1971; Mortensen and Pissarides, 1994; Postel-Vinay and Robin, 2006; Roussille and Scuderi, 2023). We found that providing aggregate statistics on salaries changes the way firms behave. Thus, the assumption that firms already know this information, while useful, is meaningfully false. Wage dispersion appears to be partly caused by firms' information frictions.

Our study also has policy implications. The regulation of salary benchmarks is an active area of antitrust policy. For instance, in 2023 the Department of Justice and the Federal Trade Commission rescinded a long-standing "safety zone" for salary benchmarks, due to concerns about anti-competitive effects (DOJ, 2023; FTC, 2023). However, a 2021 executive order mandates that agencies must also consider the procompetitive effects of benchmarks (White House, 2021). Our model provides a formal analysis of these procompetitive effects; it suggests that, in equilibrium, salary benchmarking can intensify competition and raise salaries.

Our study contributes to the fields of labor economics, personnel economics, and management by measuring the effects of salary benchmarking tools. Despite their widespread use, there is no evidence on the effects of salary benchmarking. We fill that gap by providing the first causal estimates. Our evidence contributes to the literature on business analytic tools, more generally. The existing literature is theoretical (Blankmeyer et al., 2011; Duffie et al., 2017) or descriptive (Schiemann et al., 2018)—with the notable exception of Grennan and Swanson (2020), which finds that giving U.S. hospitals access to a benchmarking database affects price negotiations for health services.

This study relates to a literature on pay transparency. There is evidence that employees have significant salary misperceptions, even about the salaries of coworkers at the same firm (Caldwell and Harmon, 2018; Caldwell and Danieli, 2021; Cullen and Perez-Truglia, 2022; Roussille, 2023; Cullen and Perez-Truglia, 2023; Jäger et al., 2024). Moreover, the provision of salary information appears to affect employee outcomes such as satisfaction, effort, and turnover (Card et al., 2012; Mas, 2016, 2017; Breza et al., 2018; Dube et al., 2019; Perez-Truglia, 2020; Cullen and Perez-Truglia, 2022; Bennedsen et al., 2022; Duchini et al., 2022; Baker et al., 2023; Cullen and Pakzad-Hurson, 2023). One widespread assumption in this literature is that transparency policies operate by changing the beliefs of employees. We contribute to this literature by showing that firms too, even the large ones, face significant information frictions. Furthermore, our evidence suggests that transparency policies may also affect the beliefs and behavior of employers, not just employees.[2]

[2]Notably, the benchmark that we study was available only to firms. However, the benchmarks available to both employers and employees could affect both sides of the market.

The rest of the paper proceeds as follows. Section 2 describes the institutional context, data and research design. Sections 3 and 4 present the empirical results. Section 5 presents a simple model of salary benchmarks. Section 6 discusses implications for research and policy.

2 Institutional Context and Data Sources

2.1 Background on Salary Benchmarking

The use of salary benchmarks dates back to the 1980s (Adler, 2022). This practice can be found in both the private and public sectors (Faulkender and Yang, 2010; Thom and Reilly, 2015) and is used for all levels of the organization, even executive pay.[3] Many HR textbooks dedicate entire chapters to the practice of salary benchmarking.[4] The following excerpt from one of these textbooks provides an illustration of the type of trade-offs that HR professionals have in mind when using salary benchmarks:

> *"Using surveys to benchmark compensation levels ensures that the pay levels determined by the organization are not extraordinarily misaligned with market practice – i.e., pay is not too low or too high. Determining the appropriate amount of compensation is a balancing act. No organization wants to waste their financial resources by paying too high relative to the market; and those who pay too low risk unwanted turnover from employees looking for a better deal elsewhere."* – Berger and Berger (2008), p. 125.

The first forms of salary benchmarks were compensation surveys administered by consulting firms. To meet these demands, some personnel management consultants grew specialized in providing market data through compensation surveys, with some notable examples being Abbott, Langer and Associates, Korn Ferry, Hayes Group, Mercer, Radford, and Willis Towers Watson. In the last decade, some tech companies have started offering online tools that allow employers and employees to find information about the market salaries in specific positions. Some of these websites, such as Glassdoor, Comparably, and LinkedIn, have become popular because they allow anyone to conduct searches for free. These websites rely

[3]In 2006, the Securities and Exchange Commission issued a new disclosure requirement, requiring firms to state whether they engaged in "any benchmarking of total compensation, or any material element of compensation, identifying the benchmark and, if applicable, its components" (Securities and Exchange Commission, 2006). In fiscal year 2015, more than 95% of the S&P 500 companies disclosed a peer group of firms that they used to benchmark executive salaries against (Larcker et al., 2019).

[4]For example, Chapter 48 from Zeuch (2016) is dedicated to the "Essentials of Benchmarking" and Chapters 9 and 10 of Berger and Berger (2008) are dedicated to "Salary Surveys" and "Benchmarking".

primarily on crowdsourcing: i.e., employees who visit the website can fill out a quick survey reporting their pay at their current or previous employers.

More recently, the largest U.S. provider of payroll services started to offer data analytics tools to their clients, including but not limited to salary benchmarking tools. Payroll data are arguably the highest-quality data one could use for salary benchmarks – any error in payroll is rapidly corrected as it impacts someone's day to day life. Payroll records are even better than tax records in terms of frequency, accuracy and detail.[5] Computing benchmarks based on payroll data has at least three key advantages over alternative sources. Compared to payroll data, survey data are subject to significant measurement error and biases due to the lack of incentives for being truthful and self-selection into the survey. Second, due to the massive sample sizes covering several millions of employees at any point, payroll records allow for benchmarks that are more precisely estimated. Third, due to the high frequency nature of the payroll data, the benchmarks can be updated more frequently.

Salary benchmarking is part of the broader phenomenon of people analytics, brought about by growth in business data capacity. HR professionals leverage data to attract and retain talent, predict employee turnover, identify talent shortages, and other aspects of workforce planning (Davenport and Shapiro, 2010). In a survey of more than 10,000 HR and business leaders in 140 countries implemented by Deloitte in 2017, 71% of companies saw people analytics as a high priority in their organizations, and recruiting was ranked as the highest priority area of focus within that (Collins et al., 2017). HR has become one of the most data-driven functions in major companies (Davenport, 2019).

2.2 Survey on Uses of Salary Benchmarking

To provide evidence on how firms use salary benchmark tools, we conducted a survey of HR professionals in collaboration with the Society for Human Resource Management (SHRM), using its Voice of Work Research Panel. From this point forward, we will refer to this survey as the SHRM survey. The sample encompasses firms of all sizes across various industries, including both the public and private sectors. We invited 9,537 panelists to the survey and had 2,696 responses from July 15 to July 20 2022, for a response rate of 28.3%. More details on the implementation of the survey, sample characteristics, and results are provided in Appendix B.[6]

The first finding is that the use of salary benchmarking is widespread: of the 2,085 re-

[5] For example, payroll records include information about the position title of the employee, which is missing from tax records. And while tax records include the gross taxable income of the employee, they do not show the critical breakdown by base salary, commissions, bonuses, etc.

[6] The full survey instrument is attached as Appendix J.

spondents who participate in setting salaries, 87.6% report using salary benchmarks.[7] Among these respondents, 1,350 complete the entire survey and constitute the main sample used for all the results that follow. Most respondents (72.3%) use multiple sources to obtain market data on salaries. The most popular sources are industry surveys and free online data sources (68.0% and 58.1% of participants, respectively, indicate that they use these). Other popular options are government data (37.1%), paid online data sources (34.4%), compensation consultants (26.3%) and payroll data services (23.2%).[8]

Our survey also explored the ways in which firms utilize benchmarks and the frequency of their use. The vast majority (97.4%) of the respondents use salary benchmarks to set the pay of new hires. There is a lot of variation in how often they use the benchmark information. Only 36.6% of the respondents report using benchmarks to set salaries for all their new hires, while the rest apply them to some, but not all, new hires. Using an open-ended question, we asked respondents why they use the salary benchmark in some cases but not others. There is a wide range of responses that vary substantially between employers. For example, some respondents said they consult the benchmark for positions in which they have less hiring experience. Setting the salary of new hires is by no means the only use of salary benchmarks. The vast majority of respondents report using the salary benchmark for their existing employees, too, and again they typically use it for some employees, but not for all of them. In addition, benchmark tools serve other purposes, such as facilitating financial planning for headcount. In light of how HR professionals use salary benchmarking, we view our intervention as a supply shock to information about competitor prices. The benchmark roll-out we study allows us to observe the incremental impact of an additional, high-quality source.

2.3 The Compensation Explorer Tool

The study builds on an ongoing collaboration with the largest payroll processing firm in America, a publicly traded firm with a current market cap of around $100 billion. This company provides payroll services for 650,000 firms, including many prominent ones, for a total of 20 million employees. In addition to providing payroll services, this company uses massive payroll data from its clients to provide business analytic tools as a subscription service. In this study, we are interested in the *Compensation Benchmark Tool*, consisting of a search engine to view detailed compensation statistics. The online tool allows the user to

[7]The magnitude of this estimate is consistent with the results from an industry survey of 5,003 U.S. firms: 96.3% of them reported that they use some form of salary benchmarking to inform their compensation strategy and structure (PayScale, 2021).

[8]Among our survey respondents, 9.5% use the compensation explorer offered by our partner organization.

browse the benchmarks in different ways. Most prominently, there is a search bar at the top of the screen.

One challenge for the creators of this tool was to aggregate data across different job titles. For example, one firm might call a job "warehouse handler," another might call the same job "inventory handler" or "material handler." The firm converts the raw position titles into a standardized taxonomy with the use of machine learning tools for probabilistic matching, and the firm directly seeks approval of matches from clients, creating new inputs for the algorithm. Our data include a match score that reflects the quality of the match between the firm-specific job title and the title in the taxonomy.[9] Until August 2020, the company used a taxonomy that spanned 2,236 distinct position titles.[10] To illustrate the granularity of this taxonomy, it includes 31 position titles for "teacher" that distinguish between preschool, primary, secondary, middle school, substitute, and special education teachers. On average, there are 3.84 unique position titles for each Occupational Information Network (O*NET) 6-digit code.

To better illustrate how the compensation explorer works, Figure 1 provides a screenshot of this online tool.[11] As soon as the user starts typing a position name in the search bar, an autocomplete function offers suggestions.[12] Once the user selects a position title, the tool provides a job description of the most common tasks for employees in that position, as well as information about the typical qualifications of the candidate.[13]

Once a position has been selected, the benchmark tool provides rich data on compensation statistics for that position. The tool displays the sample size, namely, the number

[9] We restrict our main sample to observations with match scores above the 20th percentile match score in each quarter. The results are similar without this restriction (for details, see Appendix D.3 and Appendix E.2).

[10] Starting September 2020, the company switched to a new taxonomy that expanded the number of position titles. Since our main sample stops in March 2020, our baseline results are not affected by this change.

[11] This screenshot, taken in 2020, had the company's logo and name removed. There have been changes to the tool during the study period, but the overall look and functionality remained similar.

[12] By default, users search positions from the proprietary taxonomy. Because this is the default option, a great majority (70.9%) of searches originate from the proprietary taxonomy. Additionally, a drop-down menu allows users to search using two alternative taxonomies: the client's own position titles (22.6% of searches) and the O*NET taxonomy (6.5% of searches). O*NET searches, however, are excluded from our analysis because we do not have access to data on the corresponding benchmarks.

[13] For example, the job description for an accountant is: "(i) Maintains the accounting operations for a department within the organization; (ii) Checks and verifies records, prepares invoices, vouchers, and filings; (...); (v) Undertakes responsibility for financial analysis and administration or overseeing the projects occasionally." And the corresponding qualifications are: "Requires an undergraduate degree or equivalent experience. For some jobs this may also require a graduate degree or additional certification. This is typically a knowledge worker who applies information and judgment in a specific area to achieve results and solve problems."

of organizations and the number of employees used to calculate the statistics.[14] The most prominent statistic is the median base salary, the first estimate shown on the screen and also highlighted in purple in the bottom panel. The fact that the tool highlights the median base salary is no coincidence, as conversations with the product team indicate that this is the metric their clients are most interested in, and also the type of information highlighted in HR handbooks (e.g., Berger and Berger, 2008; Zeuch, 2016).[15]

The compensation tool defines the base salary clearly and in a manner consistent with research studies using payroll data (Grigsby et al., 2021). For salaried employees, the base pay is the yearly base salary (i.e., before commissions or bonuses). For hourly employees, the annual base salary is defined as the annual equivalent of hourly pay: that is, the hourly wage multiplied by 40 hours multiplied by 52 weeks. The vast majority of the total cash compensation comes from the base salary.[16] Although the median base salary is the most salient piece of information, the tool offers more comprehensive information. As shown at the bottom of Figure 1, the tool provides a chart with various characteristics of the distribution of base salary: in addition to the median, the 10th, 25th, 75th, and 90th percentiles, as well as the average. Similarly, in addition to the base salary, the tool allows the user to learn about bonuses, overtime, and total cash compensation.

The tool also allows the user to filter by some characteristics of employers and employees. For example, users can use a drop-down menu to select a specific industry. They can also use a map to filter by geography, for example, by clicking on their own state.[17] The user can combine any number of filters as long as there are enough observations, more precisely, at least 5 firms and 10 employees, the legal limit. Although the data provider has some information on the use of the tool, the data do not include details on the filters used by the employees. However, our SHRM survey gives us complementary data. When asked about filters, the most popular choices are to filter by industry and by state (87.33% and 84.15% of the participants indicate that they typically apply these filters, respectively). In our baseline specification, we assume that, provided there is a reasonable number of observations, subjects used the

[14]The tool also indicates the quarter to which the statistics refer to, and it even shows some information about the change of the median salary during the past 12 months. The benchmarks are typically stable; for example, the median absolute quarter-over-quarter change in the benchmark is 1.12%.

[15]There is also some evidence that employees, not just employers, pay special attention to median salaries (Roussille, 2023).

[16]In addition to base salary, employees may receive other forms of compensation such as bonuses and commissions, observed in payroll and reported in the benchmark. On average, the base salary comprises 93.07% of the total cash compensation. However, our data do not include equity compensation, which can be a significant part of compensation for some employees, especially at the executive level.

[17]Given its availability of filters by location and the large sample sizes, an advanced benchmarking tool like the one we study could potentially increase firm sensitivity to the wages of their local competitors.

filters by state and industry.[18] Likewise, when a client uses the tool, we do not know whether they were looking at the median salary, the average salary or some of the other statistics. When asked in the SHRM survey, the most popular choice by far was the median (ranked first by 56.73% of the respondents).[19] Thus, we focus on the median salary for the baseline specification.

2.4 Data Sources

We have access to the following datasets:

Payroll Database: it covers all employees in a firm, including new hires. The data is of monthly frequency, covering the period from January 2017 to July 2021. This dataset includes detailed information on the position of the employee, exact hire date, and basic demographics (e.g., gender and age). The data include the full compensation breakdown, although our main focus of interest is the base salary.

Tool Usage Database: the payroll processing company tracks the web navigation of clients using the benchmark tool. For each client, this dataset shows which positions were searched for and when. Due to the firm's preexisting data storage policy, we have access to data starting on September 2019 and until August 2021.[20]

Benchmark Database: this is the database that allows us to reconstruct the search results. For each search observed in the tool usage dataset, we can obtain the corresponding information (e.g., median market benchmark) that was shown to the user at that time. Additionally, we can do counterfactual analysis: i.e., for a client who did not search for a position, we can reconstruct the benchmark they would have seen on the screen had they conducted the search. The benchmarks are updated quarterly, and we have access to the benchmarks from the first quarter of 2017 through the second quarter of 2021. This database contains the compensation benchmarks, at each point in time and for all positions.[21]

There are some additional details about the data that deserve mention. To prevent the influence of outliers, we winsorize all dependent variables in the analysis. For example, in the baseline specification, we winsorize the outcome of absolute dispersion at +/-75 percentage

[18]More precisely, in the baseline specification we assume a firm applies the filters by same state and same industry, but only if that results in at least 30 datapoints.

[19]For more details, see Appendix B.2.

[20]Due to the default setting in the tool, the company would automatically delete the usage data older than six months. For this reason, we do not have access to usage data prior to the date on which we downloaded the data for the first time.

[21]We restrict our sample to employees in positions with available benchmark information, regardless of whether the information was looked up by the firm or not.

points.[22] To minimize concerns about seasonality in hiring of some positions, the baseline specification re-weights observations to maintain the same composition across Standard Occupational Classification (SOC) groups over time.[23] Last, we complement the data provided by the payroll company with other data sources, such as the typical education levels for a position. We cannot recover the set of vacancies or job *offers* associated with a position. So, while some firms may respond to the benchmark by withdrawing a vacancy altogether, we lack sufficient data to explore this additional channel.

2.5 Sample of New Hires

Our main analysis focuses on new hires.[24] There are multiple advantages in focusing on new hires, for example, that we do not need to deal with downward wage rigidities. Furthermore, our survey of hiring managers indicates that one of the primary uses of the tool is to set salaries for new hires. In fact, this view is supported by anecdotal accounts of the partner organization. When hiring new employees, information on salary benchmarks can be used at different stages of the process. For example, information may come in handy earlier in the hiring process, to post wages in job advertisements.[25] The employer may find that information useful later in the hiring process, when producing a first offer, or when deciding how to respond to a counteroffer.[26] Indeed, according to open-ended questions from the SHRM survey, respondents mention all these different margins.

Our main sample of interest consists of new hires from January 2017 through March 2020.[27] Since we are interested in what happens around the date when the firm gains access to the tool, we restrict our sample to a window of 10 quarters around the date of onboarding: i.e., up to 5 quarters before the onboarding date, and up to 5 quarters after the onboarding date.

[22] We exclude outlier observations: employees with annual base salaries over $2,000,000 or below $1,000. Moreover, for the analysis of effects on salary levels, we winsorize the base salary at the 2.5th and 97.5th percentiles within the relevant position.

[23] More precisely, for each position type, we compute the distribution of SOC groups in the month before onboarding and re-weight all the other periods to match that distribution.

[24] In Appendix F, we present additional results for a sample of incumbent employees.

[25] As suggestive evidence that this channel is probably non-negligible, using data from Burning Glass, Hazell et al. (2021) reports that 17% of the job ads include a posted wage or wage range.

[26] As suggestive evidence that this channel may play a role, 16.4% of the companies surveyed by PayScale (2021) report that they shared their own benchmarking data with their employees.

[27] We stop in March 2020 for several reasons, most importantly because we want to avoid our baseline results from being affected by the COVID pandemic. In any case, we show that the results hold when we expand the sample to include new hires after March 2020 – for more details, see Appendix D.

2.6 Firms in the Sample

The salary benchmarking tool is only available to payroll clients that subscribe to cloud services, which launched in late 2015.[28] We observe the exact date when each client was granted access to the tool. Anecdotally, which firms are granted access to the business analytic tool and when they do so depends on many arbitrary factors. During the roll-out, account managers were instructed to introduce the tool to business clients at any opportunity, such as calls pertaining to payroll and other services. Nearly all firms that gained access to the tool did not search for the service or request it, but rather their account manager introduced them to business analytics services as part of a broader conversation. Access during the study window was priced at a "negligible" amount as an additional service for existing payroll clients, according to internal sources familiar with the business strategy during the roll-out period. The fee for the service did not vary on the basis of the number of searches or utilization in any way. Our empirical tests comparing the evolution of firm characteristics as a function of the time to adoption corroborate anecdotes that dissemination was as good as random.

Our main sample comprises 586 firms that gained access to the tool, which we call "treatment" firms. These firms had onboarding dates between December 2015 and January 2020.[29] The vast majority (96%) of treatment firms used the tool at least once. Among access firms, we have suggestive evidence that the tool was being used by a small set of employees, most likely members of the HR or compensation teams.[30]

We obtained data on an additional 1,419 firms that never gained access to the tool, which we call "control" firms. These control firms were selected to match the treatment firms in some observable characteristics: number of employees, state, and 6-digit industry codes. We assign a "hypothetical" on-boarding date to each control firm. We find the treatment firm that is most similar in observable characteristics and assign the onboarding date of that treatment firm as the hypothetical onboarding date for the control firm.[31]

We provide a comparison between our sample of firms and a representative sample of U.S. firms (for more details, see Appendix C.1). In terms of number of employees, our sample is

[28] The benchmarks themselves are based on payroll records for all clients of the payroll company, not just the ones subscribing to the cloud services.

[29] The distribution of onboarding dates is reported in Appendix C.2.

[30] For a subset of the utilization data, we observe an identifier for the person conducting the search. For 50% of the firms with access to the tool, there is a single user who searches. Even in firms with multiple users, searches are concentrated: if you take a random pair of searches, there is a 58.2% probability that they were conducted by the same user. However, these results must be taken with a grain of salt, as it is possible that one account is shared by multiple employees or that one employee is looking up the data on request from other employees.

[31] More precisely, we restrict to all treatment firms in the same industry, and then select the closest treatment firm according to the Mahalanobis distance for firm size and state.

most representative of the top quartile of firms in the United States. In terms of salaries, the employees in our sample are representative of the population of U.S. employees, with the exception that our sample has limited coverage of the bottom quartile of the distribution (earning less than $20,000 per year). Our sample also provides broad coverage of all the U.S. industries.

Table 1 presents some descriptive statistics about the firms in the sample. Column (1) shows that the average firm employs 503 employees, 45.3% of whom are women. The average employee is 34 years old and earns a salary of $46,945. Columns (2) and (3) break down these average characteristics by whether firms gained access to the tool. Due to the large sample sizes, pairwise differences are often statistically significant. However, these differences tend to be modest or negligible in magnitude. This finding should not be surprising given that we asked the partner institution to select control firms that are similar to treatment firms. Columns (4) and (5) break down the treatment firms in the top half and the bottom half based on a measure of higher versus lower utilization of the benchmark tool. Again, firms with high utilization look similar in observable characteristics to firms with low utilization. Columns (6) and (7) compare the characteristics of firms that onboarded earlier in the sample period versus firms that onboarded later. The observable differences are small, consistent with anecdotal accounts suggesting that the reasons why some firms onboarded earlier than others are largely arbitrary.

2.7 Classification of New Hires

We assign each new hire to one of the following three groups:

- Searched Positions: positions a treatment firm that were eventually searched in the compensation explorer by that firm.

- Non-Searched Positions: positions a treatment firm that were not eventually searched in the compensation explorer by that firm.

- Non-Searchable Positions: all positions in the control firms.

One potential concern with the above classification is that some Searched positions may be incorrectly attributed as Non-Searched. This may be due to the limited window of searched data or due to information spillovers.[32] For example, assume that a firm hires accountants

[32] For instance, certain positions might be classified under the Non-Searched category because they weren't searched for after the start of the usage data collection in September 2019, even though they may have been searched for before that date.

and accounting analyst and searches for the benchmark of accountant (and thus this is a Searched position) but not for accounting analyst (the Non-Searched position). Perhaps the two positions are close enough so that the firm is also using the benchmark for accountants to set pay for accounting analysts. In this case, the comparison between Searched and Non-Searched would yield a null effect of the benchmark only because the position accounting analyst is incorrectly being classified as Non-Searched. To minimize the scope for information spillovers, we exclude from the Non-Searched positions all new hires in positions "adjacent" (i.e., in the same SOC group) from those new hires that *were* searched in the same month.

The utilization data shows that while firms have access to the benchmark tool, that does not mean that all firms use it, or that they use it all the time. Consider the 534 firms who had onboarded prior to the last quarter of 2019. During that quarter, 199 (37.3%) of these firms hired in at least one position. These firms searched the benchmark for 20.8% of the positions in which they hired.[33] For this reason, there are substantially more new hires categorized as Non-Searched than as Searched. Also, since our sample includes more control firms than treatment firms, we have an even larger number of new hires in the Non-Searchable category. Our final sample includes 5,266 new hires in the Searched category, 39,686 new hires in the Non-Searched category, and 156,865 new hires in the Non-Searchable category.

In our sample of new hires, we observe 329 unique positions in the Searched category. These positions include all kinds of occupation, such as bank clerk, handpacker, and software developer. We observe a lot of overlap in the positions that different firms are searching for (for details, see Appendix C.2). For example, the 468 hires for Customer Service Representative in the Searched category are distributed across 44 different firms. We also find a lot of overlap across the Searched, Non-Searched and Non-Searchable categories: e.g., there are 468 new hires Customer Service Representative in the Searched category, there are 4,401 hires for that same position in the Non-Searched category and 4,012 in the Non-Searchable category.

Column (1) of Table 2 shows the average characteristics of the employees in the sample of new hires. The average employee is 35 years old, 50.6% of them are female, 81.1% work for an hourly wage, they have an annual starting salary of $41,359 and a median market benchmark of $41,412. The salaries differ from their corresponding median benchmarks (in absolute value) by an average of 20.4%. The last rows show the main occupation groups in the sample: 19.8% of the positions are in office and administrative support, 8.0% in management, 6.6% in production, 9.3% in transportation and material transport, 4.8% in building and ground cleaning, and the rest (51.5%) belong to other groups.

[33] More precisely, around 62.3% of these firms did not search for any of the positions in which they hired; among the remaining firms, they looked up on average 55.2% of the positions in which they hired.

Next, we can compare the characteristics between the treatment and control groups. As usual in difference-in-differences designs, the key identifying assumption is that, in the absence of treatment, the outcome of interest would have evolved similarly between treatment and control groups. Corroborating evidence for this assumption can be seen by testing whether, prior to the onboarding date, the outcome of interest evolved similarly between treatment and control. As a result, it should not matter whether the treatment and control groups are different in the baseline outcome or in other observable characteristics. However, it is always reassuring to check that the differences between the treatment and control groups are not large. Columns (2) through (4) of Table 2 break down the average characteristics for each of the three categories: Searched, Non-Searched and Non-Searchable. Perhaps the two most important characteristics are the (pre-treatment) salary and its absolute %-difference with respect to the median benchmark, because they constitute the outcome variables in the analysis that follows. The differences are economically modest. For example, the average salaries are $39,064, $42,013 and $41,405 in the Searched, Non-Searched and Non-Searchable categories, respectively. Despite the modest magnitude of the difference between the Searched and Non-Searchable groups, due to the large sample sizes, the difference is statistically significant (p-value = 0.013). The difference between the Searched and Non-Searched groups is not significant (p-value = 0.617). For the other characteristics, the pairwise differences are again almost always statistically significant, but tend to be economically small. Some exceptions include that, compared to Non-Searched and Non-Searchable positions, Searched positions have a higher proportion of female employees and a greater share of office and administrative support roles.

3 Effects on Salary Dispersion

3.1 Non-Parametric Estimates

To begin, we examine the impact of salary benchmarking on the distribution of salaries around the median benchmark. We start with a non-parametric analysis of the data by means of histograms. More precisely, we look at the distribution of the difference between the salaries chosen by the firms and the benchmarks they saw (or could have seen) in the benchmark tool. The results of this analysis are presented in Figure 2. Each panel corresponds to a different type of position (e.g., Searched). In each panel, the x-axis denotes the difference between the starting salary and the corresponding median benchmark. For example, the middle bin corresponds to salaries that are close ($\pm 2.5\%$) to the median benchmark, the bins on the left half of the figure correspond to salaries below the benchmark, and the bins on the right half correspond to salaries above the benchmark.

Panel A of Figure 2 corresponds to the Searched positions, with solid gray bins corresponding to employees who were hired before the firm gained access to the benchmark tool (i.e., when the benchmark information *was not* visible to the firm) and the hollow red bins correspond to employees hired after the onboarding date (i.e., when the benchmark information *was* visible to the firm). The comparison between the two histograms from Panel A suggests that, after onboarding, salaries are more compressed toward the median benchmark. More precisely, we observe compression from above and compression from below: there is a decline in the probability of observing salaries above the benchmark, as well as a decline in the probability of observing salaries below the benchmark.

One simple way to summarize the compression toward the benchmark is by noticing that firms are more likely to "bunch" at the benchmark: the probability that the firm chooses a salary close (±2.5%) to the median benchmark increases from 11.6% before onboarding to 22.1% after onboarding. Another way to summarize the dispersion around the benchmark is by means of the absolute mean difference. This metric suggests that, among Searched positions and before the firms gained access to the tool, the difference between the salaries and the corresponding benchmarks was on average 19.4 pp. After gaining access to the tool, the average distance from the benchmark decreased from 19.8 to 14.9 pp, a change that is highly statistically significant (p-value<0.001) and also large in magnitude (equivalent to a 24.7% drop).

For the purposes of placebo exercises, we use the Non-Searched and Non-Searchable positions as two alternative control groups. The results for Non-Searched positions are presented in Panel B of Figure 2. Because the firms never see the benchmarks for Non-Searched positions, we should not expect compression toward the benchmark. The dispersion around the median benchmark is similar in magnitude in the pre-onboarding period (20.8 pp) to the post-onboarding period (22.0 pp). Due to the large sample sizes, this difference is precisely estimated and thus statistically significant (p-value<0.001). However, the difference is small in magnitude and much smaller than the corresponding difference for the Searched category (reported in Panel A). In turn, Panel C of Figure 2 presents the results for the Non-Searchable positions. Because firms cannot see the benchmarks for the Non-Searchable positions, we should not expect compression toward the benchmark for this category. We find that dispersion around the benchmark is similar in magnitude in the pre-onboarding period (21.1 pp) as in the post-onboarding period (21.9 pp). Due to the large sample sizes, the difference is again statistically significant (p-value<0.001). However, most importantly, the difference is negligible in magnitude.

We find that salaries get compressed toward the median market pay. On the one hand, this evidence is consistent with anecdotal accounts and survey data indicating that median pay plays a prominent role. On the other hand, this result may be surprising in that firms could have chosen to be stingy, for example, by compressing around the 25th percentile of market pay instead of the median. For a more direct comparison, Appendix D.1 reproduces the analysis, but instead of using the median benchmark, it uses each of the alternative benchmarks: the average pay and the 10th, 25th, 75th and 90th percentiles. The results confirm that salaries are compressed mostly toward the median market pay.

In addition to observing compression around the benchmark, the data suggest that there is significant "bunching" at exactly the median salary.[34] The bunching at the median could reflect a genuine interest in this feature of the distribution. For example, when asked what statistics they care about, the most popular choice was the median. The median salary also plays an important role in, for example, HR handbooks (e.g., Berger and Berger, 2008; Zeuch, 2016). However, an alternative interpretation is that the bunching at the median is due to the way in which the information is presented in the benchmarking tool. As can be seen in Figure 1, the median benchmark is highlighted in multiple ways. In the chart that depicts the salary distribution, the median is depicted in purple, distinguishing it from other distribution features, such as the average and the 25th percentile, which are depicted in gray. Second, the median salaries are prominently displayed at the top of the page, whereas other distribution features are shown at the bottom of the screen only. This arrangement could lead users, especially those who pay limited attention, to focus disproportionately on the median. Furthermore, given the payroll company's expertise in pay setting, users might interpret the median's prominence as an implicit recommendation to use that figure for pay-setting.

The bunching indicates that some firms adjust their salaries completely to the benchmark. However, the evidence suggests that other firms adjust partially. There are at least two potential explanations for this finding. First, if they update in a Bayesian manner, firms should form posterior beliefs about market values by taking a weighted average between their prior belief and the signal they observe in the tool. If their beliefs adjust partially, the salaries should adjust partially too. Second, internal equity concerns may be the reason firms do not fully update toward the observed benchmark. Recent research indicates that pay equity concerns may be important in the workplace. For example, evidence indicates that employees are demoralized when they discover that they are paid less relative to their coworkers in the same position (Card et al., 2012; Breza et al., 2018; Cullen and Perez-Truglia, 2022). When firms look up a position and find out that they are under-paying or over-paying, they face a dilemma. On the one hand, they may want to adjust their offers to

[34] The bunching is even more salient in Appendix Figure D.1, which is identical to Figure 2 except that it uses narrower bins.

better align with the market benchmark. On the other hand, they might prefer to adhere to their internal benchmarks to avoid compensating new hires either more or less than their incumbent employees.[35]

3.2 Econometric Model

We continue with the difference-in-differences design. Let subscript t denote time, i index employees, and j index firms. Let $\omega_{i,j,t}$ be the starting base salary of employee i hired by firm j at time t. And let $\bar{\omega}_{i,t}$ denote the corresponding benchmark: i.e., the median base salary that the search tool indicates for the position of employee i at time t. Let $Y_{i,j,t}$ denote the outcome variable. For example, in this section the outcome of interest is the absolute difference between the employee's salary and the benchmark: $100 \cdot |\frac{\omega_{i,j,t}-\bar{\omega}_{i,t}}{\bar{\omega}_{i,t}}|$. This outcome is multiplied by 100 so that the effects can be readily interpreted as percentage points.

We have two distinct difference-in-differences designs: one based on the comparison between Searched and Non-Searched positions, and the second one based on the comparison between Searched and Non-Searchable positions. For the sake of brevity, we will use Θ_1 to refer to observations categorized as either Searched or Non-Searched, and Θ_2 to the set of observations categorized as either Searched or Non-Searchable. Let $T_{i,j}$ be a dummy variable that takes the value 1 if the employee i's position at firm j was categorized as a Searched position, and 0 if it was categorized as Non-Searched or Non-Searchable. Let $A_{j,t}$ be a dummy variable that takes the value 1 if firm j has access to the benchmark tool in period t and 0 otherwise. This variable takes the value 0 before the month of onboarding and 1 afterward.[36] Let δ_t denote year dummies, ψ_p^k denote position dummies and $X_{i,j,t}$ denote a vector of additional controls consisting of the employee's age, a dummy for gender, and a dummy for hourly pay. And let $\epsilon_{i,j,t}^k$ be the error term. Unless stated otherwise, all of the analysis in this paper uses standard errors that are clustered at the firm-position-month level. Consider the following regression specification:

$$Y_{i,j,t} = \alpha_1^k \cdot A_{j,t} \cdot T_{i,j} + \alpha_2^k \cdot A_{j,t} + \alpha_3^k \cdot T_{i,j} + X_{i,j,t}\alpha_4^k + \delta_t^k + \psi_p^k + \epsilon_{i,j,t}^k,\ \forall \{i,j,t\} \in \Theta_k \qquad (1)$$

When $k = 1$, equation (1) boils down to the first identification strategy, based on the comparison between Searched and Non-Searched groups. When $k = 2$, equation (1) boils down to the second identification strategy, based on the comparison between Searched and Non-Searched. In both cases, the difference-in-differences coefficient of interest is α_1^k, which

[35]In 44.4% of Searched observations, the employers were hiring in a position where there were no incumbent employees. For these hires, there are no internal equity concerns.

[36]In the case of control firms, this would correspond to the "hypothetical" onboarding date.

measures the effect of the benchmark tool. For instance, α_1^1 measures the difference in outcomes between Searched (treatment) and Non-Searched (control) in the post-onboarding period relative to the pre-onboarding period.

These two alternative difference-in-differences designs are based on different control groups, and as such they may offer different advantages and disadvantages. For example, one advantage of using Non-Searchable positions as control group is that it is not subject to the potential concern of misattributing Searched positions as Non-Searched positions (as discussed in Section 2.7 above).[37] Although we do not have a strong preference for one strategy versus the other, we believe that being able to compare the results across the two strategies provides a meaningful validation check for the research design.

As a formal test of pre-trends, we follow the standard practice in difference-in-differences design by introducing a "fake" treatment dummy ($A_{j,t}^{\text{fake}}$) that is identical to the true post-treatment dummy ($A_{j,t}$) except that it takes value 1 in the two quarters before the onboarding date and zero otherwise:

$$Y_{i,j,t} = \alpha_1^k \cdot A_{j,t} \cdot T_{i,j} + \alpha_2^k \cdot A_{j,t} + \alpha_3^k \cdot A_{j,t}^{\text{fake}} \cdot T_{i,j} + \alpha_2^k \cdot A_{j,t}^{\text{fake}} + \alpha_4^k \cdot T_{i,j} + X_{i,j,t}\alpha_5^k + \delta_t^k + \psi_p^k + \epsilon_{i,j,t}^k, \ \forall \{i,j,t\} \in \Theta_k \tag{2}$$

The coefficient of interest is α_3^k, which measures whether the outcomes were already diverging between the treatment and control groups before the onboarding date. Under the null hypothesis of no differences in pre-trends, we expect this coefficient to be zero. Furthermore, we can extend the econometric framework to an event-study analysis, by expanding $A_{j,t}$ into a set of dummies. Let $A_{j,t}^s$ be a dummy variable that takes the value 1 if the firm onboarded on period $t-s$. For example, $A_{j,t}^{+1}$ would take the value 1 one quarter post-onboarding, while $A_{j,t}^{-4}$ would take the value 1 four quarters prior to onboarding. And let S be the set of nonzero integers between -5 and +5, except for -1 (the reference category).[38] We expand equation (2) as follows:

$$Y_{i,j,t} = \sum_{s \in S} \alpha_{1,s}^k \cdot A_{j,t}^s \cdot T_{i,j} + \sum_{s \in S} \alpha_{2,s}^k \cdot A_{j,t}^s + \alpha_3^k \cdot T_{i,j} + X_{i,j,t}\alpha_4^k + \delta_t^k + \psi_p^k + \epsilon_{i,j,t}^k, \ \forall \{i,j,t\} \in \Theta_k \tag{3}$$

The set $\alpha_{1,s}^k \ \forall s \in S$ corresponds to the event-study coefficients. For example, $\alpha_{1,+1}^k$ would correspond to the effect one quarter post-onboarding, relative to the base category of one quarter pre-onboarding.

[37] One potential advantage of utilizing Non-Searched positions as the control group is that it could avoid the concern of capturing effects from tools other than the compensation explorer.

[38] In all the analysis, we drop observations for employees who were hired in the exact month of onboarding. Due to the coarseness of the timestamps, it would be impossible for us to distinguish between the hires that were post- vs. pre-onboarding.

3.3 Difference-in-Differences Estimates

Figure 3 presents the event-study analysis. In each of the panels, the x-axis corresponds to the time since the date of onboarding, from -5 (i.e., 5 quarters prior to the month of onboarding) to +5 (i.e., 5 quarters after the month of onboarding). The y-axis corresponds to the salary dispersion around the median benchmark, with a higher value indicating that salaries are farther away from the benchmark. The minimum value of 0 corresponds to the extreme case where all salaries are exactly equal to their respective median benchmarks. And a value of 20 would mean that the salaries differ from the benchmark, on average, by 20%. To make the interpretation of effect sizes more straightforward and intuitive, we follow Hastings and Shapiro (2018) by normalizing the y-axis. In this and all other event-study graphs, all coefficients are shifted by the same constant to match the average of the baseline outcome in the pre-treatment period. That is the reason why the coefficient for quarter -1 is the omitted category, yet its value is different from 0. Last, the left panels (A and C) of Figure 3 correspond to the comparison between Searched and Non-Searched categories, while the right panels (B and D) correspond to the comparison between Searched and Non-Searchable categories.

The results from Figure 3 indicate that the effects on salary dispersion coincide precisely with the timing of access to the benchmark: the dispersion with respect to the benchmark was stable in the quarters before the firm gained access to the tool, dropped sharply in the quarter after the firm gained access, and remained stable at the lower level afterward. More precisely, Panel A of Figure 3 shows the evolution of the outcome separately for the Searched positions (denoted in red dots) and Non-Searched positions (blue squares). For the Searched positions, the dispersion with respect to the benchmark was stable at around 19.8 pp prior to the onboarding, but then dropped sharply to around 14.9 pp in the quarter after onboarding and remained stable at that lower level afterwards. In contrast, the dispersion in Non-Searched positions was stable around 20.8 pp prior to onboarding, and remained stable at a similar level (22.1 pp) after the onboarding date. Panel C of Figure 3 corresponds to the difference between the two series from Panel A. This difference-in-differences estimate suggests that the benchmark tool reduced the salary dispersion from 19.8 pp to 14.8 pp (p-value$<$0.001), equivalent to a 25.3% reduction.

Regarding the second identification strategy, Panel B of Figure 3 provides a comparison between Searched (denoted in red dots) and Non-Searchable (purple squares) positions. While the outcome dropped sharply after onboarding for Searched positions, it remained stable around the date of onboarding for Non-Searchable positions. Panel D of Figure 3 corresponds to the difference between the two series in Panel B. The difference-in-differences estimate suggests that the benchmark tool reduced the salary dispersion from 19.8 pp to 13.6

pp (p-value<0.001). The drop in dispersion from Panel D (6.2 pp) is close in magnitude to the corresponding drop from Panel C (5 pp) – furthermore, these two effects are statistically indistinguishable from each other. The similarity of results across both identification strategies reinforces the validity of the research design.

In Panel C Figure 3, a pattern is evident that suggests diminishing effects over time. However, there are two significant caveats associated with this finding. First, the estimation of this pattern lacks precision (notably, the confidence interval for each individual post-treatment quarter overlaps with the average effect over the entire post-treatment period). Second, this pattern of decreasing effect is not observed in other specifications, for example, when using Non-Searchable as the control group (Panel D of Figure 3) or when examining different outcomes (e.g., retention).

3.4 Robustness Checks

Table 3 presents the difference-in-differences estimates in table form, which summarizes the difference-in-differences results in fewer coefficients. This simpler approach maximizes statistical power and is also more practical for comparing results across specifications. Panel A of Table 3 presents the post-treatment coefficients (α_1^k from equation (1)). Column (1) of Table 3 corresponds to the baseline specification. The post-treatment coefficients are negative and statistically significant: -4.775 (p-value<0.001) when using Non-Searched positions as control group, and -6.149 (p-value<0.001) when using Non-Searchable positions as control. In turn, Panel B presents the corresponding "pre-treatment" coefficients (α_3^k from equation (1)). Consistent with the assumption of no differences in pre-trends, the pre-treatment coefficients in column (1) are close to zero (-0.346 for the comparison with Non-Searched positions and -0.310 for the comparison with Non-Searchable positions), statistically insignificant (p-values of 0.749 and 0.604, respectively) and precisely estimated.

Each of columns (2) through (12) of Table 3 is identical to column (1) except for one change to the baseline specification. Column (2) uses an alternative version of the dependent variable based on the log-difference: $100 \cdot |log(\omega_{i,j,t}) - log(\bar{\omega}_{i,j,t})|$. Just like in column (1), the outcome from column (2) is multiplied by 100 so that it can be interpreted (approximately) in percentage points. The results from column (2) are qualitatively and quantitatively consistent with the results from column (1). In column (3), we measure dispersion with a dummy variable that takes the value 100 if the salary is more than 10% away from the median benchmark and 0 otherwise. Again, the results are qualitatively and quantitatively similar between columns (1) and (3). For example, the first post-treatment coefficient from column (1) suggests that, relative to baseline, the dispersion dropped by 24.1% ($= \frac{4.775}{19.812}$), while the corresponding coefficient from column (3) suggests a decrease of 25.5% ($= \frac{16.270}{63.732}$).

The specification of column (4) of Table 3 is different from column (1) in that it is winsorized at $\pm$ 100% instead of $\pm$ 75%. Column (5) uses heteroskedasticity-robust standard errors instead of clustered standard errors. Column (6) does not include any of the additional control variables. Column (7) excludes position fixed effects. Column (8) includes firm fixed effects instead of position fixed effects. Column (9) excludes positions in which tipping may play a major role (e.g., waiter or waitress). Column (10) restricts the sample to include only the positions that appear in the list of 329 Searched positions. Column (11) does not re-weight observations by SOC groups. Column (12) restricts the sample to individuals aged 21 to 60 years. In all these alternative specifications, the results are qualitatively and quantitatively similar to those in column (1).

In Appendix D, we present some additional results and robustness checks. For instance, Appendix D.3 show that the results are robust for a range of additional specifications, such as extending the sample after March 2020. And Appendix D.2 shows that there are no significant effects on the composition of new hires.

We present two additional exercises to corroborate the validity of the quasi-experimental findings. For the sake of brevity, these additional results are reported in the Appendix and summarized below.

The first piece of evidence, reported in Appendix B.4, consists of a survey experiment that we embedded in the SHRM survey. We ask the participants to pick two positions for which they are planning to hire in the future, and we elicit the annual base salary they are willing to offer for these new hires. Next, we provide them (hypothetical) information on the median salary benchmark for that position. Participants receive a benchmark that is 15% above their initial salary offer or 15% below. After the respondent receives the benchmark information, we re-elicit the salary they are willing to offer for that position. The results of this survey experiment are largely consistent with the results presented above. More precisely, the experiment shows that the salary offers get compressed toward the benchmark, both from above and from below.

The second piece of evidence, presented in Appendix H, utilizes quasi-random shocks to salary benchmarks in some specific positions. Drawing inspiration from Derenoncourt et al. (2021), we identify a unique instance where large firms abruptly raise the base salary for a specific position by 10% or more. We first show that this shock is sudden and localized: the salary benchmark displayed in the tool rises sharply for that position, but not for other closely related positions. Through an event-study analysis, we demonstrate that, among other firms with access to the benchmarking tool, salaries for affected positions converge to the new benchmark provided they searched the affected benchmark. In contrast, the convergence occurs at a much slower pace for firms that did not search for the affected position or firms that did not have access to the tool.

3.5 Magnitude of the Effects

The effect of benchmarking on salary dispersion documented above is not only highly statistically significant but also large in magnitude. Next, we discuss some reasons why those results may under-estimate or over-estimate the true magnitude of the effects.

On the one hand, our results may lead to under-estimation of the effects of benchmarking due to multiple sources of attenuation bias. First, the tool we study is not the only source of data on market salaries. Firms in the treatment and control groups may be using other sources of data on market salaries in addition to the benchmark tool that we study. Therefore, our estimates should be interpreted as the effect of adding one additional source of benchmarking information.[39] Second, we do not observe precisely which filters the clients are using in the benchmarking tool, and we do not track whether they focus on one particular statistic or another (e.g., median vs. mean). This means that the benchmark we measure is subject to measurement error, thus introducing attenuation bias. Third, in some cases, we may incorrectly assume that the act of looking up the benchmark was related to setting pay for a new hire in that position, when, in reality, it may be to negotiate with an incumbent employee. Likewise, when multiple people are hired in a particular firm-position cell, our specification implicitly assumes that the firm will use that information for everyone who gets hired in that position going forward. However, perhaps the manager was looking that information up for one specific new hire (e.g., someone with an outside offer), and perhaps the manager forgets the information shortly thereafter. Such misclassifications of a new hire in or out of the Searched group may introduce attenuation bias as well.

On the other hand, it is possible that our results over-estimate the importance of salary benchmarks. To the extent that the effects can be heterogeneous across positions, we estimate a treatment effect on the treated. In other words, we estimate the effects of salary benchmarking for positions that end up being searched. Had they been searched, the effects could have been different for positions that were not searched. For example, following the logic of rational inattention, it could be argued that firms look up the positions for which they value information the most. If they value the information the most, they are arguably more likely to use it. In that case, our estimates for the positions that are looked up may overestimate the strength of information frictions for the average position. However, the fact that we estimate the effects of treatment on the treated is not necessarily a limitation. On

[39] In Appendix section C.4, we compare our proprietary salary benchmark with an free public benchmark using popular positions. We show that there are significant discrepancies between the propietary benchmarks and the free benchmarks, although there does not seem to be a systematic positive or negative bias.

the contrary, for some purposes, the treatment effects on the treated may be most relevant. For example, from the perspective of policy implications, the counterfactual of interest is not what would happen if all firms were "forced" to look up every position, but what would happen if all firms had access to look up the positions they want. In that sense, the treatment effects on the treated are the right object of interest.

To assess the extent to which our results were surprising, we also conducted a forecast survey with a sample of 97 experts, most of whom are economics professors specializing in labor economics. The experts received a brief explanation of the context and then made predictions about the effects of the benchmark tool. In summary, most experts expressed low confidence in their own forecasts, and a minority of experts were able to predict key results, such as the effect on salary dispersion documented above. The complete design of the forecast survey and the results are presented in Appendix I.

3.6 Heterogeneity Analysis

The above analysis estimates the average effects of salary benchmarking across all sorts of positions, which may mask substantial heterogeneity.

A key distinction often highlighted in interviews with HR professionals is between low-skill and high-skill positions. On the one hand, low-skill positions involve standardized tasks, minimal training, and can be easily monitored. As one HR practitioner put it, candidates for a low-skill position are "viewed as interchangeable" (Adler, 2020). As a result, firms may want to look up the market rate and offer exactly that amount to all candidates. According to anecdotal accounts, once a candidate is deemed qualified for the job, his or her pay is a function of the job, not its individual characteristics. Low-skill candidates are given take-it-or-leave-it offers, and the candidate's efforts to ask for more are not only rejected, but are even considered inappropriate (Adler, 2020). On the other hand, in high-skill positions, there can be large differences in quality from one candidate to another. HR professionals emphasize the importance of tailoring offers to specific candidates (Adler, 2020). The firm may still look up and use the salary benchmark as a starting point, but there are other factors that can come into play, such as the line manager's opinion of the candidate, the candidate's own salary history, outside offers, and salary expectations. Consistent with this view, survey data suggest that, relative to low-skill candidates, high-skill candidates are substantially more likely to engage in salary negotiations (Hall and Krueger, 2012).

In our sample, we categorize positions as low-skill or high-skill using information on education, age, and earnings. In the first step, we identify the positions in O*NET Job

Zones 1 and 2, which generally require no more than a high school diploma.[40] In the second step, we exclude positions in which the average worker is older than 31 years or has an annual salary greater than $30,000. Approximately 42% of the sample is classified as low-skill, and the remaining 58% as high-skill. Some examples of low-skill positions are bank teller, hand packer and receptionist; some examples of high-skill positions are ophthalmic technician, production operations engineer, and software developer.[41]

Figure 4 breaks down the baseline results from Figure 2 by low-skill and high-skill positions. The panels on the left hand side of Figure 4 (A, C and E) correspond to the low-skill positions, while panels on the left hand side (B, D and F) correspond to high-skill positions. The top panels (A and B) correspond to the Searched positions. A comparison between these two panels indicates stark differences by skill level. Even before the firms had access to the tool (gray bins), there was more dispersion among the high-skill positions (Panel B) than among the low-skill positions (Panel A). This evidence is consistent with the idea of standardization, according to which employees in low-skill positions are seen as interchangable. Most importantly, the drop in salary dispersion is markedly sharper for low-skill positions than for high-skill positions. Among low-skill positions (Panel A), dispersion drops from 14.5 pp to 8.7 pp (p-value<0.001), corresponding to a 40% drop. For high-skill positions (Panel B), dispersion falls from 24.0 pp to 20.5 pp (p-value=0.021), corresponding to a drop of just 14.6%.[42] For the placebo tests, panels C through F of Figure 4 reproduce the analysis for Non-Searched and Non-Searchable positions. As expected, the differences in dispersion between post-onboarding and pre-onboarding salaries are always small in magnitude.

Appendix D.4 shows some additional results related to heterogeneity by skill. For the sake of brevity, the full event-study analysis for low-skill and high-skill positions is presented in Appendix D.4 – the conclusions remain unchanged. We also provide an alternative split of positions in terms of the heterogeneity by skills. For each position, we compute a measure of "market dispersion," namely the difference between the 90th and 10th percentiles of the market benchmarks (as shown in the benchmarking tool). Intuitively, if there is a lot of variation in salaries within a position, that would suggest a high variation in skills. The correlation between the skill classification and the market dispersion classification is high, but far from perfect.[43] Most importantly, the results for the heterogeneity by market dispersion are similar to, and consistent with, the results for the split by skill.

Our preferred interpretation of the heterogeneity by skill is that employers rely more on

[40] For 27% of observations there is no job zone classification available. In those cases, we impute education using data from Zippia.com on the share of employees with more than a high school degree.

[41] For more details and examples, see Appendix C.3.

[42] In Appendix D.4 we report the heterogeneity results using the difference-in-differences framework.

[43] Among the low-skill positions, 81% are classified as having low market dispersion; among the high-skill positions, 75% are classified as having high market dispersion.

salary benchmarking for low-skill positions than for high-skill ones. In high-skill roles, although the median market salary may serve as a starting point, other factors often become more significant as employers tailor the offer to the individual candidate. Another way to view this is that benchmarks are less informative for high-skill positions. In low-skill positions, candidates are seen as interchangeable, so the firm only needs to determine the median pay and offer that to every candidate. For high-skill positions, however, the information on the median pay may fall short from ideal. For example, rather than a single benchmark for "software developer," a firm might prefer two distinct benchmarks: one for "below-average software developer" and another for "above-average software developer," to be used depending on the perceived quality of the candidate. Indeed, employers may attempt to overcome this limitation of the benchmarking data by leveraging information on the distribution of salaries. For example, in hiring a below-average software developer, the firm might offer a salary at the 25th percentile of market salaries; for an above-average candidate, the offer might be at the 75th percentile. Unfortunately, we lack the necessary data to further investigate this hypothesis.[44]

In addition to the heterogeneity by skill, we explore other sources of heterogeneity. If the incentive to look up salary information is to keep up with the competition, we may expect the effects to be stronger in more competitive labor markets. We split the sample using measures of monopsonistic power created by other researchers (Azar et al., 2022). The results, which are reported in Appendix D.4, provide suggestive evidence that the effects of salary benchmarking are stronger in more competitive labor markets. Given the literature on the effects of negotiations on the gender pay gap (Bear, 2019), another natural question is whether there are differences in how salary benchmarking affects female and male employees. Appendix D.4 shows that we do not observe any significant differences by gender.

4 Effects on Average Salary and Retention

The above evidence suggests that the use of salary benchmarks has a significant effect on the salary dispersion. Next, we explore the effects on the average salary and on the retention rate.

[44] The tool usage data does not include details on which feature of the market salary distribution the firm examined (e.g., the 25th percentile), nor do we have insights into the firm's assessment of each new hire's quality.

4.1 Effects on Average Salary

To estimate the effects on the average salary, we use the same identification strategy as in Section 3 above. The key difference is that, instead of using the salary *dispersion* as the dependent variable, we use the salary *level*. The event-study results are presented in Figure 5. This figure is identical to Figure 3, except that the y-axis is the salary level (in logs). Figure 5 suggests that salary benchmarking has an insignificant effect on the average salary. Panel A of Figure 5 corresponds to the comparison between Searched (denoted in red dots) and Non-Searched (blue squares) positions. During the pre-onboarding period, the salary level was stable in both Searched and Non-Searched positions. In the post-onboarding period, both the Searched and Non-Searched positions continued at their pre-onboarding levels. Panel C of Figure 5 corresponds to the difference between the two series in Panel A. This difference-in-differences estimate suggests that there is no significant effect of salary benchmarking on the salary level. More precisely, access to the tool had an effect on the salary level that is virtually zero (-0.002 log points, or equivalent to an effect of just 0.2%),[45] and statistically insignificant (p-value=0.756).

Regarding the second identification strategy, Panel B of Figure 5 corresponds to the comparison between Searched positions (shown as red diamonds) and Non-Searchable positions (purple circles). Again, the salary level evolved similarly before and after the onboarding date, both for Searched and Non-Searchable positions. Panel D of corresponds to the difference between the two series in Panel B. This difference-in-differences comparison indicates that access to the tool had a slight positive effect on the average salary (0.017 log points, equivalent to a 1.7% increase), but the effect is imprecisely estimated and therefore statistically insignificant (p-value = 0.308). The similarity of the results across both identification strategies lends credence to the validity of the findings. Moreover, as reported in Appendix E, these results are robust to a wide range of alternative specifications.[46]

Given that the effects of benchmarking on salary dispersion are largely concentrated in low-skill positions, we can explore this same heterogeneity for salary levels. Figure 6 reproduces the results from Figure 5, but for the subsample of low-skill positions. The evidence points to a modest increase in average salary. Depending on whether the control group consists of Non-Searched positions (Panels A and C) or Non-Searchable positions (Panel B and D), the gains in average salary are estimated at 5.0% (p-value=0.014) and 6.7% (p-value=0.001), respectively. By comparison, in high-skill positions, there is no evidence of

[45] To be more precise, the effect is 0.2002% ($= 100 \cdot (exp(0.002) - 1)$). Since the approximation error is so small, in the remainder of the article we treat log-point effects and %-effects as interchangeable.

[46] A natural question is whether employees of one gender may have benefited more from benchmarking. Appendix E.4 shows that we do not find any evidence of significant gender differences.

significant effects on the salary level.[47] The effects on average salary are largely consistent with the non-parametric analysis presented in Section 3 above. For example, consider panel A of Figure 4, which shows the results for the low-skill positions. During the pre-onboarding period (gray bins), the distribution of salaries is skewed toward the left of the benchmark, meaning that firms were systematically under-paying employees. Thus, when salaries are compressed toward the benchmark, the compression from the bottom dominates, and the average salary goes up.[48]

4.2 Effects on Retention

It may seem puzzling at first that benchmarking leads firms to increase the average salary in low-skill positions. A possible interpretation is that employers raise salaries because, while it increases labor costs, it has some benefits, such as improving retention rates.[49] To test this hypothesis, we estimate the effects of salary benchmarking on retention of new hires.

Figure 7 is identical to Figure 6, except that, instead of the salary level, the dependent variable is the probability that the employee is still working at the firm 12 months after the hiring date.[50] Figure 7 suggests that, for low-skill positions, the gains in salaries were followed by an increase in retention rates. In contrast, for high-skill positions, for which we did not observe a significant change in salary levels, we did not observe a change in retention rates either.[51] The magnitude of retention gains is also worth discussing. Depending on whether the Non-Searched or Non-Searchable positions are used as control group (panels C and D of Figure 7, respectively), the gains in retention rates for low-skill positions are estimated at 6.6 pp (p-value=0.101) and 6.8 pp (p-value=0.029), respectively. These effects correspond to 16.1% and 16.6% of the baseline retention rates, respectively. For comparison, the corresponding gains in average salary are estimated at 5.0% and 6.7%, respectively. These effects on salary levels and retention imply labor supply elasticities of 3.22 ($= \frac{16.1}{5.0}$) and 2.48 ($= \frac{16.6}{6.7}$), respectively. These estimates are consistent with the range of estimates found in

[47] More precisely, the average salary drops by 2.9% and 1.6%, depending on the control group used, but these effects are statistically insignificant (p-values of 0.119 and 0.288, respectively). These results are reported in Appendix E.1.

[48] In turn, Figure 2 shows that when considering the whole sample, the compression from below and from above is similarly strong, so the negative and positive effects largely cancel each other out.

[49] One obvious expected benefit is that positions should be filled more quickly. Since we do not have data on job offers that were not accepted, unfortunately we cannot measure the effect on acceptance rates. There may be other expected benefits from higher salaries in addition to acceptance rates and retention, such as higher employee morale.

[50] It is worth noting that, for the employees hired in the later period (between March 2019 and March 2020) their 12-month horizon of retention will partially overlap with the COVID pandemic (beginning in April 2020).

[51] The results for the full sample and the high-skill subsample are presented in Appendix E.6.

the literature. For example, the meta-analysis of Sokolova and Sorensen (2021) reports a (weighted) mean of separation-based labor supply elasticities of 3.05.

5 A Model of Salary Benchmarks

Motivated by the evidence, we propose a simple model that can fit our main findings. We use the model as a lens to interpret our empirical results, and to explore implications for policy makers.

Salary benchmarks are aggregate statistics on the salary distribution. Standard labor market models assume that, in equilibrium, firms face no uncertainty about such statistics. For instance, the Diamond-Mortensen-Pissarides search model assumes that each firm-worker pair splits the surplus from the job in fixed proportions, with full information about the job's productivity and each side's outside option (Diamond, 1971; Mortensen and Pissarides, 1994). Similarly, the model of Postel-Vinay and Robin (2006) assumes that all wages and job offers are perfectly observed, and that firms set wages that best-respond to the steady-state wage distribution. As another example, the model of Roussille and Scuderi (2023) assumes that firms best respond to the distribution of competing offers for each worker, conditional on that worker's observable characteristics. In such models, firms already know everything that a benchmark may teach them. Direct survey questions of HR professionals suggest the contrary: salary benchmarks reveal new information in an environment ripe with information frictions.[52] Our empirical analysis indicates that firms adjust their behavior in response to benchmark information, challenging the assumption that firms are already aware of all the insights a benchmark could provide.

A model of benchmarks must allow for the possibility that firms are uncertain about the prevailing wage distribution. We study such a model, using techniques from auction theory, in particular, from Milgrom and Weber (1982). Our model shows that information frictions can cause wage dispersion, even in competitive markets. To isolate this new mechanism, we assume away standard causes of wage dispersion: Workers are identical, firms have identical amenities, firms have no monopsony power, and efficiency wages play no role. In our model, each firm faces a trade-off: offering a high wage means paying more, but offering a low wage risks leaving the position unfilled. Firms have different private information, and thus different beliefs about the population distribution of wages. This aggregate uncertainty leads to wage dispersion. Furthermore, in this simple model, aggregate uncertainty is the *only* cause of

[52]For example, HR professionals face hurdles to even access internal information about pay of similar employees, and they frequently do not have access to external offers – see Appendix B.3 for direct evidence from the SHRM survey.

wage dispersion. When a benchmark resolves the uncertainty, the wage dispersion vanishes.

There is a unit mass of firms. Each firm has a single open position; the firm's **value** is its marginal revenue from filling that position. Each firm knows its own value, but is uncertain about other firms' values. Formally, the state S is a random variable; given realization $S = s$, the population distribution of firm values is given by the CDF $F_s : [\underline{v}, \overline{v}] \to [0, 1]$. We assume that $0 \leq \underline{v}$, that $\overline{v} < \infty$, and that F_s is atomless and strictly increasing. The mass of workers is a random variable Q, with $\text{supp}(Q) \subseteq [0, 1]$. Each firm i observes Q (the supply of workers) and its own value V_i. If a firm with value V_i hires a worker at wage w, its payoff is $V_i - w$; firms that do not hire have payoff 0.

One can interpret the state as follows: each firm knows, based on internal data, how many product orders go unfilled while the position remains empty. But it does not know the situation at other firms. Thus, when one firm receives many excess orders (has a high V_i), it partly attributes this to idiosyncratic variation, and partly to an aggregate change of the whole population of firms, captured by the distribution F_s.[53]

All firms simultaneously make offers, and the firms that offer the highest Q wages hire workers. Formally, let G be the CDF of the wages offered in the population. Each firm is infinitesimal, so it treats G as exogenous. If there exists w such that $G(w) = 1 - Q$, then the firm hires if and only if it offers wage $w' \geq \inf\{\underline{w} : G(\underline{w}) = 1 - Q\}$. Otherwise, it must be that G jumps past $1 - Q$, that is, there exists $\underline{w}$ such that $G(\underline{w}) > 1 - Q$ and $G(w) < 1 - Q$ for all $w < \underline{w}$. In that case, if the firm's offer exceeds $\underline{w}$, then it hires for sure, and if its offer is exactly $\underline{w}$, then we ration by breaking ties randomly.[54]

A function $\eta : [\underline{v}, \overline{v}] \times \text{supp}(Q) \to \mathbb{R}_{\geq 0}$ is a **no-benchmark equilibrium** if for every (v, q), offering wage $\eta(v, q)$ maximizes the firm's expected payoff conditional on $(V_i, Q) = (v, q)$, when all other firms behave according to η.[55]

We define the **cutoff** to be the random variable $C \equiv F_S^{-1}(1 - Q)$; this is the $(1 - Q)$-quantile of the value distribution in state S. We assume that (V_i, C, Q) have the same joint distribution for all i. We denote the conditional cumulative distribution function $H(c \mid v, q) \equiv P(C \leq c \mid V_i = v, Q = q)$, with the corresponding density $h(c \mid v, q)$.

We assume that the random variables (V_i, C) are **affiliated** conditional on Q.[56] That is, let $m(v, c \mid q)$ be the joint density of V_i and C conditional on $Q = q$, and let $\vee$ denote the component-wise maximum and let $\wedge$ denote the component-wise minimum. We assume that

[53]In our model, firms have rational expectations based on their own private information. The idea that firms do not perfectly observe aggregate conditions arises also in the Lucas islands model (Lucas, 1972).

[54]We specify this for completeness; rationing does not arise in the equilibria we characterize.

[55]Throughout we restrict attention to functions η that are measurable with respect to the first argument.

[56]Affiliation is a standard technical condition. For a textbook treatment, see Krishna (2009), p. 285-288. In our model, affiliation ensures the existence of monotone pure-strategy equilibria. In Section Appendix B.5, we provide empirical support for the assumption of affiliated firm values in the labor market context.

for all (v, c), all (v', c'), and all q, we have

$$m(v, c \mid q)m(v', c' \mid q) \leq m((v, c) \vee (v', c') \mid q)m((v, c) \wedge (v', c') \mid q). \tag{4}$$

This would be implied, for example, if Q and S are independent, if $\text{supp}(S)$ is ordered, and F_s has the property of monotone likelihood ratio with respect to s (Milgrom and Weber, 1982, p. 1099). Affiliation implies non-negative correlation. For multivariate normal distributions, affiliation is equivalent to non-negative correlation.

If the function η is strictly increasing in the firm's value, then a firm with value $V_i = v$ facing supply $Q = q$ hires with probability $P(C \leq v \mid V_i = v, Q = q) = H(v \mid v, q)$. Let us define $\tau_q \equiv \inf\{v : H(v \mid v, q) > 0\}$. We assume that for all q and all $v > \tau_q$, there exists $\epsilon > 0$ such that $H(v \mid v + \epsilon, q) > 0$. This assumption ensures that there exists an equilibrium $\eta(v, q)$ that is continuous in v.

Every no-benchmark equilibrium involves wage dispersion. The law of one price does not hold, even though there are many firms and many identical workers. We now state this formally.

Theorem 5.1. *For any no-benchmark equilibrium η and any q, the function $\eta(v, q)$ is not constant in v for $v > \tau_q$.*

The intuition for Theorem 5.1 is that if, at equilibrium, all workers are hired at the same wage, then high-value firms would make profits upon hiring, but would sometimes fail to hire because of ties. Such firms could profitably deviate by slightly raising their offer, a contradiction. The proof is given in Appendix A.1.

We characterize a monotone no-benchmark equilibrium, leaning on techniques from Milgrom and Weber (1982). Firms with higher values offer higher wages, because when one firm has a high value, it infers that *other* firms are also likely to have high values, and thus that other firms will offer high wages.[57]

Theorem 5.2. *The function w^* is a no-benchmark equilibrium, where:*

$$w^*(v, q) \equiv \begin{cases} v & v \leq \tau_q \\ \int_{\tau_q}^{v} \alpha dL(\alpha \mid v, q) & v > \tau_q \end{cases} \tag{5}$$

for

$$L(\alpha \mid v, q) \equiv \exp\left(-\int_{\alpha}^{v} \frac{h(\beta \mid \beta, q)}{H(\beta \mid \beta, q)} d\beta\right).[58] \tag{6}$$

[57] Jäger et al. (2024) proposed a model in which workers have biased beliefs about the wage distribution, anchored on their current wage. In contrast, in our model it is firms that are uncertain and their beliefs are Bayesian posteriors derived from a common prior.

[58] We adopt the convention that $\frac{h(v|v,q)}{H(v|v,q)} = 0$ if v is not in the support of the conditional distribution of C.

Moreover, $w^*(v, q)$ so defined is continuous in v, increasing in v, and we have $w^*(v, q) \leq v$.

The proof is given in Appendix A.2.

The only uncertainty firms face concerns the aggregate demand for workers, as captured by the state S. If S is persistent over time, then under mild assumptions observing the distribution of accepted offers in an earlier period suffices to identify S. Thus, we model access to a salary benchmark as learning the state with certainty.

Suppose that one firm covertly observes S, while all other firms continue to offer wages according to the no-benchmark equilibrium w^*. The informed firm knows the cutoff C, because Q is public and C depends only on S and Q. Figure 8 illustrates the function w^*, fixing the realization of the worker supply $Q = q$. Suppose that the cutoff realization is $C = c$; then if the informed firm's value is in the interval $[w^*(c, q), c)$, it will fail to hire at its original offer but would be willing to hire at the marginal firm's offer, $w^*(c, q)$.

Theorem 5.3. *For arbitrary realizations $C = c$ and $Q = q$, it is a best-response for the informed firm i:*

1. *to offer a wage too low to be accepted if $V_i < w^*(c, q)$,*
2. *to raise its offer from $w^*(V_i, q)$ to $w^*(c, q)$ if $w^*(c, q) \leq V_i \leq c$,*
3. *and to lower its offer from $w^*(V_i, q)$ to $w^*(c, q)$ if $c < V_i$.*

Proof. By inspection. □

Theorem 5.3 indicates that when a firm compresses the wage offers in response to a benchmark, that is not necessarily an indication of monopsony power. That compression arises even in our model with many firms, each of which is effectively a price taker.

Clearly, when a firm responds to a benchmark, that is evidence that the market was not (originally) at a full-information equilibrium. But Theorem 5.2 and Theorem 5.3 together show that such responses are at least consistent with *incomplete-information* equilibrium, when firms face aggregate uncertainty.

Theorem 5.3 predicts that a firm that gains access to a benchmark will compress its offer to an atom at $w^*(C, Q)$, increasing offers that would otherwise be too low and lowering offers that would otherwise be too high. On the one hand, this prediction could rationalize the empirical finding of bunching at the median benchmark. On the other hand, there is no theoretical reason why this atom should be exactly at the median of past accepted offers. Formally, suppose that the state S is persistent, and we divide firms randomly between two sub-markets at times $t = 1, 2$, with worker supply Q_1 and Q_2. At time 2, observation of the distribution of the past accepted offers will suffice (under mild assumptions) to identify

S, and thus the relevant cutoff $C_2 \equiv F_S^{-1}(Q_2)$, but the informed firm's offer $w^*(C_2, Q_2)$ could be above or below the median accepted offer at time 1. Compression to the median in particular might be due to the salience of the median in the benchmark's user interface, or to a behavioral heuristic.[59]

Theorem 5.2 and Theorem 5.3 imply that the benchmark causes a form of compression in *accepted* offers. That is, even when we condition on the firm's offer being accepted, the lower end of the distribution does not fall, and the upper end of the distribution does not rise. We now state this formally.

Corollary 5.4. *For any worker supply q, we have*

$$\begin{aligned} &\lim_{v\downarrow\tau_q} E\left[w^*(C,Q) \mid w^*(C,Q) \le V_i, (V_i,Q) = (v,q)\right] \\ &- \lim_{v\downarrow\tau_q} E\left[w^*(V_i,Q) \mid C \le V_i, (V_i,Q) = (v,q)\right] \ge 0 \end{aligned} \tag{7}$$

and

$$\begin{aligned} &\lim_{v\uparrow\overline{v}} E\left[w^*(C,Q) \mid w^*(C,Q) \le V_i, (V_i,Q) = (v,q)\right] \\ &- \lim_{v\uparrow\overline{v}} E\left[w^*(V_i,Q) \mid C \le V_i, (V_i,Q) = (v,q)\right] \le 0. \end{aligned} \tag{8}$$

Proof. Inequality (7) follows by $w^*(v,q)$ continuous in v and increasing in v, and $w^*(\tau_q, q) = \tau_q$. Inequality (8) follows by $w^*(v,q)$ continuous in v and increasing in v, and $w^*(v,q) \le v$. □

5.1 Extensions

We discuss some simple extensions of the model, to gain a clearer understanding of some additional results from the empirical analysis. A first relevant result is that salaries get more compressed towards the median benchmark in low-skill positions than in high-skill positions (Section 3.6). Our preferred interpretation is that for low-skill jobs workers may be more homogeneous within each job title, whereas high-skill jobs might nest several distinct kinds of employees within a job title—for instance, software developers specialize in different programming languages. Formally, suppose that there is a finite set of categories Φ, and a unit mass of firms for each category. There are category-specific value distributions $(F_s^{\phi})_{\phi\in\Phi}$ and worker supplies $(Q^{\phi})_{\phi\in\Phi}$. Then Theorem 5.2 and Theorem 5.3 apply to each category separately, and offers from informed firms can exhibit dispersion within a job title.[60]

[59] For instance, the median minimizes the sum of absolute differences between the new hire's wage and the wages recorded in the benchmark tool. It is thus the inequity-minimizing wage in the model of Fehr and Schmidt (1999) with parameters $\alpha_i = \beta_i > 0$.

[60] While this is our preferred explanation, there may be other explanations for the heterogeneity by skill. For instance, some high-skill positions may pay efficiency wages due to non-contractible aspects of performance.

Another relevant result is that the effects of benchmarking on salary dispersion are stronger in more competitive labor markets (Appendix D.4). The stark compression in Theorem 5.3 arises in part because the informed firm has no market power; it fills the position if and only if its offer is at least $w^*(C,Q)$. One way to introduce market power is to assume that the informed firm has an amenity shock A_i, identical across workers, so that the firm hires if and only if its offer exceeds $w^*(C,Q) - A_i$. Then the informed firm chooses w to maximize

$$(V_i - w)H(w - w^*(C,Q)) \tag{9}$$

where H is the CDF of the amenity shock. Suppose H is continuously differentiable and has support on $(-\infty, +\infty)$. Then the objective function (9) is continuously differentiable, and its derivative with respect to w is strictly increasing in V_i. It follows that the firm's optimal offer is strictly increasing in V_i wherever it has an interior solution (Edlin and Shannon, 1998). In this sense, market power can lead to wage dispersion even when benchmarks resolve aggregate uncertainty.

A third key finding is that salary benchmarks raise the average salary and retention in low-skill positions (Section 4.2). We have not explicitly modeled employee retention. However, the payoffs in our model are related to the payoffs of the search process studied by Burdett and Mortensen (1998). Firms choose permanent wage offers and workers search by sampling randomly from the set of offers. Workers search even while employed, so firms have retention concerns. As the sampling rate goes to infinity, the payoffs from that dynamic process converge to those in our static model. On that interpretation, an uninformed firm with value V_i will (eventually) hire a worker whom they retain permanently if and only if $C \leq V_i$, whereas an informed firm with value V_i will do the same if and only if $w^*(C,Q) \leq V_i$. Since we have $w^*(C,Q) \leq C$, the benchmark raises the retention rate.

Next, we discuss some additional extensions of the model that capture more realistic features of the institutional context. In practice, we observe that there are some sources of benchmarking data that are less precise but freely available to all firms (Section 2). This can be incorporated to the model as follows. Let us represent those benchmarks as a random variable B that is partially informative about the state S. All the firms have common knowledge of B. If our technical conditions hold conditional on B, then all the results extend straightforwardly.[61]

Our baseline model assumes that firms offer identical amenities. However, in real labor markets, there can be large differences in amenities between firms. Suppose that each firm has a firm-specific amenity A_i that is exogenous and known to that firm and to all workers.

[61] That is, (V_i, C) are affiliated conditional on Q and B, and there exists a distribution $H(c \mid v, q, b) \equiv P(C \leq c \mid V_i = v, Q = q, B = b)$ with corresponding density, and so on.

The firms that hire are those that offer the highest total remuneration (i.e., wages plus amenities). Let us define a new random variable $\overline{V}_i \equiv V_i + A_i$, equal to the revenue from filling the position plus the amenity. Moreover, assume that $\overline{V}_i$ is a sufficient statistic for the firm's private information about aggregate conditions, and satisfies the same technical assumptions as V_i in the baseline model, with $\overline{C}$ defined analogously. Then the function w^* can be reinterpreted as the total remuneration offered in equilibrium, as a function of $\overline{V}_i$ and Q.[62] On that interpretation, Theorem 5.3 predicts that learning the state causes compression in total remuneration.

5.2 Equilibrium Effects

Policy makers may be especially interested in the equilibrium effects of salary benchmarking. Our empirical findings relate to partial equilibrium effects, namely, when a single firm gains covert access to the benchmark. However, we can employ the model to examine the equilibrium effects, at least theoretically.

Suppose that the state S is common knowledge between firms, and hence the cutoff C is common knowledge. A **benchmark equilibrium** is a function $\eta : [\underline{v}, \overline{v}] \times \text{supp}(Q) \times \text{supp}(C) \to \mathbb{R}_{\geq 0}$ such that for each (v, q, c), the offer of a wage $\eta(v, q, c)$ maximizes firm i's expected payoff conditional on $(V_i, Q, C) = (v, q, c)$, when all other firms behave according to η.

In any benchmark equilibrium, all workers must be hired at the same wage and workers are hired by firms with values above the cutoff. Thus, the prevailing wage is equal to C, so that the marginal firm is indifferent between hiring and not hiring. For instance, it is a benchmark equilibrium to set

$$\tilde{w}(v, q, c) = \begin{cases} v & v < c \\ c & v \geq c \end{cases}. \tag{10}$$

Moreover, every benchmark equilibrium is outcome-equivalent to $\tilde{w}$. Thus, we see compression even in equilibrium, now to an atom at $C \geq w^*(C, Q)$.

Observe that under the benchmark equilibrium $\tilde{w}$, each firm chooses an offer that best corresponds to the *realized* distribution of offers. Furthermore, the realized distribution of *wages* suffices to derive that best response, so $\tilde{w}$ captures the idea that the wage distribution is common knowledge.

Does the benchmark raise wages in equilibrium? Under the benchmark equilibrium $\tilde{w}$,

[62] Observe that the payoff to a firm with value V_i and amenity A_i of hiring at total remuneration $w + A_i$ is $V_i - w = V_i + A_i - w - A_i = \overline{V}_i - (w + A_i)$. In particular, firms with the same $\overline{V}_i$ but different A_i will offer different wages, but the same total remuneration.

firms make offers that are sometimes higher than and sometimes lower than under the no-benchmark equilibrium w^*. However, an argument using the linkage principle enables us to sign the expected change in wages (Milgrom and Weber, 1982), as we now state.

Theorem 5.5. *For any worker supply q and any firm value $v > \tau_q$, we have*

$$w^*(v,q) \leq E_C\left[\tilde{w}(v,q,C) \mid C \leq v, V_i = v, Q = q\right]. \tag{11}$$

As a corollary, expected wages are higher under the benchmark equilibrium than under the no-benchmark equilibrium, that is

$$E_{V_i,Q,C}\left[w^*(V_i,Q) \mid C \leq V_i\right] \leq E_{V_i,Q,C}\left[\tilde{w}(V_i,Q,C) \mid C \leq V_i\right]. \tag{12}$$

The proof is in Appendix A.3.

To build intuition for Theorem 5.5, let Acme be a firm with some arbitrary value $V_i = v > \tau_q$, facing the no-benchmark equilibrium w^*. To hire, Acme does not need to make an offer that exceeds C, the marginal firm's value; it only needs to beat the marginal firm's offer, which is $w^*(C,q)$. In equilibrium, Acme hires if and only if $C \leq v$. By affiliation, whenever $C \leq v$, the marginal firm believes that the demand for workers is relatively weak, compared to the belief of Acme. This drives down the marginal firm's offer, to Acme's benefit. Thus, Acme enjoys information rents; its expected profit under w^* exceeds its expected contribution to social surplus, which is $E_{V_i,C}\left[\max\{V_i - C, 0\} \mid V_i = v, Q = q\right]$.[63] In this way, aggregate uncertainty blunts labor-market competition between firms.

In contrast, the benchmark makes the cutoff C common knowledge, leading to intense wage competition between firms with values near the cutoff. If the marginal firm makes an offer strictly below C, then a firm with value just below C could profitably deviate to hire workers. Thus, under the benchmark equilibrium $\tilde{w}$, the prevailing wage is equal to C, and each firm's profit is equal to $\max\{V_i - C, 0\}$, its contribution to social surplus. Firms no longer have information rents, so expected firm surplus is lower under $\tilde{w}$ than under w^*. The total surplus is equal under $\tilde{w}$ and w^*, because the same set of firms hire workers. It follows that the expected worker surplus is higher under $\tilde{w}$ than under w^*, and thus that expected wages are higher as well.[64]

[63]Observe that C is the opportunity cost to society of Acme hiring a worker.

[64]This accounting exercise implicitly assumes that workers are risk-neutral. But recall that w^* results in wage dispersion, while $\tilde{w}$ does not. Thus, Theorem 5.5 implies that wages under $\tilde{w}$ second-order stochastically dominate wages under w^*, and therefore are preferred by any worker with an increasing concave utility function. Allowing for worker risk aversion yields another argument in favor of benchmarks, namely that they increase worker surplus by reducing wage uncertainty.

6 Conclusion

While U.S. legislation forbids employers from exchanging compensation information directly, it permits the use of aggregated data through third parties, a method known as salary benchmarking. In partnership with the leading provider of payroll services and salary benchmarks, we explore the introduction of a novel benchmarking tool. Employing an event-study methodology, we present evidence that access to this tool significantly influences firm behavior. Notably, the salaries of new hires are more compressed toward the median market benchmark displayed in the tool. This effect on salary dispersion is particularly strong in low-skill positions.

Using a theoretical model, we discuss implications for the study of labor markets. Standard models of the labor market assume that each firm knows how much other firms are paying, at least in the aggregate. Our evidence is inconsistent with this full-information assumption, because firms substantially changed their behavior in response to information on market pay. These results suggest that we need models of labor markets with richer information assumptions that allow for aggregate uncertainty about the salaries paid by other firms. As a step in that direction, we proposed a competitive labor market model with aggregate uncertainty about the demand for workers. Our model highlights a novel mechanism for salary dispersion. In equilibrium, firms pay different salaries because they have different posterior beliefs about the distribution of salaries.

Furthermore, our empirical and theoretical analysis has policy implications. In the United States, salary benchmarks are regulated by the Department of Justice (DOJ) and the Federal Trade Commission (FTC). From 1993 to 2011, the DOJ and the FTC released a series of antitrust policy statements that created a "safety zone" for salary benchmarks. That is, agencies would not challenge benchmarks managed by a third party, provided that the data were anonymized, sufficiently aggregated, and more than three months old (Bloom, 2014). In 2021, the Biden administration issued an executive order that urged the DOJ and the FTC to "prevent employers from collaborating to suppress wages or reduce benefits by sharing wage and benefit information with one another" (White House, 2021). In 2023, both agencies rescinded the policy statements that created the safety zone, stating that they were "overly permissive on certain subjects, such as information sharing" (DOJ, 2023; FTC, 2023).

Our labor market model indicates that, at equilibrium, salary benchmarks can lead to higher pay, as resolving uncertainty prompts firms near the hiring margin to compete more fiercely with one another. Thus, our model provides a formal analysis of the pro-competitive argument for salary benchmarks highlighted by policymakers. Our empirical findings cannot directly address the equilibrium effects of salary benchmarking because we estimate the partial equilibrium effect of providing benchmark information to an additional firm. Bearing this limitation in mind, we do not find evidence that salary benchmarking suppresses wages. Access to the benchmark information does not lower the average salary of new hires. On the contrary, for low-skill positions, we observe an increase in the average salary and retention rate.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High Wage Workers and High Wage Firms. *Econometrica 67*(2), 251–333.

Adler, L. (2020). What's a Job Candidate Worth? Status and Evaluation in Pay-Setting Process. *Working Paper*.

Adler, L. (2022). From the Job's Worth to the Person's Price: Changes in Pay-Setting Practices since 1950. *Doctoral dissertation, Harvard University Graduate School of Arts and Sciences.*.

Azar, J., I. Marinescu, and M. Steinbaum (2022). Labor market concentration. *Journal of Human Resources 57*(S), S167–S199.

Baker, M., Y. Halberstam, K. Kroft, A. Mas, and D. Messacar (2023). Pay Transparency and the Gender Gap. *American Economic Journal: Applied Economics, forthcoming*.

Bear, J. (2019). Gender gaps in negotiation: Implications for individuals and organizations. Forthcoming.

Bennedsen, M., E. Simintzi, M. Tsoutsoura, and D. Wolfenzon (2022). Do Firms Respond to Gender Pay Gap Transparency? *Journal of Finance 77*, 2051–2091.

Berger, L. A. and D. Berger (2008). *The Compensation Handbook*. New York: McGraw-Hill.

Blankmeyer, E., J. LeSage, J. Stutzman, K. Knox, and R. Pace (2011). Peer-group dependence in salary benchmarking: a statistical model. *Managerial and Decision Economics 32*(2), 91–104.

Bloom, M. (2014). Information exchange: Be reasonable. Federal Trade Commission. Retrieved from https://www.ftc.gov/news-events/blogs/competition-matters/2014/12/information-exchange-be-reasonable.

Breza, E., S. Kaur, and Y. Shamdasani (2018). The Morale Effects of Pay Inequality. *The Quarterly Journal of Economics*.

Bruckner, A. and C. Goffman (1976). Differentiability through change of variables. *Proceedings of the American Mathematical Society 61*(2), 235–241.

Burdett, K. and D. T. Mortensen (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 257–273.

Caldwell, S. and O. Danieli (2021). Outside Options in the Labor Market. *Working Paper*.

Caldwell, S. and N. Harmon (2018). Outside Options, Bargaining and Wages: Evidence from Coworker Networks. *Working Paper*.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics 36*(S1), S13–S70.

Card, D., A. Mas, E. Moretti, and E. Saez (2012). Inequality at Work: The Effect of Peer Salaries on Job Satisfaction. *American Economic Review 102*(6), 2981–3003.

Collins, L., D. Fineman, and A. Tsuchida (2017). People analytics: Recalculating the route. *Rewriting the rules for the digital age: 2017 Deloitte Global Human Capital Trends*.

Cullen, Z. and B. Pakzad-Hurson (2023). Equilibrium Effects of Pay Transparency in a Simple Labor Market. *Econometrica 91*, 765–802.

Cullen, Z. and R. Perez-Truglia (2022). How Much Does Your Boss Make? The Effects of Salary Comparisons. *Journal of Political Economy 30*(3), 766–822.

Cullen, Z. and R. Perez-Truglia (2023). The Salary Taboo: Privacy Norms and the Diffusion of Information. *Journal of Public Economics 222*, 104890.

Davenport, T. (2019). Is HR the Most Analytics-Driven Function? *Harvard Business Review Digital Article*.

Davenport, T. and J. Shapiro (2010). Competing on talent analytics. *Harvard Business Review 88*(10), 52–58.

DellaVigna, S. and E. Linos (2022). Rcts to scale: Comprehensive evidence from two nudge units. *Econometrica 90*(1), 81–116.

Derenoncourt, E., C. Noelke, D. Weil, and B. Taska (2021). Spillover Effects from Voluntary Employer Minimum Wages. *NBER Working Paper No. 29425*.

Diamond, P. A. (1971). A model of price adjustment. *Journal of economic theory 3*(2), 156–168.

DOJ (2023). Justice Department Withdraws Outdated Enforcement Policy Statements. *Press Release, February 3, 2023*.

Dube, A., L. Giuliano, and J. Leonard (2019). Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior. *American Economic Review 109*(2), 620–663.

Duchini, E., S. Simion, and A. Turrell (2022). Pay Transparency and Cracks in the Glass Ceiling. *CAGE Working Paper No. 482.*.

Duffie, D., P. Dworczak, and H. Zhu (2017). Benchmarks in Search Markets. *The Journal of Finance 72*(5), 1983–2044.

Edlin, A. S. and C. Shannon (1998). Strict monotonicity in comparative statics. *Journal of Economic Theory 81*(1), 201–219.

Faulkender, M. and J. Yang (2010). Inside the black box: The role and composition of compensation peer groups. *Journal of Financial Economics 96*(2), 257–270.

Fehr, E. and K. M. Schmidt (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics 114*(3), 817–868.

FTC (2023). Federal Trade Commission Withdraws Health Care Enforcement Policy Statements. *Press Release, July 14, 2023*.

Grennan, M. and A. Swanson (2020). Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining. *Journal of Political Economy 128*(4), 1234–1268.

Grigsby, J., E. Hurst, and A. Yildirmaz (2021). Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data. *American Economic Review 111*(2), 428–471.

Hall, R. and A. Krueger (2012, 10). Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-the-Job Search. *American Economic Journal: Macroeconomics 4*(4), 56–67.

Hastings, J. and J. M. Shapiro (2018). How Are SNAP Benefits Spent? Evidence from a Retail Panel. *American Economic Review 108*(12), 3493–3540.

Hazell, J., C. Patterson, H. Sarsons, and B. Taska (2021). National Wage Setting. *Working Paper.*.

Jäger, S., C. Roth, N. Roussille, and B. Schoefer (2024). Worker Beliefs About Outside Options. *Quarterly Journal of Economics*. forthcoming.

Kaur, S. (2019). Nominal Wage Rigidity in Village Labor Markets. *American Economic Review 109*(10), 3585–3616.

Kline, P., N. Petkova, H. Williams, and O. Zidar (2019, 03). Who Profits from Patents? Rent-Sharing at Innovative Firms*. *The Quarterly Journal of Economics 134*(3), 1343–1404.

Krishna, V. (2009). *Auction theory.* Academic press.

Krueger, A. B. and L. H. Summers (1988). Efficiency wages and the inter-industry wage structure. *Econometrica*, 259–293.

Larcker, D., C. McClure, and C. Zhu (2019). Peer Group Choice and Chief Executive Officer Compensation. *Stanford University, Graduate School of Business Working Paper No. 3767*.

Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of economic theory 4*(2), 103–124.

Mas, A. (2016). Does Disclosure affect CEO Pay Setting? Evidence from the Passage of the 1934 Securities and Exchange Act. *Working Paper*.

Mas, A. (2017). Does Transparency Lead to Pay Compression? *Journal of Political Economy 125*(5), 1683–1721.

Milgrom, P. R. and R. J. Weber (1982). A Theory of Auctions and Competitive Bidding. *Econometrica 50*(5), 1089–1122.

Mortensen, D. T. and C. A. Pissarides (1994). Job creation and job destruction in the theory of unemployment. *The review of economic studies 61*(3), 397–415.

Murphy, K. M. and R. H. Topel (1990). Efficiency wages reconsidered: Theory and evidence. In *Advances in the Theory and Measurement of Unemployment*, pp. 204–240. Springer.

PayScale (2021). 2021 Compensation Best Practices Report. Technical report.

Perez-Truglia, R. (2020). The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment. *American Economic Review 110*, 1019–54.

Postel-Vinay, F. and J.-M. Robin (2006). Equilibrium Wage Dispersion with Worker and Employer Heterogeneity. *Econometrica 70*(6), 2295–2350.

Rosen, S. (1986). The theory of equalizing differences. *Handbook of labor economics 1*, 641–692.

Roussille, N. (2023). The Central Role of the Ask Gap in Gender Pay Inequality. *Working Paper*.

Roussille, N. and B. Scuderi (2023). Bidding for talent: A test of conduct in a high-wage labor market.

Schiemann, W. A., J. H. Seibert, and M. H. Blankenship (2018). Putting human capital analytics to work: Predicting and driving business success. *Human Resource Management 57*(3), 795–807.

Securities and Exchange Commission (2006). SEC final rules 33-8732a, Item 402(b)(2)(xiv).

Sokolova, A. and T. Sorensen (2021). Monopsony in Labor Markets: A Meta-Analysis. *Industrial & labor relations review 74*(1), 27–55.

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. Von Wachter (2019). Firming up inequality. *Quarterly Journal of Economics 134*(1), 1–50.

Thom, M. and T. Reilly (2015). Compensation Benchmarking Practices in Large U.S. Local Governments. *Public Personnel Management 44*(3), 340–355.

White House (2021). Fact Sheet: Executive Order on Promoting Competition in the American Economy. *Statements and Releases from the White House, July 9, 2021*.

Zeuch, M. (2016). *Handbook of Human Resources Management.* Berlin: Springer.

Figure 1: Screenshot of the Salary Benchmarking Tool

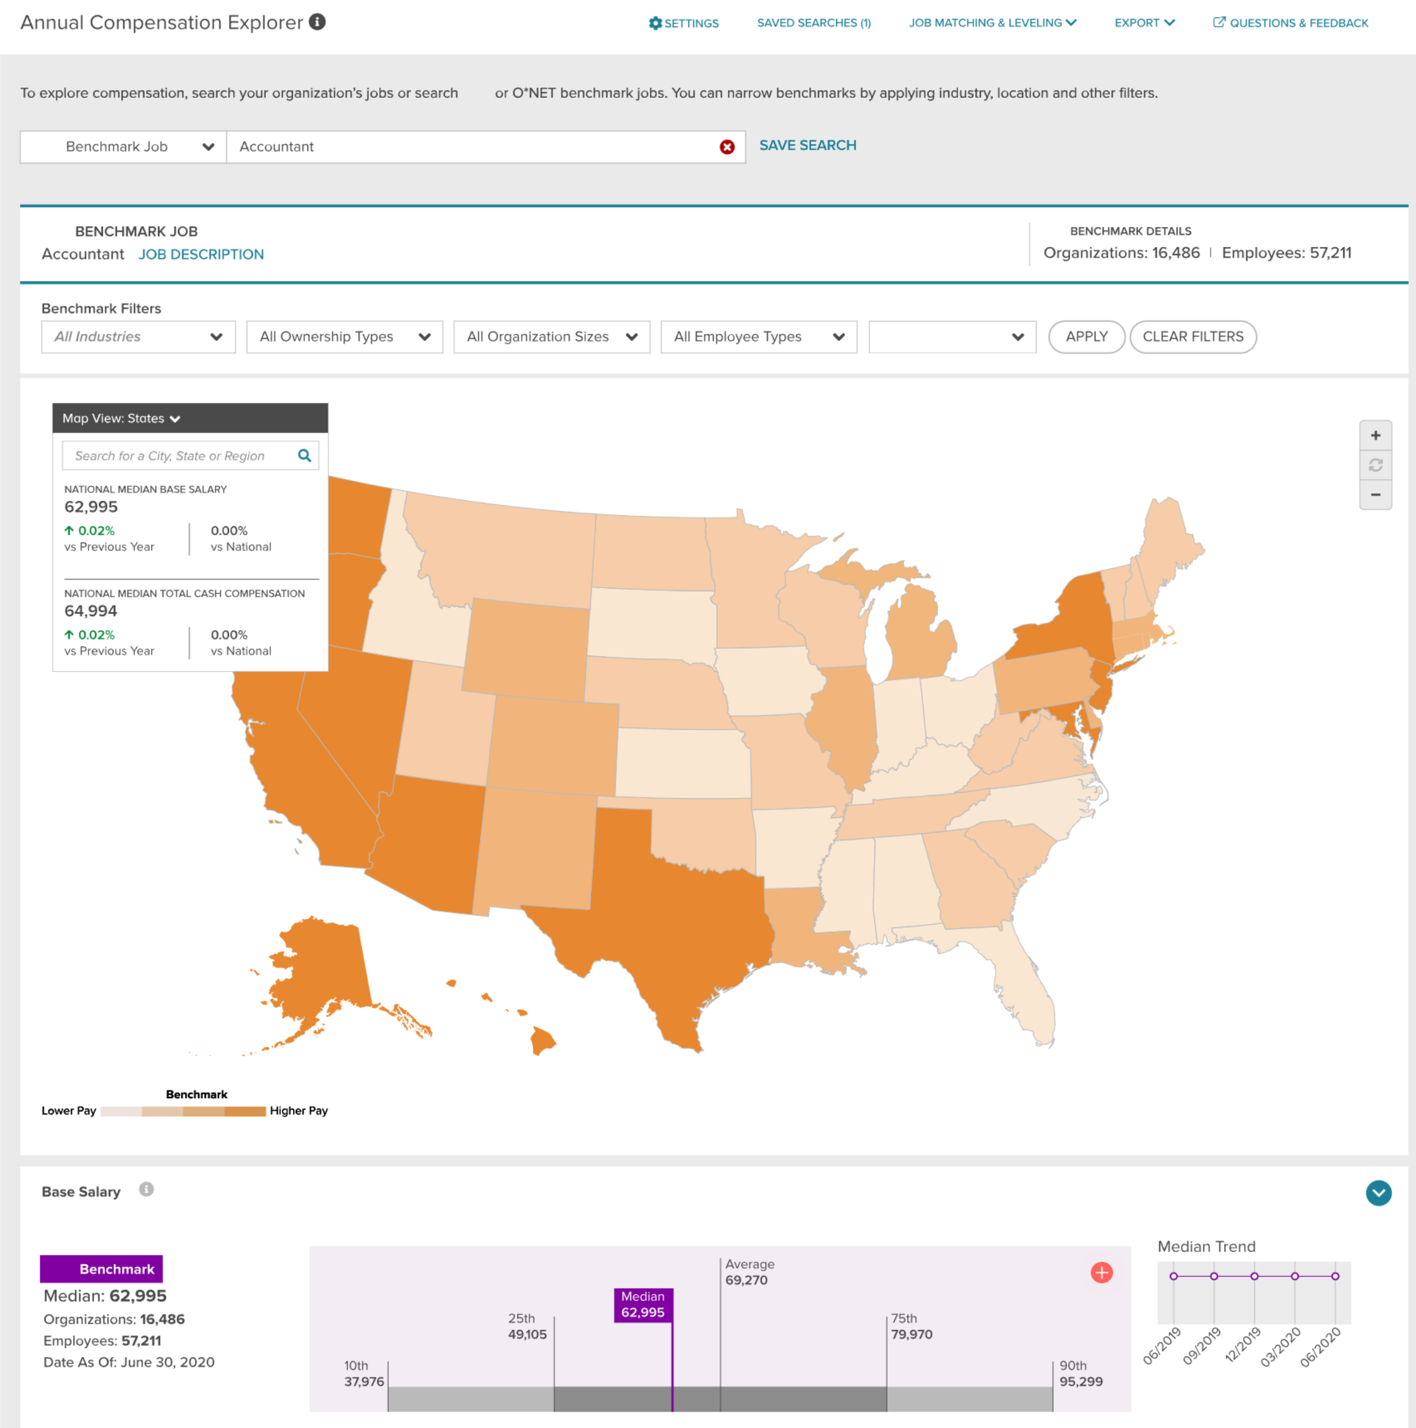

Notes: This is a screenshot of the pay benchmarking tool. It has been slightly altered to conceal the identity of the firm. This is the top of the screen. If you scroll down, you can see panels similar to the bottom panel titled *Base Salary* but for *Bonus*, *Overtime*, and *Total Compensation*.

Figure 2: The Effects of Benchmarking on Salary Dispersion Around the Benchmark: Non-Parametric Analysis

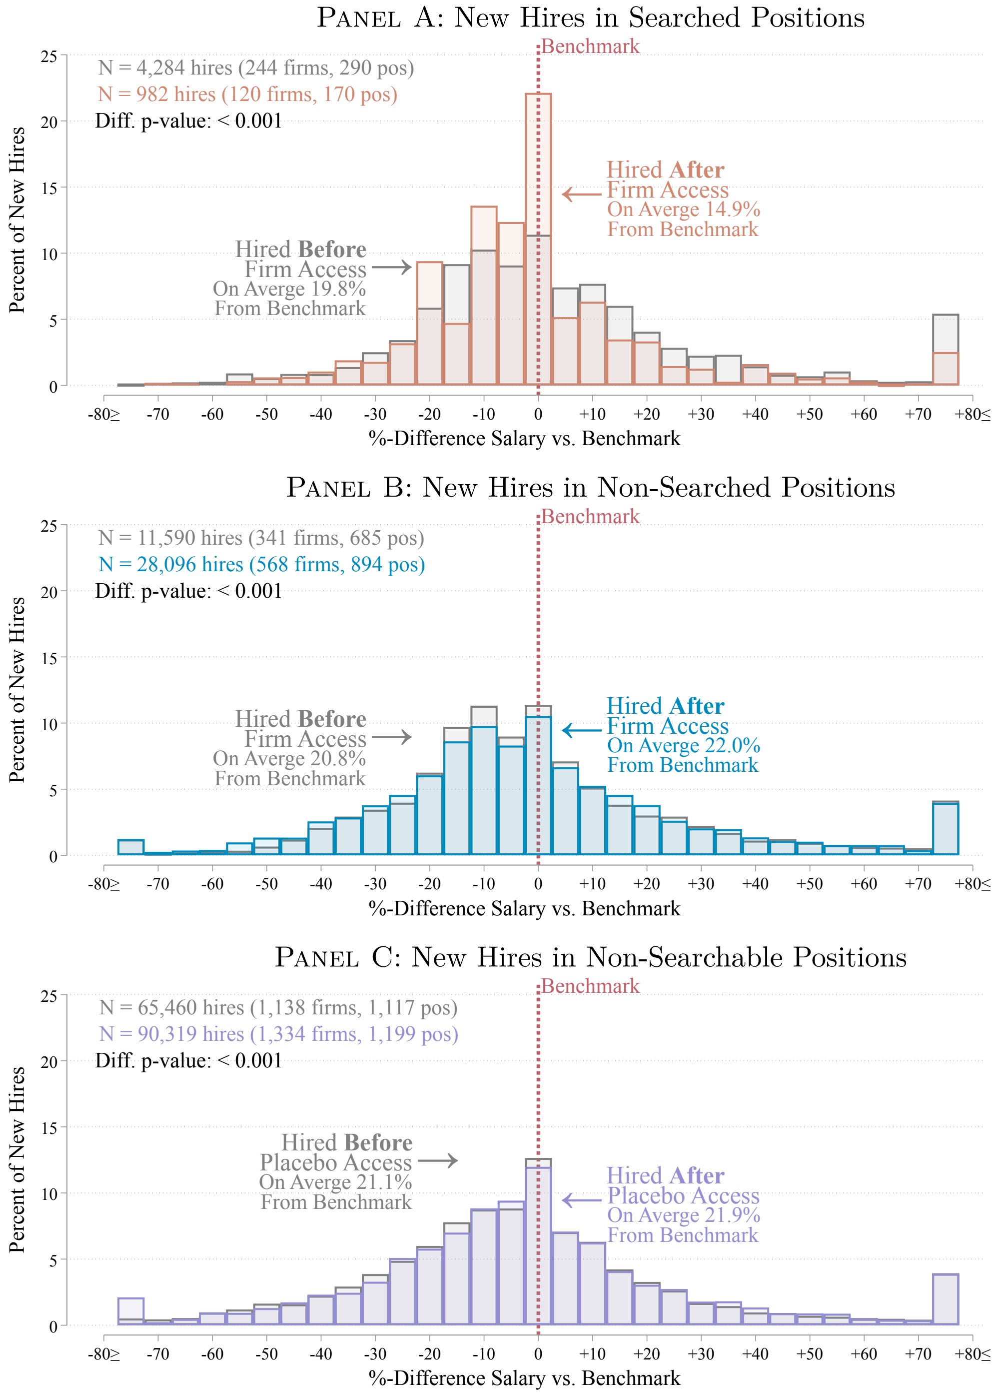

Notes: Histograms of the starting base salary relative to the corresponding external benchmark (winsorized at ± 75%). Each panel corresponds to a different set of positions: Panel A for Searched positions (i.e., positions in firms with access to the benchmark tool that are eventually searched for by the firm), Panel B for Non-Searched positions (i.e., positions in firms with access to the benchmark tool that are not eventually searched for by the firm), and Panel C for Non-Searchable positions (i.e., positions in firms without access to the benchmark tool). In each panel, the solid and hollow bins correspond to the observations before and after the firm gains access to the benchmark tool, respectively (and in Panel C, that date corresponds to the "hypothetical" onboarding date assigned to the firm that never gains access to the tool).

Figure 3: Event-Study Analysis: Effects on Pay Dispersion Around the Benchmark

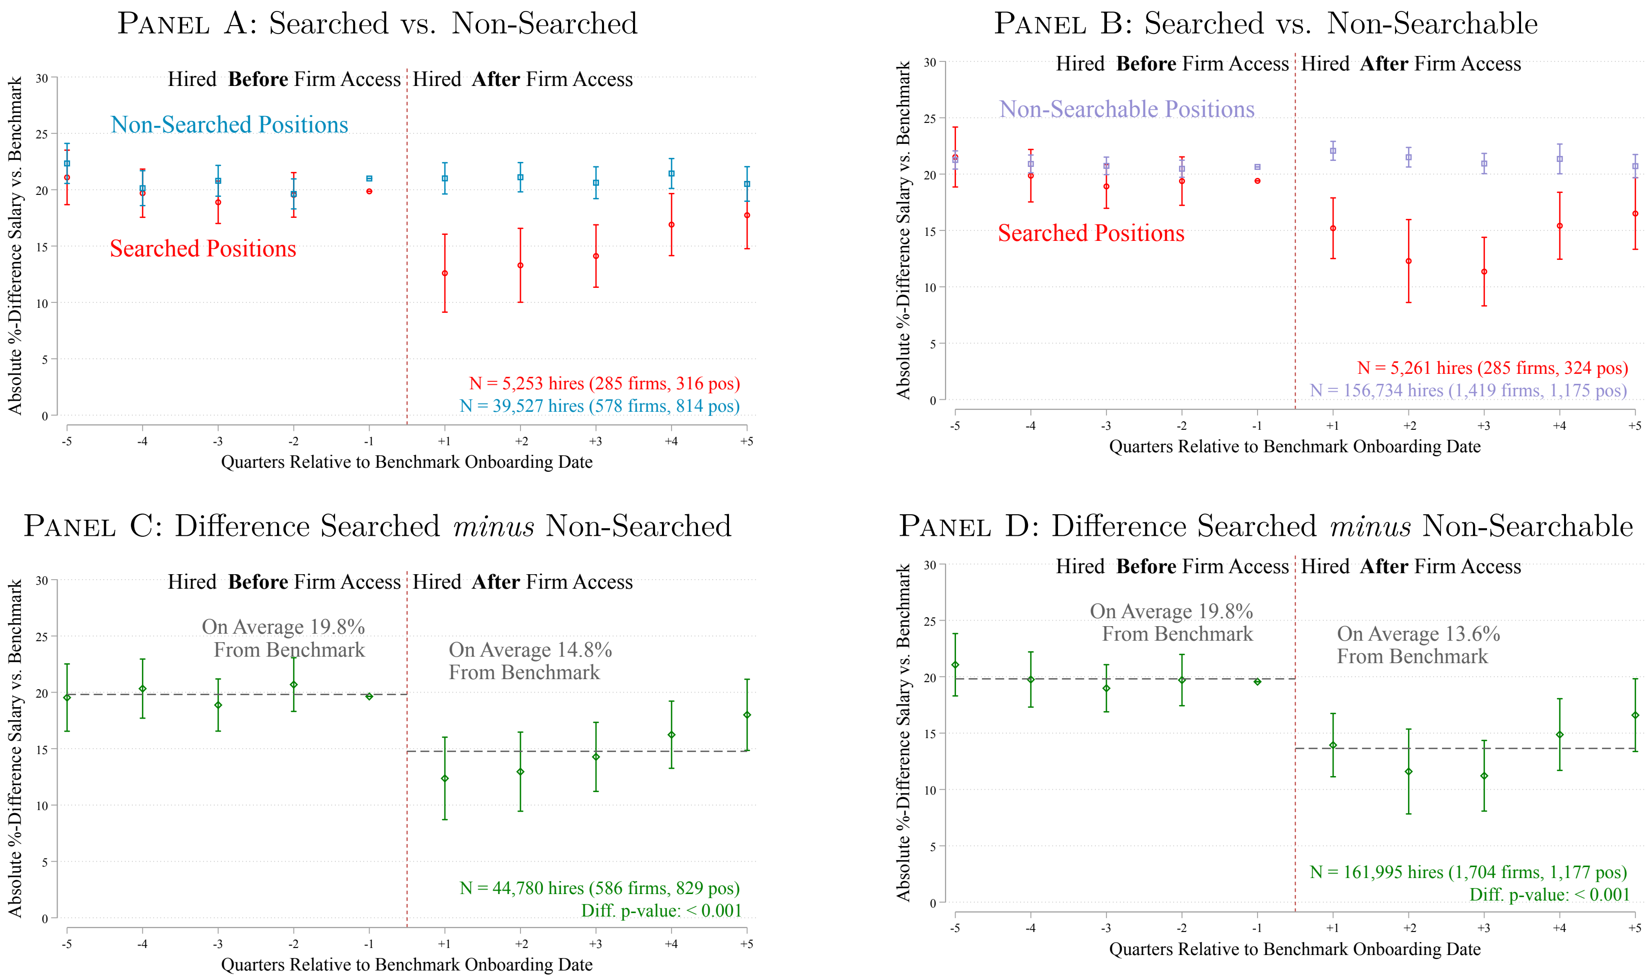

Notes: Point estimates with 90% confidence intervals in brackets, using standard errors clustered at the firm-position-month level. Panels A and C are based off one regression for Searched and Non-Searched positions, while Panel A presents the estimates for each position type, and Panel C presents the difference. Panels B and D are analogous for Searched vs. Non-Searchable positions. All coefficients are shifted such that the pre-treatment coefficients average to the pre-treatment mean of the absolute dispersion outcome. Coefficients in panels C and D refer to parameters $\alpha^k_{1,s}\ \forall s \in S$ from equation (3) (see Section 3.2 for details).

Figure 4: Heterogeneity by Skill: Non-Parametric Analysis

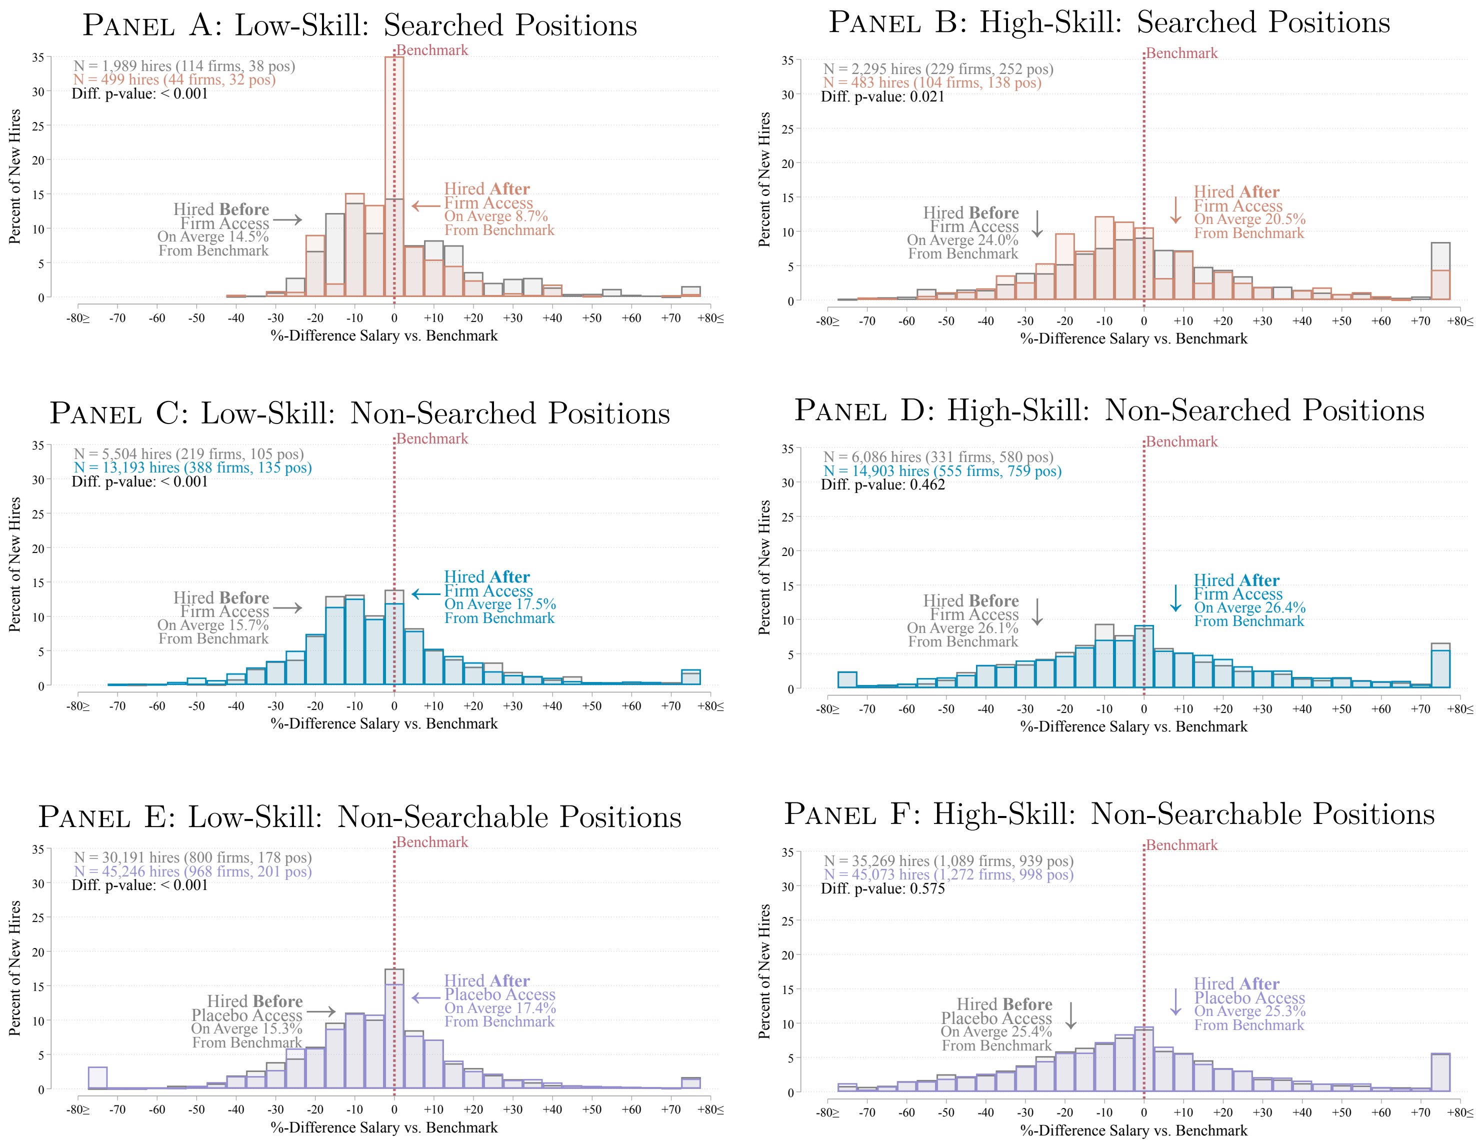

Notes: All figures are a reproduction of the corresponding panel of Figure 2 for low-skill positions (left) and high-skill positions (right).

Figure 5: Event-Study Analysis: The Effects on Salary Levels

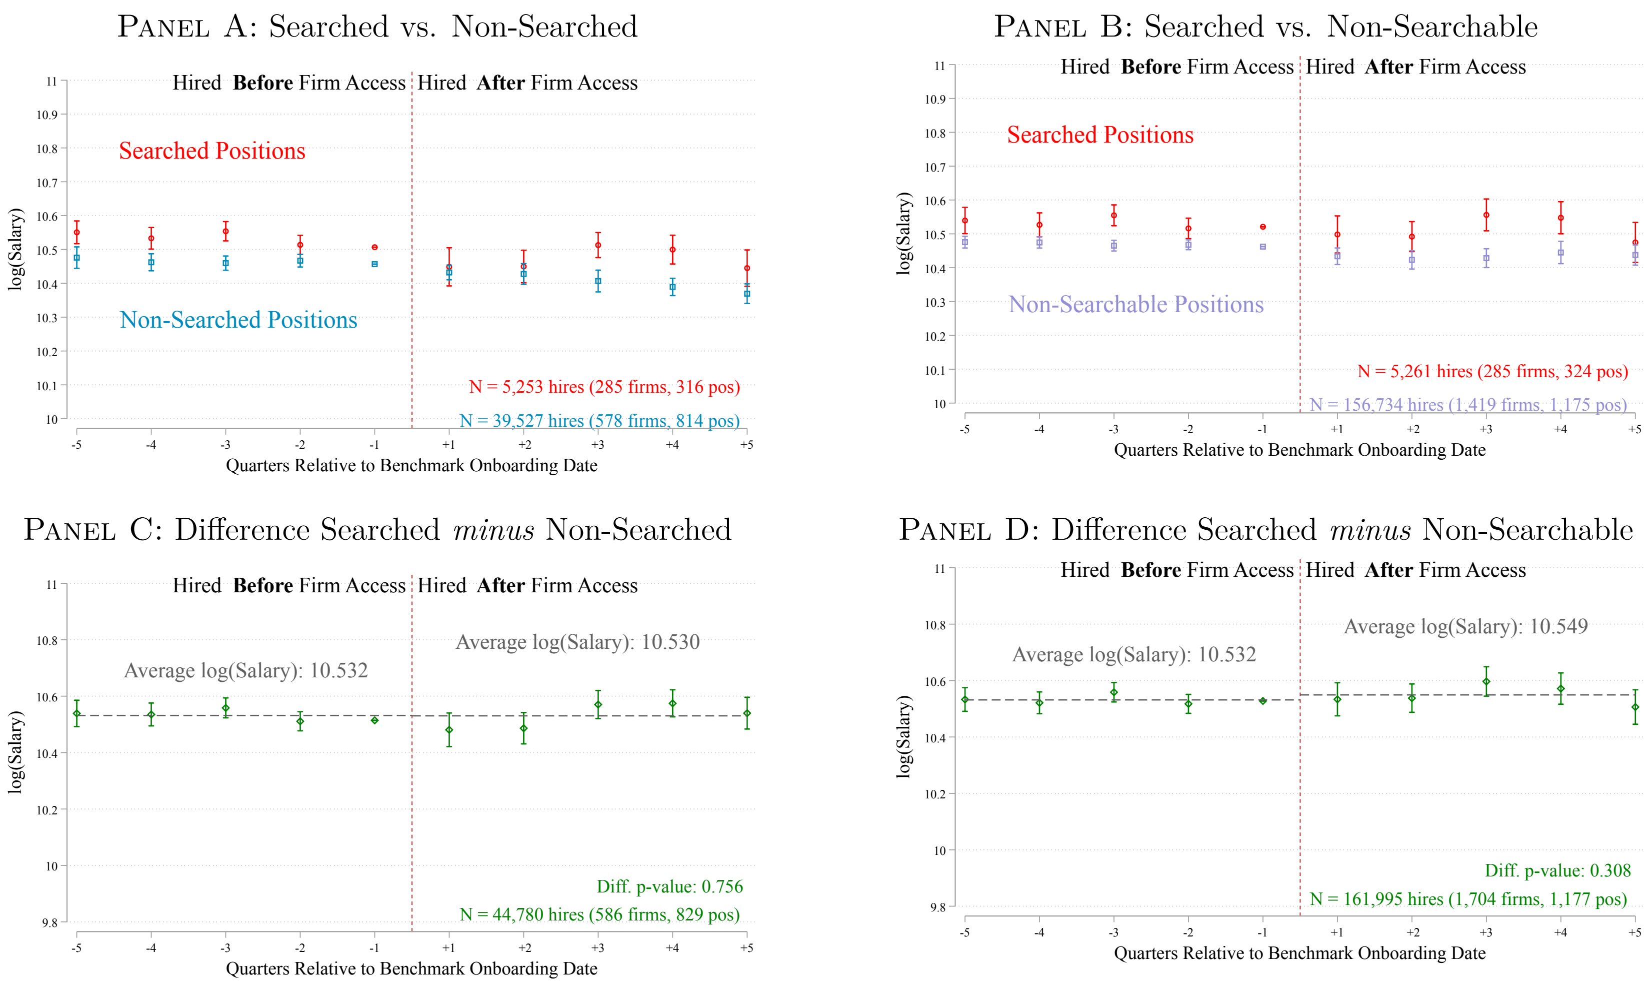

Notes: Point estimates with 90% confidence intervals in brackets, using robust standard errors. Panels A and C are based off one regression for Searched and Non-Searched positions, while Panel A presents the estimates for each position type and Panel C presents the difference. Panels B and D are analogous for Searched vs. Non-Searchable positions. All coefficients are shifted such that the pre-treatment coefficients average to the pre-treatment mean of log salary. Coefficients in panels C and D refer to the parameters $\alpha^k_{1,s}$ $\forall s \in S$ from equation (3) (see Section 3.2 for details).

Figure 6: The Effects of Salary Benchmarking on Salary Levels: Low-Skill Subsample

Panel A: Searched vs. Non-Searched

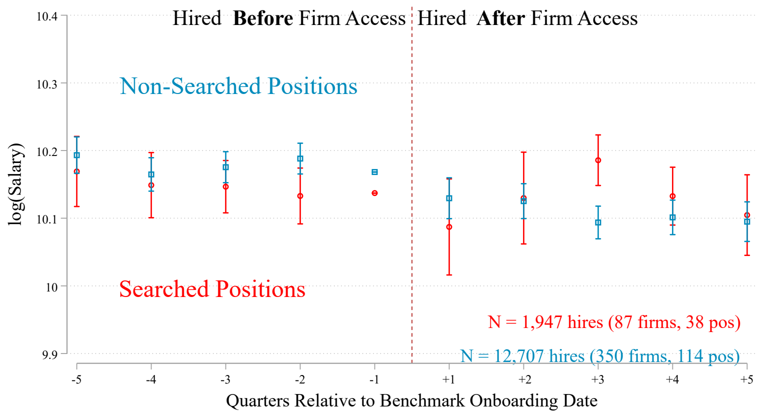

Panel B: Searched vs. Non-Searchable

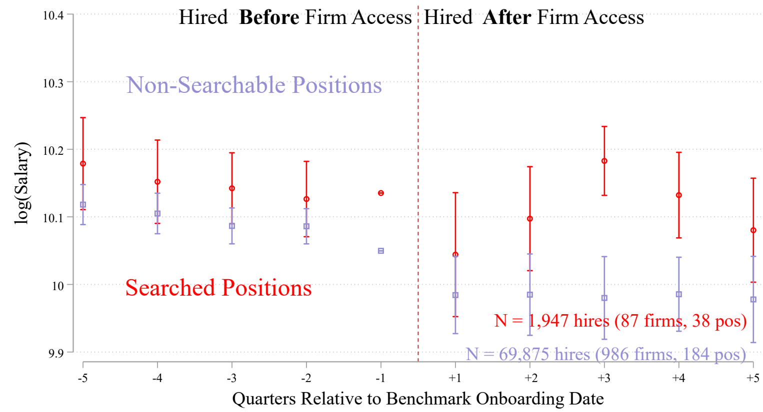

Panel C: Difference Searched *minus* Non-Searched

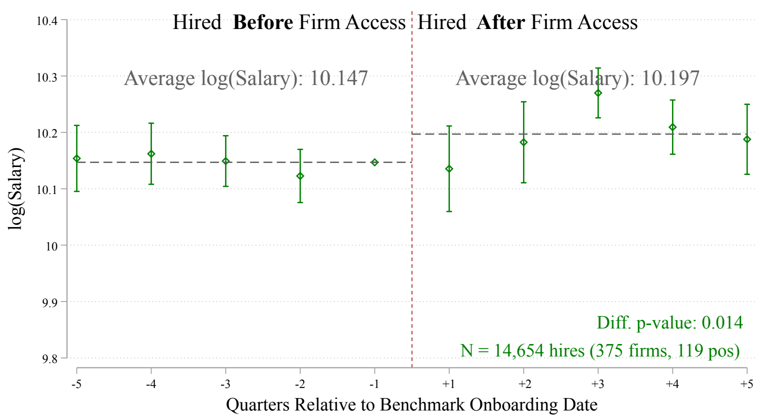

Panel D: Difference Searched *minus* Non-Searchable

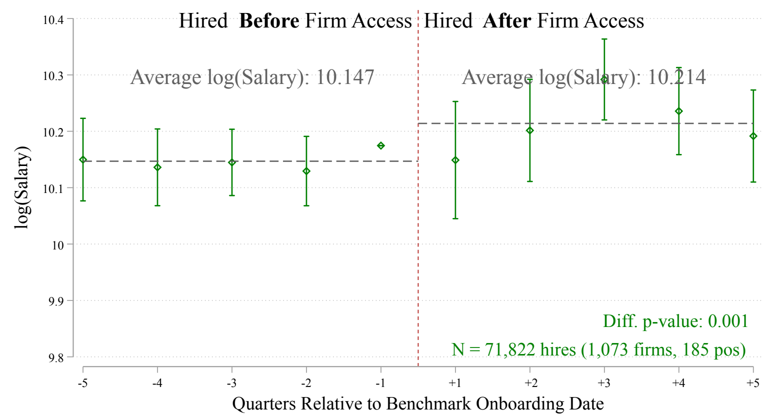

Notes:

Figure 7: The Effects of Salary Benchmarking on Retention Rates: Low-Skill Subsample

Panel A: Searched vs. Non-Searched

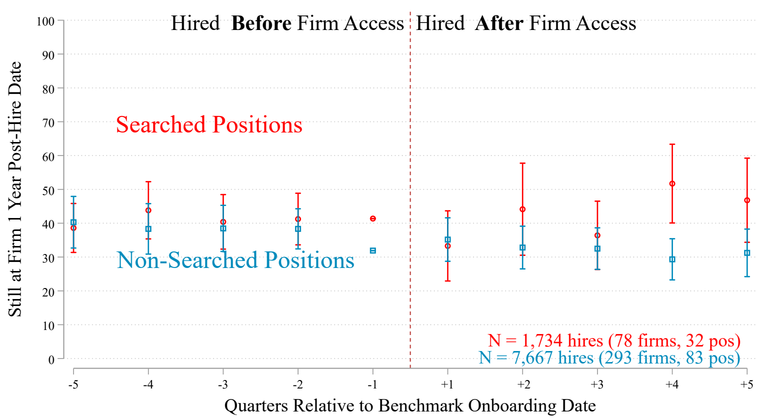

Panel B: Searched vs. Non-Searchable

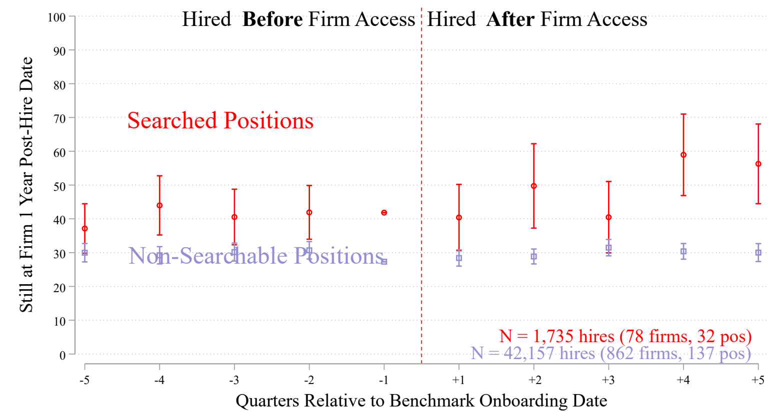

Panel C: Difference Searched *minus* Non-Searched

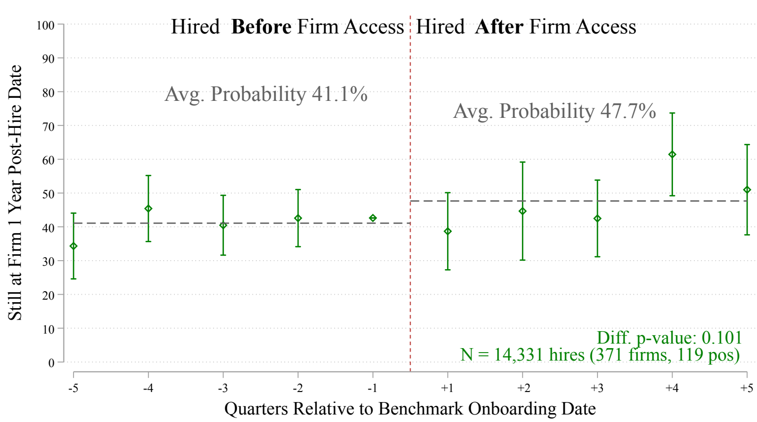

Panel D: Difference Searched *minus* Non-Searchable

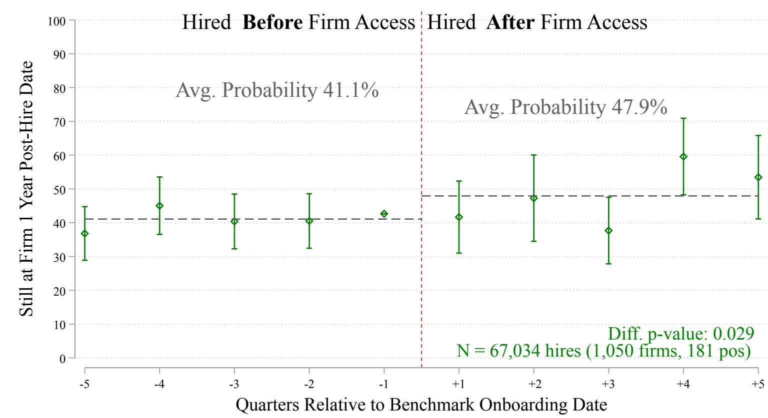

Notes:

Figure 8: Illustrative example of $w^*(v,q)$, for given $Q = q$.

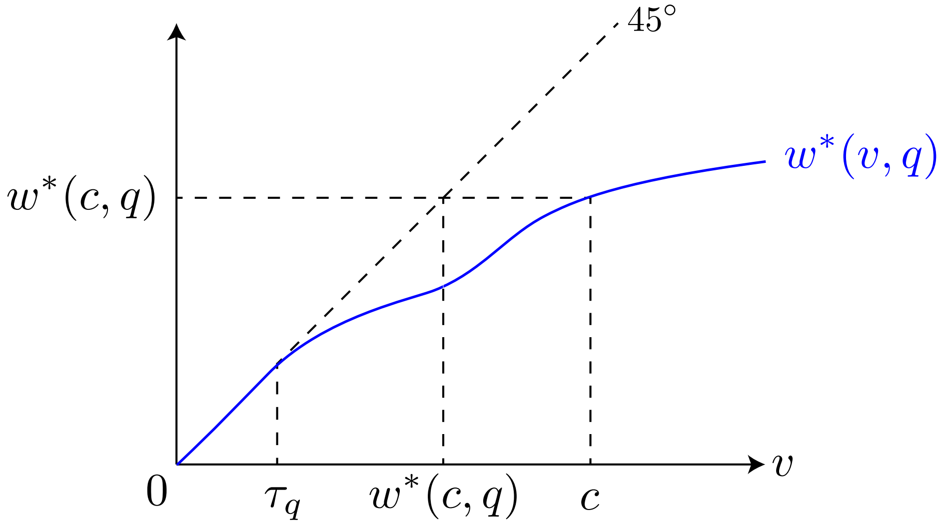

Notes: Suppose a firm with value v learns that the cutoff is c. If $v \in [w^*(c,q), c)$, then it is a best-response to raise its offer to $w^*(c,q)$. On the other hand, if $w^*(c,q) < v$, then it is a best-response to lower its offer to $w^*(c,q)$. By Theorem 5.2, $w^*(v,q)$ is increasing and continuous in v, and bounded above by v.

Table 1: Summary Statistics for Firms with vs. without Access

		Has Access?		By Usage		Early Access	
	(1) All	(2) No	(3) Yes	(4) Higher	(5) Lower	(6) Early	(7) Late
Average Firm Characteristics							
Average Employment	503.3	509.8	483.2	525.7	444.5	482.1	483.9
	(28.1)	(33.2)	(52.1)	(50.3)	(88.4)	(50.9)	(79.1)
Turnover Rate (%)†	2.424	2.438	2.382	2.392	2.374	2.680	2.192
	(0.061)	(0.070)	(0.126)	(0.159)	(0.192)	(0.279)	(0.101)
Business Services Sector (%)	17.27	16.73	18.94	14.62	22.87	18.57	19.18
	(0.99)	(1.13)	(2.07)	(2.71)	(3.07)	(3.30)	(2.67)
Hospitality Sector (%)	2.62	2.83	1.95	2.34	1.60	2.14	1.83
	(0.42)	(0.50)	(0.73)	(1.16)	(0.92)	(1.23)	(0.91)
Retail & Wholesale Trade Sector (%)	12.04	11.97	12.26	16.37	8.51	12.14	12.33
	(0.85)	(0.98)	(1.73)	(2.84)	(2.04)	(2.77)	(2.23)
Health Care Sector (%)	8.47	7.95	10.03	11.70	8.51	10.00	10.05
	(0.73)	(0.82)	(1.59)	(2.46)	(2.04)	(2.54)	(2.04)
Banking Sector (%)	7.16	7.13	7.24	7.02	7.45	6.43	7.76
	(0.68)	(0.78)	(1.37)	(1.96)	(1.92)	(2.08)	(1.81)
Other Sector (%)	52.44	53.38	49.58	47.95	51.06	50.71	48.86
	(1.31)	(1.51)	(2.64)	(3.83)	(3.66)	(4.24)	(3.39)
Average Employee Characteristics							
Salary (annual)†	46,945	46,439	48,488	45,232	51,449	48,445	48,515
	(794)	(956)	(1,356)	(1,632)	(2,103)	(2,366)	(1,634)
External Benchmark (annual)†	47,643	47,008	49,579	46,491	52,389	48,744	50,114
	(652)	(752)	(1,307)	(1,650)	(1,977)	(1,931)	(1,754)
Abs. %-Diff. Salary vs. Benchmark†	22.16	22.46	21.26	19.41	22.95	21.28	21.25
	(0.38)	(0.45)	(0.68)	(0.84)	(1.04)	(1.10)	(0.87)
Age	34.40	34.30	34.72	34.36	35.04	35.09	34.48
	(0.18)	(0.22)	(0.32)	(0.42)	(0.48)	(0.55)	(0.39)
Share Female (%)	45.29	46.39	41.92	44.74	39.36	40.00	43.15
	(1.29)	(1.48)	(2.57)	(3.78)	(3.51)	(4.09)	(3.32)
Share High Education (%)	56.92	55.30	61.84	57.89	65.43	62.86	61.19
	(1.28)	(1.49)	(2.53)	(3.74)	(3.42)	(4.00)	(3.27)
Share Hourly (%)	71.89	73.08	68.25	71.35	65.43	71.79	65.98
	(1.17)	(1.33)	(2.44)	(3.47)	(3.44)	(3.77)	(3.20)
Base Salary as Share of Total Comp. (%)	95.70	95.92	95.03	95.18	94.89	95.10	94.98
	(0.20)	(0.25)	(0.33)	(0.45)	(0.47)	(0.48)	(0.43)
Number of Firms	2,005	1,419	586	183	403	183	403

Notes: Average characteristics in the main sample of new hires, with robust standard errors in parentheses. Variables marked with † are computed using only pre-onboarding data. *Higher Usage* are firms that search at least once and *Lower Usage* are firms with access that never search. *Early* are firms that are given access before the median date. *Late* are firms that are given access after the median date. *Turnover Rate* is defined as number of employee departures in a month over the number of employees employed at the firm during that month. *Business Services Sector* through *Other Sector* correspond to the distribution of industry sectors. *Salary* is the annual base salary at the time of hire. *External Benchmark* is the median annual base salary benchmark in the position of the new hire during the quarter of the hire date.

Table 2: Summary Statistics by Position Type

		by Position Type		
	(1)	(2)	(3)	(4)
	All	Searched	Non-Searched	Non-Searchable
Salary (annual \$)†	41,359	39,064	42,013	41,405
	(146)	(462)	(390)	(166)
External Benchmark (annual \$)†	41,412	38,649	41,092	41,672
	(113)	(409)	(295)	(128)
Abs. %-Diff. Salary vs. Benchmark†	20.36	17.36	21.03	20.45
	(0.08)	(0.28)	(0.21)	(0.09)
Age	34.77	34.53	34.54	34.83
	(0.05)	(0.22)	(0.13)	(0.06)
Share Female (%)	50.63	60.14	51.01	49.87
	(0.20)	(0.83)	(0.53)	(0.23)
Share High Education (%)	42.21	34.49	42.28	42.76
	(0.20)	(0.80)	(0.52)	(0.23)
Share Hourly (%)	81.11	82.94	80.13	81.16
	(0.16)	(0.64)	(0.42)	(0.18)
Base Salary as Share of Total Comp. (%)	93.07	93.47	91.58	93.32
	(0.06)	(0.19)	(0.14)	(0.06)
Occupation Groups				
Office and Administrative Support (%)	19.84	32.44	28.97	17.23
	(0.16)	(0.79)	(0.48)	(0.17)
Building and Grounds Cleaning (%)	4.77	5.22	2.58	5.14
	(0.09)	(0.38)	(0.17)	(0.10)
Management (%)	8.04	8.10	9.21	7.81
	(0.11)	(0.46)	(0.31)	(0.12)
Production (%)	6.59	6.48	6.35	6.64
	(0.10)	(0.42)	(0.26)	(0.11)
Transportation and Material Moving (%)	9.30	6.62	9.72	9.42
	(0.12)	(0.42)	(0.31)	(0.13)
Other (%)	51.47	41.14	43.16	53.75
	(0.20)	(0.83)	(0.52)	(0.23)
Number of Firms	2,005	285	578	1,419
Number of Positions	1,406	329	973	1,306
Observations	201,817	5,266	39,686	156,865

Notes: Average characteristics in the main sample of new hires, with robust standard errors in parentheses. Variables marked with † are computed using only pre-onboarding data. *Salary* is the annual base salary at the time of hire. *External Benchmark* is the median annual base salary benchmark in the position of the new hire during the quarter of the hire date. Variables under *Occupation Groups* correspond to a new hire's SOC group.

Table 3: The Effects of Benchmarking on Salary Dispersion

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	$\lvert\%\Delta\rvert$	$\lvert log\Delta\rvert$	$\lvert\%\Delta\rvert > 10$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$
Panel A: Post-treatment												
Searched vs. Non-Searched	-4.775***	-5.155***	-16.270***	-5.148***	-4.775***	-4.786***	-5.324***	-4.950***	-4.421***	-4.887***	-4.880***	-4.564***
	(1.143)	(1.266)	(3.626)	(1.338)	(0.906)	(1.198)	(1.282)	(1.286)	(1.153)	(1.165)	(1.276)	(1.178)
Searched vs. Non-Searchable	-6.149***	-7.118***	-13.861***	-6.836***	-6.149***	-6.128***	-7.494***	-7.450***	-5.714***	-6.163***	-5.044***	-5.934***
	(1.070)	(1.211)	(3.681)	(1.220)	(0.824)	(1.076)	(1.233)	(1.576)	(1.078)	(1.087)	(1.231)	(1.127)
Panel B: Pre-treatment												
Searched vs. Non-Searched	-0.346	-0.129	-5.872	-0.233	-0.346	-0.488	-1.646	-2.062*	-0.714	-0.144	-2.205	-0.199
	(1.167)	(1.313)	(3.690)	(1.289)	(0.751)	(1.185)	(1.514)	(1.200)	(1.133)	(1.199)	(1.528)	(1.174)
Searched vs. Non-Searchable	-0.310	0.156	-4.221	-0.513	-0.310	-0.318	0.021	-1.029	0.241	-0.247	-0.754	-0.500
	(1.055)	(1.175)	(3.246)	(1.184)	(0.643)	(1.057)	(1.375)	(1.116)	(1.046)	(1.069)	(1.342)	(1.105)
Winsorizing at +/- 100%				✓								
No Clustering					✓							
No Additional Controls						✓						
No Position FE							✓					
Firm FE								✓				
Exclude High-Tip Jobs									✓			
Searched Positions Only										✓		
No Re-weighting											✓	
Ages 21-60												✓
Mean Dep. Var. (Baseline)	19.812	20.590	63.732	21.004	19.812	19.812	19.812	19.812	19.430	19.812	19.802	19.903
Observations												
Searched	5,253	5,253	5,253	5,253	5,253	5,253	5,266	5,262	5,105	5,253	5,331	4,611
Non-Searched	39,527	39,527	39,527	39,527	39,527	39,527	39,686	39,673	37,841	34,954	39,810	34,338
Non-Searchable	156,734	156,734	156,734	156,734	156,734	156,734	156,865	156,817	148,521	127,145	157,018	135,051

Notes: Significant at *10%, **5%, ***1%. Standard errors clustered at the firm-position-month level in parentheses. Each column corresponds to two regressions: one for Searched vs. Non-Searched new hires and one for Searched vs. Non-Searchable new hires. Post-treatment coefficients in Panel A refer to parameters α_1^k from equation (2), while pre-treatment coefficients in Panel B refer to parameters α_3^k from equation (2) (see Section 3.2 for details). All columns include year fixed effects. In columns (1) and (4)–(12) the dependent variable is the absolute percent difference between the annual base salary and median benchmark (Δ). The dependent variable in column (2) is the log of Δ and in column (3) is a dummy that equals 100 if $|\%\Delta|$ is greater than 10% and zero otherwise. We multiply $\%\Delta$ and $log(\Delta)$ by 100 so that the effects can be interpreted as percentage points. Δ is winsorized to $\pm$ 75 except in column (4) where it is winsorized to $\pm$ 100. All columns except (6) include additional controls (female dummy, high education dummy, hourly dummy, age, position tenure). Column (7) excludes position fixed effects. Column (8) includes firm fixed effects instead of position fixed effects. Column (9) excludes the three positions where gross pay most exceeds base pay: Waiter/Waitress, Chauffeur, and Bartender/Mixologist. Column (10) restricts the sample to only positions of Non-Searched or Non-Searchable new hires in positions that are searched and hired by firms in the data.

Online Appendix (For Online Publication Only)

What's My Employee Worth? The Effects of Salary Benchmarking

Cullen, Li and Perez-Truglia (March 24, 2024)

A Proofs of Theorems

A.1 Proof of Theorem 5.1

Lemma A.1. *In any no-benchmark equilibrium η, if a firm with value v hires with positive probability when $Q = q$, then $\eta(v, q) \leq v$.*

Proof. Otherwise the firm with value v can profitably deviate to offer 0. □

Let ω be an arbitrary constant. Suppose that $\eta(v, q) = \omega$ for all $v > \tau_q$; we now argue that η is not a no-benchmark equilibrium. There are two cases to check; $\omega > \tau_q$ and $\omega \leq \tau_q$.

If $\omega > \tau_q$, then some firm makes an offer strictly higher than its value and hires with positive probability, so η is not an equilibrium by Lemma A.1.

Suppose instead that $\omega \leq \tau_q$. Since we assumed that C has a density conditional on $Q = q$ and $V_i = v$, it follows that $\tau_q < \overline{v}$ and, moreover, that under η the firm with value $\overline{v}$ hires with probability strictly less than 1. If offering ω results in a positive probability of hiring, then the firm with value $\overline{v}$ is sometimes rationed under the tie-breaking rule, so it can profitably deviate to offer $\omega + \epsilon$ for small $\epsilon > 0$. If offering ω results in zero probability of hiring, then the firm with value $\overline{v}$ can profitably deviate to offer $\tau_q + \epsilon$ for small $\epsilon > 0$, which by Lemma A.1 hires with probability 1. This completes the proof.

A.2 Proof of Theorem 5.2

Lemma A.2. *$H(c \mid v, q)/h(c \mid v, q)$ is decreasing in v.*

Proof. This follows by affiliation, as in the proof of Lemma 1 in Milgrom and Weber (1982). □

By the usual argument, $L(\cdot \mid v, q)$ defines a probability distribution on $[\tau_q, v]$ that increases stochastically in v, that is $L(\alpha \mid v, q)$ is decreasing in v (Krishna, 2009, p. 95). Thus $w^*(v, q)$ is increasing in v and we have $w^*(v, q) \leq v$.

Lemma A.3. *$w^*(v, q)$ is continuous in v.*

Proof. By construction, $w^*(\cdot, q)$ is continuous at v for all $v \leq \tau_q$. In particular,

$$\lim_{v \downarrow \tau_q} \int_{\tau_q}^{v} \alpha dL(\alpha \mid v, q) = \tau_q \tag{A.1}$$

Suppose that $w^*(\cdot, q)$ is discontinuous at v for some $v > \tau_q$. Then there exists $v > \tau_q$ such that for all $\epsilon > 0$ the first expression of (A.2) is infinite.

$$\begin{aligned} \int_{v}^{v+\epsilon} \frac{h(\beta \mid \beta, q)}{H(\beta \mid \beta, q)} d\beta &\leq \int_{v}^{v+\epsilon} \frac{h(\beta \mid v+\epsilon, q)}{H(\beta \mid v+\epsilon, q)} d\beta \\ &= \ln H(v+\epsilon \mid v+\epsilon, q) - \ln H(v \mid v+\epsilon, q). \end{aligned} \tag{A.2}$$

The inequality is by Lemma A.2. The last expression of (A.2) is infinite only if $H(v \mid v+\epsilon, q) = 0$. But we assumed that for all q and all $v > \tau_q$, there exists $\epsilon > 0$ such that $H(v \mid v+\epsilon, q) > 0$, a contradiction. $\square$

A firm facing w^* can hire with positive probability only if it offers a wage of at least τ_q, so if its value is $v \leq \tau_q$, then it is optimal to offer a wage equal to v.

It remains to verify that no firm with value $v > \tau_q$ can profitably deviate. Suppose for the moment that $w^*(v, q)$ is differentiable in v. For $v > \tau_q$, the wage function w^* defined in (5) solves the differential equation

$$w_1^*(v, q) = (v - w^*(v, q)) \frac{h(v \mid v, q)}{H(v \mid v, q)} \tag{A.3}$$

with boundary condition $w^*(\tau_q, q) = q$. Thus we have

$$(\hat{v} - w^*(\hat{v}, q)) \frac{h(\hat{v} \mid \hat{v}, q)}{H(\hat{v} \mid \hat{v}, q)} - w_1^*(\hat{v}, q) = 0 \text{ for all } \hat{v} > \tau_q \tag{A.4}$$

If the firm with $v > \tau_q$ has any profitable deviation, then it has a profitable deviation to an offer in the set $\{w^*(\hat{v}, q) : \hat{v} \geq \tau_q\}$. Thus it remains to verify that

$$v \in \underset{\hat{v} \geq \tau_q}{\operatorname{argmax}} (v - w^*(\hat{v}, q)) H(\hat{v} \mid v, q). \tag{A.5}$$

Taking the derivative of the objective with respect to $\hat{v}$ yields

$$\begin{aligned} &-w_1^*(\hat{v}, q) H(\hat{v} \mid v, q) + (v - w^*(\hat{v}, q)) h(\hat{v} \mid v, q) \\ &= H(\hat{v} \mid v, q) \left[(v - w^*(\hat{v}, q)) \frac{h(\hat{v} \mid v, q)}{H(\hat{v} \mid v, q)} - w_1^*(\hat{v}, q) \right] \end{aligned} \tag{A.6}$$

The right-hand side of (A.6) has the same sign as $\hat{v} - v$, by (A.4) and Lemma A.2. Thus, the firm's objective is maximized at $\hat{v} = v$.

Suppose that $w^*(v, q)$ is not differentiable in v. Since $w^*(v, q)$ is continuous in v and of bounded variation in v, and $[\underline{v}, \overline{v}]$ is bounded, there exists a homeomorphism $\lambda : [\underline{v}, \overline{v}] \to [\underline{v}, \overline{v}]$ such that $\overline{w}(\theta, q) \equiv w^*(\lambda(\theta), q)$ is differentiable in θ (Bruckner and Goffman, 1976). Moreover, we can pick λ to be strictly increasing. Let us define the random variables $\Theta_i \equiv \lambda^{-1}(V_i)$ and

$\overline{C} \equiv \lambda^{-1}(C)$. Affiliation is preserved under increasing transformations (Milgrom and Weber, 1982, Theorem 5), so $(\Theta_i, \overline{C})$ are affiliated conditional on Q. Let $\overline{H}(\bar{c} \mid \theta, q) \equiv P(\overline{C} \leq \bar{c} \mid \Theta_i = \theta, Q = q)$. Now we reformulate (A.5) into the equivalent statement

$$\theta \in \operatorname*{argmax}_{\hat{\theta} \geq \lambda^{-1}(\tau_q)} (\lambda(\theta) - \overline{w}(\hat{\theta}, q))\overline{H}(\hat{\theta} \mid \theta, q). \tag{A.7}$$

The argument we used to prove (A.5) for differentiable w^* applies *mutatis mutandis* to establish (A.7). It follows that w^* is a no-benchmark equilibrium.

A.3 Proof of Theorem 5.5

This closely parallels the proof of Theorem 15 of Milgrom and Weber (1982). Let us consider an incentive-compatible mechanism M in which the firm observes V_i and Q, and then reports V_i to the mechanism. Depending on the firm's report and the cutoff C, this results in some probability of hiring and some wage conditional on hiring. Let A be the mechanism that corresponds to the no-benchmark equilibrium w^*, and let B be the mechanism that corresponds to the benchmark equilibrium $\tilde{w}$.

Let $R(\hat{v}, v, q) = vH(\hat{v} \mid v, q)$ denote the expected gross revenue to the firm when it reports value $\hat{v}$, its value is v, and the supply of workers is q.

Let $K^M(\hat{v}, v, q)$ denote the conditional expected wage paid by the firm in mechanism M when it reports value $\hat{v}$, its value is v, the supply of workers is q, and the firm hires a worker. Observe that $K^A(\hat{v}, v, q) = w^*(\hat{v}, q)$ and $K^B(\hat{v}, v, q) = E[C \mid C \leq \hat{v}, V_i = v, Q = q]$.

Facing mechanism $M \in \{A, B\}$, the firm chooses $\hat{v}$ to maximize

$$R(\hat{v}, v, q) - K^M(\hat{v}, z, q)H(\hat{v} \mid v, q). \tag{A.8}$$

By incentive compatibility, the first-order condition holds at $\hat{v} = v$, so for $v > \tau_q$ we have

$$0 = R_1(v, v, q) - K_1^M(v, v, q)H(v \mid v, q) - K^M(v, v, q)h(v \mid v, q). \tag{A.9}$$

The boundary condition is $K^M(\tau_q, \tau_q, q) = \tau_q$.

By inspection, the partial derivative $K_2^A(\hat{v}, v, q)$ is equal to zero. By affiliation and Theorem 5 of Milgrom and Weber (1982), we have $K_2^B(\hat{v}, v, q) \geq 0$. If $K^B(v, v, q) < K^A(v, v, q)$ for some v and q, then by (A.9) we have

$$\frac{d}{dv}K^B(v, v, q) = K_1^B + K_2^B \geq K_1^A + K_2^A = \frac{d}{dv}K^A(v, v, q). \tag{A.10}$$

By Lemma 2 of Milgrom and Weber (1982), it follows that $K^B(v, v, q) \geq K^A(v, v, q)$ for all $v > \tau_q$. Thus, for all $v > \tau_q$ we have

$$w^*(v, q) = K^A(v, v, q) \leq K^B(v, v, q) = E_C\left[\tilde{w}(v, q, C) \mid C \leq v, V_i = v, Q = q\right]. \tag{A.11}$$

The corollary follows by the law of iterated expectations.

Supplemental Material

What's My Employee Worth? The Effects of Salary Benchmarking

Cullen, Li and Perez-Truglia (March 24, 2024)

B SHRM Survey

B.1 Pool of Respondents

Table B.1 presents the average characteristics of the sample of respondents. Most of the participants work in the private sector (80.81%), set salaries for both new hires and current employees (91.04%), and have 6 years or more of experience setting salaries (63.26%). The sample includes firms of all sizes and industries.

B.2 Use of Salary Benchmarking Tools

Table B.2 provides statistics on the use of salary benchmark tools. The most common uses of benchmark tools are to set the salary ranges for specific job titles and to change the salaries of current employees – 89.78% and 76.81% of participants, respectively, indicate that they use them for these purposes. Other popular uses are to set precise salaries for new employees (54.07%), in salary negotiations (53.11%), to determine salary in job advertisements (40.89%), and to plan ahead for headcount (25.33%). When asked about when they typically use benchmarks in relation to new hires, most (67.19%) answer that they use them before they publicize the position to include the expected salary in a job advertisement. Some managers also indicate that they use them right before making an offer to the candidate (34.96%), during negotiations after the candidate received the initial offer (22.30%), and when the candidate presents an outside offer (12.44%). When asked the same question but in relation to existing employees, the majority (74.30%) indicate that they use them when adjusting the salary ranges for positions. Many also respond that they use them when the employee goes through an annual review (48.06%), when the employee is up for promotion (47.48%) and when the employee presents an outside offer (33.11%).

Table B.2 shows that the majority (61.67%) of HR managers in our sample report using benchmarks to set salaries for the majority or all of their current employees, 25.89% only for some, 10.01% for a minority and 1.82% for none of their current employees. Similarly, the majority (64.36%) indicate that they use benchmarks to set salaries for a majority or all their new hires, 26.06% for some of their new hires, 6.99% for a minority of hires, and 2.58% for none of their new hires. When asked about why they use benchmarks for some but not all new hires, 68.91% answer they search for the benchmarks only for some positions and then apply them to all employees, 17.85% that they only search for specific employees within a given position, and 18.10% provide an open-ended response. Among open-ended responses, one common reason is that they search only for new positions in the firm. For instance, one respondent wrote: "For new hires, we primarily use the salaries of current employees rather than salary benchmarks. If it's a position we don't have that for, then we use the salary benchmarks.". Another reason is that they only search for positions that are more competitive or challenging to fill. For example, one respondent wrote: "Certain positions that are becoming more competitive in our industry or market." Some responses also mention a range of other reasons, such as the role of unions or budget constraints.

In relation to the sources used to obtain salary benchmarks, Table B.2 shows that the most popular options are industry surveys and free online data sources (68% and 58.07% of participants, respectively, indicate that they use these sources). Other popular options are government data (37.11%), paid online data sources (34.37%), compensation consultants (26.30%), and payroll data services (23.19%). Among HR managers in our sample, 48.59% have used Glassdoor as their salary benchmark source and 9.48% our partner firm's benchmarking tool. Table B.2 also shows that when looking up benchmark information, a strong majority apply filters for state and industry (84.15% and 87.33%, respectively). Other popular filters are firm size (48%), revenue size (38.96%), and hourly vs. salaried (37.11%). Figure B.1 presents the piece of information they usually care most about when looking for the benchmark for a position. Most (56. 73%) ranked the median salary first. The second most popular piece of information is the average salary (32.59% ranked it 1st), and only the remaining 10.68% choose the 10th percentile, 25th percentile, 75th percentile, or the 90th percentile as the one they care about the most.

B.3 Other Information Frictions

For reference, we included a survey question to assess whether there may be frictions in accessing internal information. More precisely, we asked respondents if they could access information on the median salary that their own company pays to employees in a specific position. Most (79.78%) of the participants indicated that they could access the information easily, 19.26% indicated that it would require some work, and less than 1% indicated that they do not have access. In addition, we asked HR managers how frequently their employees share with them the offers they receive from other firms. Table B.1 shows that 1.7% of the participants responded *Always,* 17.33% *Often*, 61.85% *Sometimes*, 15.56% *Rarely* and 3.56% answered *Never.* This evidence suggests that employers have limited information on competing offers.

B.4 Survey Experiment on the Effects on Salary Dispersion

We embedded an experiment in the SHRM survey. With these data, we can provide complementary experimental evidence on the effects on salary dispersion documented in Section 3. For each participants, we face them with two sequential scenarios. In each scenario, respondents pick a positions for which they plan to hire in the future. Second, we ask them for the annual base salary they are willing to pay for a new hire in that position. Third, we provide them with information on a hypothetical benchmark for that position. After they receive the feedback, we re-elicit the salary they are willing to pay for that position. Participants receive one of two types of feedback: half of the scenarios receive a benchmark that is 15% *below* what they were originally planning to offer, while the other half receive a benchmark that is 15% *above*.

Figure B.2 presents the distribution of the percentage change in the salary offer after receiving feedback. Panel A shows the salary update where the respondent was shown a benchmark 15% above the initial offer, and Panel B corresponds to scenarios where the respondent was shown a benchmark 15% below the initial offer. Figure B.2 shows that, consistent with the results from the payroll data, salaries get compressed toward the benchmark, both from above (Panel A) and below (Panel B). More precisely, when individuals learn that the benchmark is above their planned salary, they react by revising their offer upward; when they find that the benchmark is below their planned salary, they react by revising their offer downward.

An interesting similarity exists between the results of the survey experiment and the results from the natural experiment presented in Section 3. In the natural experiment, it looks like some individuals react to the benchmark information by fully updating (i.e., by bunching at exactly the median benchmark), while other individuals react by updating partially. We observe a similar pattern in the survey experiment. For example, panel A of Figure B.2 indicates that 27.9% of subjects fully updated to the feedback (i.e., they revised their offers upward by exactly 15%) while 44.2% updated only partially (i.e., they revised their offers upward by a number between 0% and 15%). On the other hand, there is one difference between the results of the survey experiment and the results from the natural experiment. In the natural experiment, the compression from above is similarly strong as the compression from above. In the survey experiment, the compression from above (Panel B) is stronger than the compression from below (Panel A). However, there are two natural explanations for this discrepancy. The simplest explanation involves downward wage rigidities: respondents might implicitly assume that the initial salary they selected was already conveyed to the hypothetical candidate, leading to a reluctance to lower the salary amount after receiving the feedback. Alternatively, social desirability bias provides another explanation. Due to the survey experiment's non-incentivized nature, respondents may wish to appear generous in the eyes of the researchers by being open to increasing salaries but reluctant to reducing them.

B.5 Affiliated Values

The model presented in Section 5 relies on a key assumption of affiliated values. We conducted another survey experiment to provide a test of this assumption. We begin by asking participants to choose the position for which the benchmark is most useful to them and to think of their two closest competitors (firm A and firm B) that also hire in that position. Second, we elicit the maximum salary they think firm A would be willing to pay a new hire in that position. Third, we give them information on a hypothetical salary that firm A is actually willing to pay. The participants are assigned to receive information that the salary is 15% below or 15% above their initial guess. After receiving the information on the salary firm A is willing to offer, we elicit the salary they think firm B would be willing to pay to hire in that position (posterior salary).

The results of this survey experiment are reported in Figure B.3, which shows the distribution of the %-change in their guess after receiving feedback. Panel A presents this outcome when the feedback is 15% above the initial guess, and Panel B when the feedback is 15% below. Consistent with the assumption on affiliated values, Figure B.3 shows compression towards the competitor's salary, both from above (Panel A) and from below (Panel B). Intuitively, when individuals learn that firm A values the worker more (less) than they thought initially, they update upward (downward) what they think firm B would be willing to pay.

Figure B.1: Ranking of Salary Benchmark Statistics

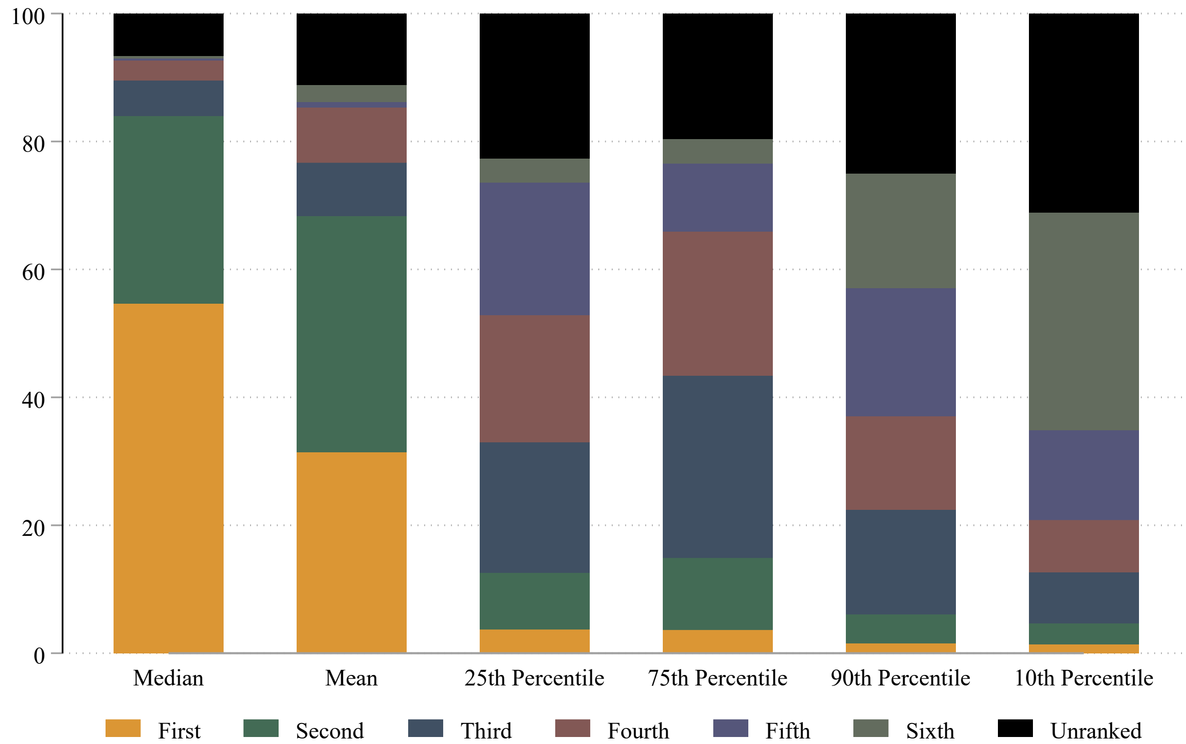

Notes: Responses from a sample of 1,301 HR managers that participated in the SHRM survey. The survey question asks participants to rank these pieces of information from most important to least important.

Figure B.2: Salary Updating with Benchmark Information

Panel A: When benchmark is 15% above prior

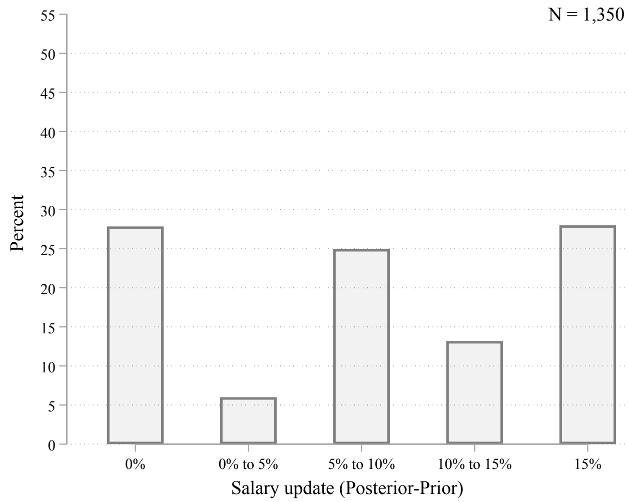

Panel B: When benchmark is 15% below prior

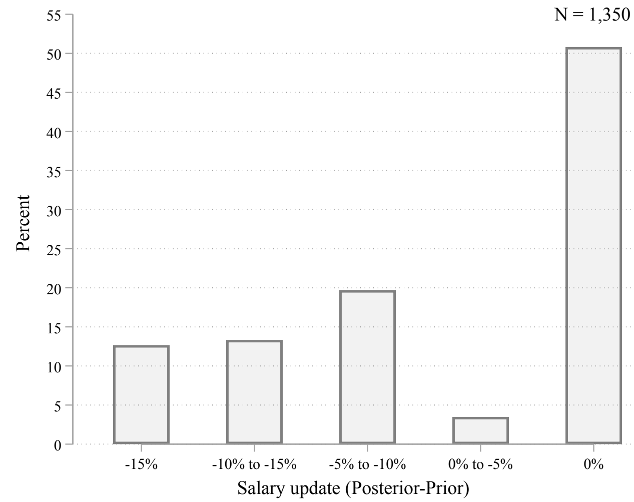

Notes: Salary updating with benchmark information in experimental setting. Participants are asked to choose two positions they expect to hire in the future and to indicate the base salary for the new hires, before and after receiving information on an hypothetical benchmark for each position. The hypothetical benchmark is 15% above their prior salary for one position, and 15% below for the other. Participants are randomly assigned to receive first a higher or lower benchmark. Panel A shows the base salary update for positions where the salary benchmark was 15% above the prior salary, and Panel for positions where the benchmark was 15% below.

Figure B.3: A Test of Affiliated Values

Panel A: When competitor's salary is 15% above prior   Panel B: When competitor's salary is 15% below prior

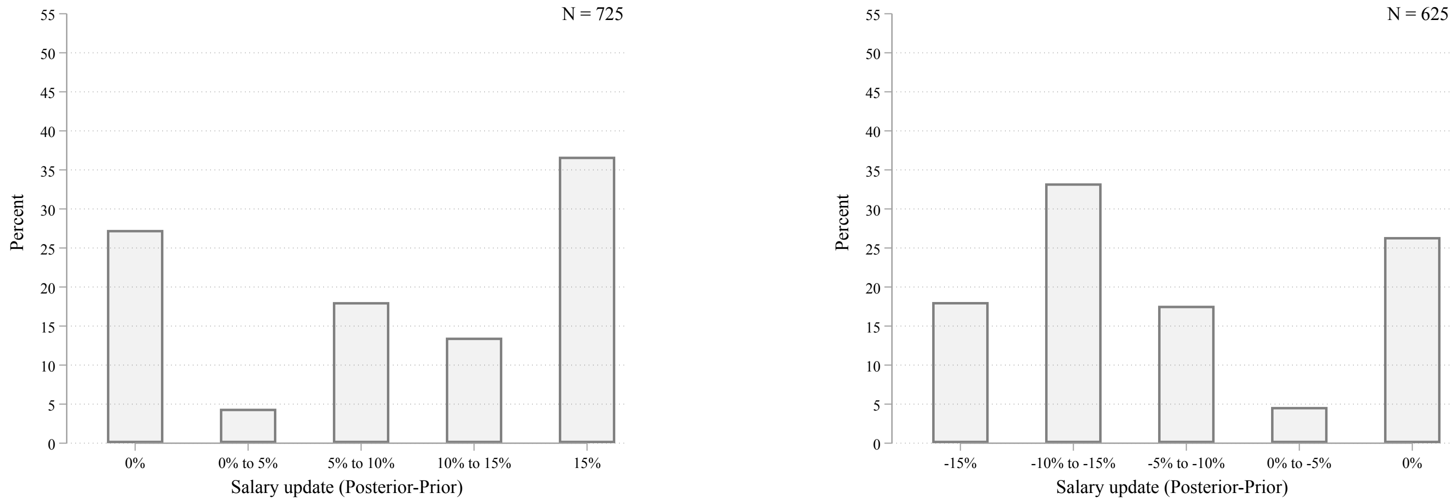

Notes: A test of affiliated values in an experimental setting. Participants are asked to choose the most relevant position for them and to think of their two closest competitors (A and B) that hire in that position. We elicit their beliefs on the maximum amount competitor A will be willing to pay to hire a worker in the position (prior salary), and then we provide them information on the actual (hypothetical) salary that competitor A would be willing to pay. We then elicit the maximum amount they think the competitor B will be willing to pay for to hire a worker in that position (posterior salary). Participants were randomly assigned to receive as information a competitor's salary 15% above their prior or 15% below. Panel A shows the salary update when the competitor's salary was 15% above the prior salary, and Panel B when it was 15% below.

Table B.1: Characteristics of the Sample

Questions	Share of responses (%)
How many employees does your company have?	
1-49	22.15
50-99	23.48
100-999	36.67
1000-4999	9.70
5000 or more	8.00
Are you working in the private sector or the public sector?	
Private sector	80.81
Public sector	19.19
What main industry do you operate in?	
Agriculture, Forestry, Fishing and Hunting	0.45
Mining, Quarrying, and Oil and Gas Extraction	0.45
Utilities	0.96
Construction	4.08
Manufacturing	15.43
Wholesale Trade	3.04
Retail Trade	4.01
Transportation and Warehousing	1.71
Information	3.49
Finance and Insurance	7.12
Real Estate and Rental and Leasing	2.15
Professional, Scientific, and Technical Services	18.99
Management of Companies and Enterprises	0.07
Administrative and Support and Waste Management and Remediation Services	3.41
Educational Services	5.04
Health Care and Social Assistance	11.80
Arts, Entertainment, and Recreation	2.23
Accommodation and Food Services	1.56
Other Services (except Public Administration)	7.86
Public Administration	6.16
How many years of experience do you have setting salaries?	
Less than 1 year	2.44
1-5 years	34.30
6-10 years	25.70
11+ years	37.56
Do you participate in salary settings for:	
New hires	6.44
Current employees	2.52
Both	91.04
Do you typically set salaries for:	
Higher-education positions	13.93
Lower-education positions	6.07
Both	80.00
Do you have access to the median salary that your company pays employees in a specific position?	
No, I could not access that data even if I wanted to	0.96
Yes, I can access it easily	79.78
Yes, but it would take some work	19.26
If your employees get an offer from another company, do they share the terms of the offer with you?	
Never	3.56
Rarely	15.56
Sometimes	61.85
Often	17.33
Always	1.70

Notes: Characteristics of the sample of 1,350 HR managers that participated in the SHRM survey.

Table B.2: Use of Salary Benchmarking Tools

Questions	Share of responses (%)
What do you use the salary benchmark for? (Select all that apply) [N = 1,350]	
To set precise salaries for new hires	54.07
To change salaries for current employees	76.81
In salary negotiations	53.11
To set salary ranges for specific job titles	89.78
To determine salary in job advertisement	40.89
To plan ahead for headcount	25.33
How frequently do you use salary benchmarks to set salaries for new hires? [N = 1,316]	
Never	2.58
A minority of hires	6.99
Some of the hires	26.06
A majority of hires	27.81
For every hire	36.55
When do you use salary benchmarks in relation to new hires? (Select all that apply) [N = 1,316]	
Before I publicize the position to include the expected salary in a job advertisement	67.19
Right before I make an offer to the candidate	34.96
After the candidate receives the offer, if the candidate wants to negotiate	22.30
When the candidate presents an outside offer	12.44
How frequently do you use salary benchmarking to change salaries for current employees? [N = 1,263]	
Never	1.82
For a minority of employees	10.61
For some of my employees	25.89
For a majority of my employees	21.77
For all my employees	39.90
When do you use salary benchmarks with current employees? (Select all that apply) [N = 1,240]	
When the employee goes through an annual review	48.96
When the employee is up for promotion	47.48
When the employee presents an outside offer	33.11
When adjusting the salary ranges for positions	74.30
Why do you use salary benchmarks for some but not all employees? (Select all that apply) [N = 801]	
I search for some specific employees (within a given position)	17.85
I search only in some specific positions (and apply it to employees in those positions)	68.91
Other	18.10
Which sources do you use to obtain salary benchmarks? (Select all that apply) [N = 1,350]	
Free online data sources	58.07
Paid online data sources	34.37
Industry surveys	68.00
Government data	37.11
Compensation consultants	26.30
Payroll data services	23.19
Have you ever used Glassdoor as your salary benchmark source? [N = 1,350]	
Yes	48.59
No	51.41
Have you ever used [the Firms's] Data Cloud Compensation Explorer as your salary benchmark source? [N = 1,350]	
Yes	9.48
No	90.52
Which filters would you typically apply when searching for a position salary benchmark? [N = 1,350]	
Industry	87.33
State	84.15
Firm size	48.00
Revenue size	38.96
Hourly vs. salaried	37.11
None of the above, the position-level filter is sufficient	0.00
Which piece of information you typically care about the most? (% that ranked first) [N = 1,301]	
Median salary	56.73
Average salary	32.59
10th percentile	1.46
25th percentile	3.84
75th percentile	3.77
90th percentile	1.61
For which positions are salary benchmarks most useful, higher-education or lower-education? [N = 1,350]	
Not useful for either group	0.30
Most useful for higher-education positions	21.78
Most useful for lower-education positions	3.93
Equally useful for both groups	74.00

Notes: Responses to survey questions on the use of salary benchmarking tools. The subject pool consists of 1,350 HR managers that participated in the SHRM survey. The relevant sample size, which may be different from question to question, is reported in brackets. For questions where more than one option could be selected (i.e., "select all that apply"), we report the share of total responses that selected that option, and thus the percentages can add up to more than 100%.

C Additional Details about the Payroll Data

C.1 Comparison with a Representative Sample of U.S. Firms

Table C.1 provides a comparison between the firms in our sample and a representative sample of U.S. firms (Song et al., 2019). We match the sample restrictions of Song et al. (2019) by excluding firms with fewer than 20 employees, and employees outside of 20 to 60 years of age. In terms of size, measured in the number of employees, our sample is most representative of the top quartile of firms in the United States. This is probably because businesses with fewer than 100 employees do not have enough scale to justify the use of data analytics services. In terms of salaries, the employees in our sample are representative of the population of U.S. employees, with the exception that our sample has limited coverage of the bottom quartile of the distribution (earning less than $20,000 per year).

Table C.2 provides some statistics on the distribution of industries, given by the first 2 digits of the firm's main 6-digit NAICS code. Columns (1) and (2) compare the distribution of sectors in our sample (column (1)) with the U.S. distribution according to Census data (column (2)). Columns (3) and (4) are the same as columns (1) and (2), except that they are based on the number of employees instead of the number of firms. We should not expect our sample to be perfectly representative of the U.S. industries. For example, as discussed above, the firms in our sample are larger than the U.S. average, and as a result they will be more representative of industries with larger firms. While not perfectly representative of the U.S. average, our sample provides broad coverage of the U.S. industries. Some industries, such as manufacturing and finance, are somewhat overrepresented, while some other industries, such as construction and accommodation and food services, are somewhat underrepresented.

C.2 Timing of Adoption and Utilization

Figure C.1 shows the timing of onboarding of the firms in our dataset. The onboarding dates range from December 2015 to January 2020. Table C.3 presents a list of the most popular position titles in the Searched position. Mechanically, this list includes a lot of positions at the lower end of the company's hierarchies, as those are the positions in which the companies hire most often. The important lesson from Table C.3 is that there is a lot of overlap between the different position types. For instance, while bank tellers fall into the Searched category for some firms, there are plenty of bank tellers in the Non-Searched and Non-Searchable categories as well. The same is true for each of the other position titles listed in this table.

C.3 Classification of Positions by Skill Level

In Table C.4, we provide more details on the classification of the positions as high-skill and low-skill. This table reports some characteristics of the 35 most searched position titles. Column (1) indicates if a position is classified as high-skill, as defined in Section 2.4. Column (2) indicates whether a position is high-education or low-education. Column (3) indicates if the mean age among new hires in a position is equal to or above 31. Column (4) indicates if the mean salary among new hires in a position is equal to or above $30,000. A position is categorized as high-skill if: it is categorized as high-education (column (2)); or if it is low-age (column (3)) and low-earnings (column (4)).[65]

C.4 Comparison between the Proprietary Benchmark and a Free Online Benchmark

In Figure C.2, we compare our salary benchmark, constructed from proprietary payroll data, with one popular free online salary benchmark. The alternative benchmark is constructed from anonymous users who voluntarily provide their salary information. According to the SHRM survey, 48.59% of HR managers report having used this free online source for their salary benchmarking. To compare these two benchmarks, we use the salient figure provided by the free online tool with the similarly salient median salary reported in our proprietary salary benchmark.[66] We select the 100 most searched positions in our proprietary benchmark tool, and collect the salary benchmark available for these positions in the free online website in 2019 using The Wayback Machine. For each position, we compare the proprietary and the free online benchmark for the same quarter of 2019. Figure C.2 presents a histogram of the percent difference between the proprietary salary benchmark and the free online benchmark. While the propietary benchmark and the online source are similar on average, there is a significant deal of variation. Most of the positions in the public benchmark are more than 15% off from the propietary benchmark.

[65] According to the benchmark data, among low-skill positions, 95.2% of the total cash compensation comes in the form of base salary. For high-skill positions, the corresponding figure is 92.9%.

[66] On the free online tool, the salient figure is referred to as the "average salary". However, in the technical notes they report that it is calculated as the median salary.

Figure C.1: Onboarding Dates for the Firms in Our Sample

Panel A: Treated Firms

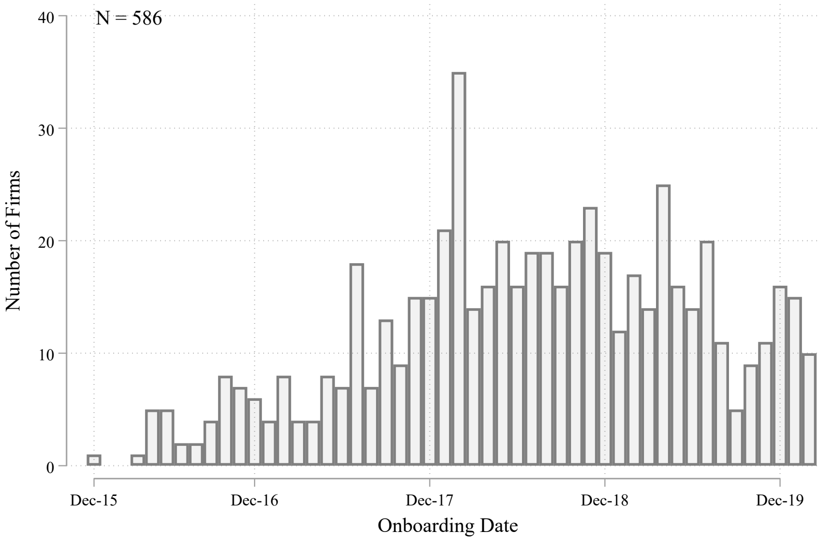

Panel B: Control Firms

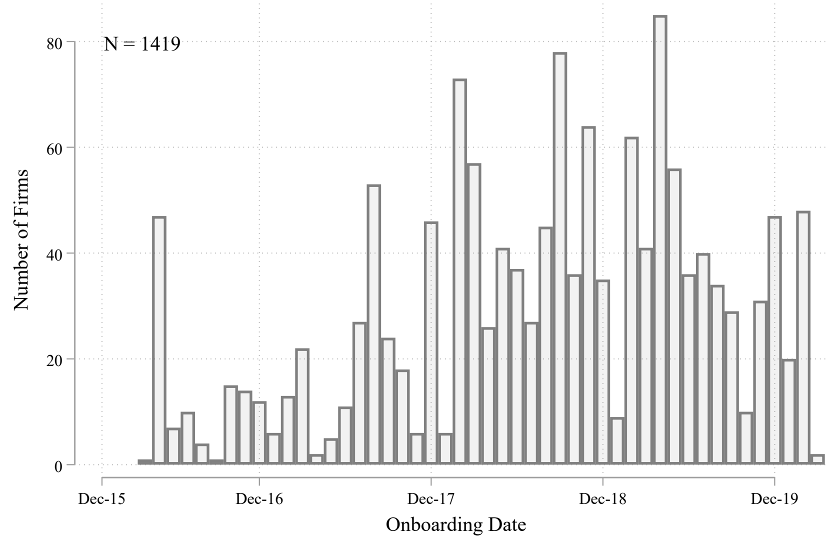

Notes: Panel A presents the onboarding dates of the 586 firms in our main data sample. Panel B presents the "hypothetical" onboarding date assigned to the firm that never gains access to the tool. See section 2.6 for more details.

Figure C.2: Comparison Between Proprietary vs. Public Online Benchmarks

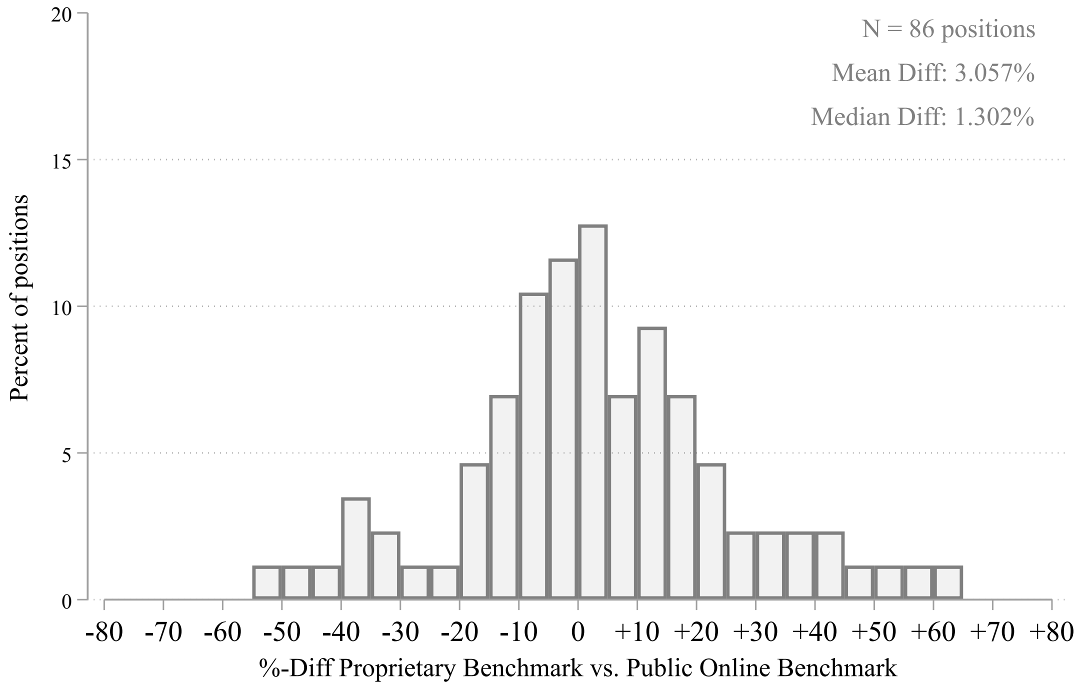

Notes: Histogram of the percent difference between the proprietary median annual base salary benchmark relative to a public online benchmark of median salary. The graph includes salary benchmarks from 2019 for 86 of the 100 most searched positions for which the public benchmark was available. For each position, we calculate the difference between the proprietary and the public online benchmark as (proprietary-public)*100/public. The salaries from the public online benchmark are obtained using The Wayback Machine.

Table C.1: Comparison of Firms in Our Sample vs. Representative Sample of U.S. Firms

	Percentile				
	10th	25th	50th	75th	90th
Number of Employees					
Our Sample	68	109	225	529	1,159
U.S. Representative Sample	22	26	39	79	189
Salary (Annual $)					
Our Sample	20,071	25,468	38,177	64,604	105,689
U.S. Representative Sample	9,820	19,200	36,000	63,200	104,000

Notes: *U.S. Representative Sample* corresponds to the statistics of firms taken from the most recent year (2013) of Song et al. (2019). *Our Sample* denotes to the sample of 2,051 firms in our dataset for the earliest period for which data are available (January 2016). To make the statistics more comparable across the two samples, we match the sample restrictions from Song et al. (2019) by excluding firms with less than 20 employees and employees younger than 20 years old or older than 60 years old. Our *Salary* statistics are based off the distribution of individual annual base salaries across employees in all firms. Song et al. use earnings. To make the two samples more comparable, we converted the salary statistics in our sample to 2013 dollars using the PCE deflator published by the Bureau of Economic Analysis.

Table C.2: Comparison of Sector Representation in Our Sample vs. U.S. Employees & Firms

	Firms (%)		Employees (%)	
	(1)	(2)	(3)	(4)
Sector	Our Sample	U.S.	Our Sample	U.S.
Agriculture, Forestry, Fishing and Hunting	0.25	0.37	0.37	0.13
Mining, Quarrying, and Oil and Gas Extraction	0.44	0.32	0.11	0.45
Utilities	0.44	0.10	0.34	0.50
Construction	2.33	11.58	0.51	5.08
Manufacturing	22.22	4.10	21.94	9.12
Wholesale Trade	8.87	4.92	14.24	4.76
Retail Trade	3.90	10.70	7.82	12.21
Transportation and Warehousing	2.20	3.05	1.25	3.78
Information	2.77	1.32	3.71	2.73
Finance and Insurance	13.91	3.94	11.10	4.98
Real Estate and Rental and Leasing	3.02	5.11	1.58	1.67
Professional, Scientific, and Technical Services	11.83	13.39	8.56	6.93
Management of Companies and Enterprises	1.01	0.45	1.29	2.69
Administrative and Support and Waste Management	4.59	5.74	6.58	9.25
Educational Services	2.64	1.54	2.51	2.87
Health Care and Social Assistance	11.33	10.81	13.42	15.74
Arts, Entertainment, and Recreation	0.57	2.15	0.40	1.84
Accommodation and Food Services	1.95	8.91	1.59	10.96
Other Services (except Public Administration)	5.73	11.50	2.70	4.30

Notes: Percent of firms and employees in each sector in our sample vs. in the U.S. The NAICS code *Public Administration* excluded from statistics of our sample because the Census does not report data for that code.

Table C.3: Most Common Searched Position Titles

Position Title	(1) Searched	(2) Non-Searched	(3) Non-Searchable
Bank Teller	539 [12]	287 [24]	1,976 [87]
Customer Service Representative	468 [44]	4,401 [170]	4,012 [385]
Security Guard	286 [6]	139 [44]	6,263 [95]
Hotel Cleaner	208 [2]	379 [5]	1,058 [17]
Legal Associate Specialist	163 [1]	7 [4]	14 [9]
Hand Packer	155 [4]	234 [17]	1,957 [55]
Patient Care Coordinator	117 [3]	103 [14]	133 [29]
Receptionist	93 [15]	310 [86]	2,911 [238]
Cook	86 [6]	334 [21]	1,606 [85]
Waiter/Waitress	84 [7]	1,113 [18]	2,986 [87]
Delivery Driver	79 [5]	34 [9]	744 [26]
Dish Washer/Plate Collector/Table Top Cleaner	69 [5]	187 [18]	1,350 [67]
Medical Assistant	69 [10]	370 [17]	889 [55]
Welder	66 [8]	112 [27]	652 [59]
Cashier	65 [2]	175 [11]	2,706 [48]
Registered Nurse	64 [11]	244 [22]	2,699 [110]
Assembler	60 [9]	606 [26]	3,823 [90]
Other Housekeeper and Related Worker	59 [5]	173 [17]	948 [63]
Software Developer/Programmer	59 [23]	403 [78]	1,285 [173]
Warehouse Laborer	59 [10]	761 [43]	3,025 [116]
Mammographer	55 [1]	9 [1]	3 [2]
Nursing Assistant	51 [4]	662 [13]	7,346 [65]
Bartender/Mixologist	49 [2]	228 [12]	611 [46]
Production Operations Engineer	49 [1]	41 [16]	68 [29]
Licensed Practical Nurse	48 [9]	189 [23]	1,605 [69]
Sales Manager	48 [18]	166 [67]	693 [181]
General Practitioner/Physician	46 [2]	143 [17]	340 [28]
Lawyer	43 [5]	17 [10]	268 [52]
Ophthalmic Technician	42 [2]	4 [1]	34 [4]
Business Development Specialist	41 [2]	124 [27]	447 [41]
Warehouse Manager	40 [7]	133 [23]	430 [72]
Other Social Work and Counseling Professional	39 [1]	1 [1]	32 [9]
Building Caretaker/Watchman	38 [2]	288 [59]	917 [139]
Operations Officer	37 [2]	73 [18]	108 [36]
Shipping Clerk	37 [4]	39 [19]	218 [63]

Notes: New hires in each position [firms hiring in each position]. Tabulations across all new hires for the 35 Searched *Position Titles* with the most new hires.

Table C.4: Position Characteristics

Position Title	(1) High Skill	(2) High Educ.	(3) $\bar{Age} \geq 31$	(4) $\bar{Inc.} \geq \$30K$
Customer Service Representative	Y	N	Y	Y
Security Guard	Y	N	Y	Y
Legal Associate Specialist	Y	Y	Y	Y
Patient Care Coordinator	Y	N	Y	Y
Medical Assistant	Y	Y	Y	Y
Welder	Y	N	Y	Y
Registered Nurse	Y	Y	Y	Y
Software Developer/Programmer	Y	Y	Y	Y
Mammographer	Y	Y	Y	Y
Production Operations Engineer	Y	Y	Y	Y
Licensed Practical Nurse	Y	Y	Y	Y
Sales Manager	Y	Y	Y	Y
General Practitioner/Physician	Y	Y	Y	Y
Lawyer	Y	Y	Y	Y
Ophthalmic Technician	Y	Y	Y	Y
Business Development Specialist	Y	Y	Y	Y
Warehouse Manager	Y	N	Y	Y
Other Social Work and Counseling Professional	Y	Y	Y	Y
Operations Officer	Y	Y	Y	Y
Bank Teller	N	N	N	N
Hotel Cleaner	N	N	Y	N
Hand Packer	N	N	Y	N
Receptionist	N	N	N	N
Cook	N	N	Y	N
Waiter/Waitress	N	N	N	N
Delivery Driver	N	N	Y	N
Dish Washer/Plate Collector/Table Top Cleaner	N	N	N	N
Cashier	N	N	N	N
Assembler	N	N	Y	N
Other Housekeeper and Related Worker	N	N	Y	N
Warehouse Laborer	N	N	Y	N
Nursing Assistant	N	N	Y	N
Bartender/Mixologist	N	N	Y	N
Building Caretaker/Watchman	N	N	Y	N
Shipping Clerk	N	N	Y	N

Notes: List of 35 position titles with the most hires in the Searched group. A position is classified as low-skill if: (i) it is classified as low education; (ii) the average age is below 31 years; (iii) the average salary is less than $30,000.

D Effects on Salary Dispersion: Additional Results and Robustness Checks

D.1 Dispersion Around Different Benchmarks

Figure D.1 reproduces Figure 2 using a narrower bin width (+/- 1%). This figure highlights bunching at exactly the median. Panel A of Figure D.2 replicates Panel A of Figure 2, which presents the dispersion of salaries around the median benchmark in Searched positions both before and after onboarding. In panels B through F of Figure D.2, we replicate the analysis but, instead of analyzing the difference between the salaries and the median benchmark, we look at the difference with respect to the other points of the benchmark distribution (e.g., the 10th percentile benchmark). Panel B shows the change in percent difference from the 10th percentile benchmark, both before (denoted by solid gray bins) and after (hollow red bins) the onboarding date. Before onboarding, salaries in the Searched positions were on average 34.3 pp from the 10th percentile benchmark. After the onboarding date, the average distance to the benchmark fell to 29.2 pp (p-value<0.001). While this corresponds to a reduction in salary dispersion, Panel B shows that this decrease is driven by a higher share of salaries 10-20% above the 10th percentile benchmark, rather than a higher share of salaries at the 10th percentile benchmark itself. The same result holds when we analyze compression around the other points of the distribution: the 25th percentile benchmark (Panel C), the average benchmark (Panel D), the 75th percentile benchmark (Panel E), and the 90th percentile benchmark (Panel F).

D.2 Effects on Composition of New Hires

The types of new hires that join a firm after a firm gains access to the salary benchmark may shift. To test this hypothesis, Table D.1 presents difference-in-differences estimates using the characteristics of the employee (instead of their starting salary) as the dependent variable. In column (1), the dependent variable is an indicator for whether the employee is female. In column (2), the outcome is an indicator variable for whether the employee was hired through an hourly contract. In column (3), the dependent variable is the employee's age in years. The coefficients on the post-treatment coefficients are close to zero and statistically insignificant, suggesting that access to benchmarking did not have a significant effect on the composition of new hires. However, this result must be taken with a grain of salt for at least two reasons. First, we do not have sufficient precision to rule out modest effect sizes. Second, these results cover just three basic employee characteristics (gender, contract type and age).

D.3 Other Robustness Checks

Table D.2 and Table D.3 present the same difference-in-differences regressions as in Table 3, but with different clustering for the standard errors. Table D.2 clusters at the firm level and Table D.3 clusters at the firm-position level. The results are largely robust across all these different specifications.

Table D.4 presents additional robustness checks for our difference-in-differences estimates for the effect on absolute dispersion from the benchmark. Column (1) replicates the baseline specification from column (1) of Table 3. Columns (2) through (7) each change a different aspect of the baseline specification. In columns (2) and (3), we use alternative benchmark filters to compute dispersion from the benchmark. In our main specification we use benchmarks for a given position title filtered by state and sector when that benchmark is based on more than 30 employees, and no filters otherwise. In column (2), we use only the unfiltered benchmark and in column (3) we use the filtered benchmark only when it based on more than 100 employees and no filters otherwise. In column (4), we include job titles with low "match scores" (the quality of the mapping between a firm-specific job title and a position title) that are filtered out of the baseline specification. Column (5) uses only Non-Searched and Non-Searchable new hires after September 2019, the start of our search data. Column (6) includes data from August 2020 through July 2021.[67] Finally, column (7) excludes all HR positions (2.28% of the sample). Intuitively, HR employees may be looking their own salaries up due to curiosity, or for their own salary planning, rather than to negotiate with the new hires in HR positions. The results from the baseline specification in column (1) of Table D.4 are qualitatively and quantitatively similar to the results in each of the alternative specifications in columns (2) through (7).

D.4 Additional Heterogeneity Analysis

Figure D.3 presents the event-study results for low-skill and high-skill positions. The panels in the top row of Figure D.3 (i.e., panels A and B) correspond to the low-skill positions, while the panels below (i.e., panels C and D) correspond to high-skill positions. Before firms had access to the tool (i.e., the left side of each panel), for the Searched positions, there was more dispersion among high-skill positions (22.6%) than in low-skill positions (14.2%). The difference-in-differences comparison for Searched vs. Non-Searched low-skill positions (Panel A) suggest the benchmark tool reduced the salary dispersion from 14.2 pp to 7.6 pp (p-value<0.001), equivalent to a 46.5% reduction. For Searched vs. Non-Searchable low-

[67] Starting September 2020, the payroll company switched to a new taxonomy that expanded the number of position titles. Since our main sample stops at March 2020, our baseline results are not affected by this change. However, the results from this specification include both the old and the new taxonomies.

skill positions (Panel B), salary dispersion drops from 14.2 pp to 7.6 pp (p-value<0.001), or a 46.5% reduction. For high-skill positions, the difference-in-differences comparison for Searched vs. Non-Searched (Panel C) suggests that the benchmark tool reduced the salary dispersion from 22.6 pp to 18.9 pp (p-value = 0.019), equivalent to a 16.4% reduction. Finally, comparing Searched vs. Non-Searchable high-skill positions (Panel D) suggests a drop in dispersion from 22.6 pp to 17.2 pp (p-value < 0.001), equivalent to a 23.9% reduction.

Figure D.4 displays the event-study analysis by low-dispersion and high-dispersion of market-level salaries. For each combination of position, state and industry, we calculate a metric known as "market dispersion," which is the difference between the 90th and 10th percentiles of market benchmarks, as indicated in the benchmarking tool. Intuitively, significant variations in salaries within a cell imply a substantial variation in skills. The results from Figure D.4 are qualitatively similar to our high- vs. low-skill heterogeneity analysis, though less stark.

Figure D.6 analyzes heterogeneity by monopsonistic power, using the HHI measure computed by Azar et al. (2022). The original measures are based on job ads data from LightCast (previously BurningGlass) and are at the SOC-commuting zone level. We aggregate to the 2-digit SOC major groups and state level and merge with our data. We split observations by whether the HHI index is above or below the median value in our entire sample. The results from Figure D.6 constitute suggestive evidence indicating that the impact of salary benchmarking is more pronounced in more competitive labor markets.

Figure D.5 breaks down the effect on salary dispersion by gender. Panel A and Panel B reproduce Figure 3 with the male subsample, and Panel C and Panel D reproduce Figure 3 with the female subsample. we find differences by gender that are economically small and statistically insignificant.

Figure D.1: The Effects of Benchmarking on Dispersion around the Benchmark: Non-Parametric Analysis

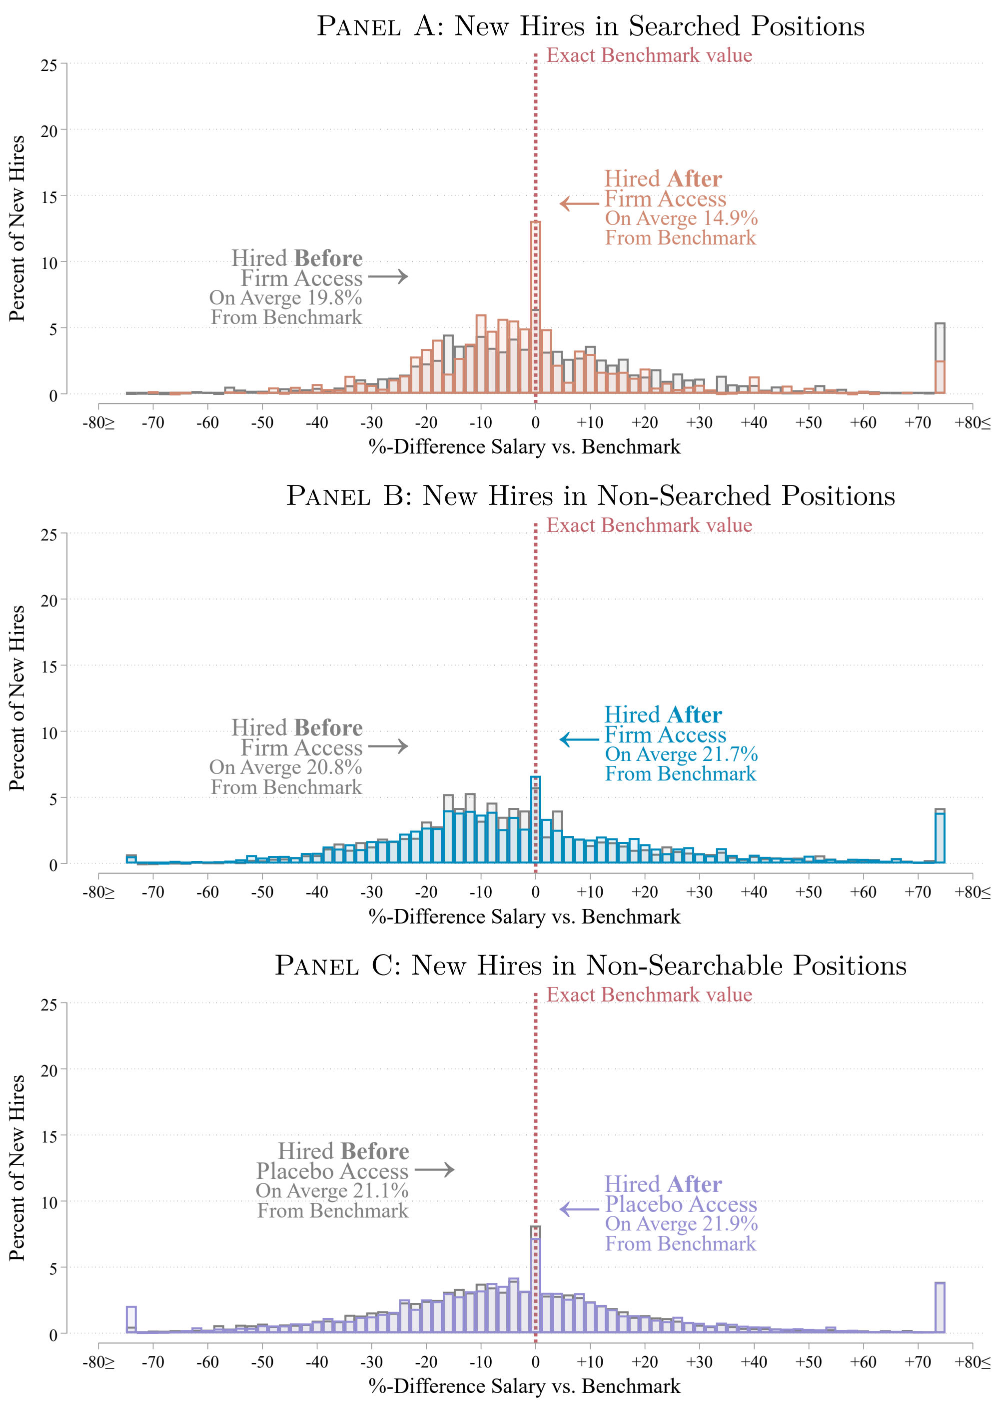

Notes: It reproduces Figure 2 using a narrower bin width (+/- 1%).

Figure D.2: Effects of the Compensation Benchmark on New Hire Salaries Relative to Different Percentiles of the Benchmark: Non-Parametric Analysis

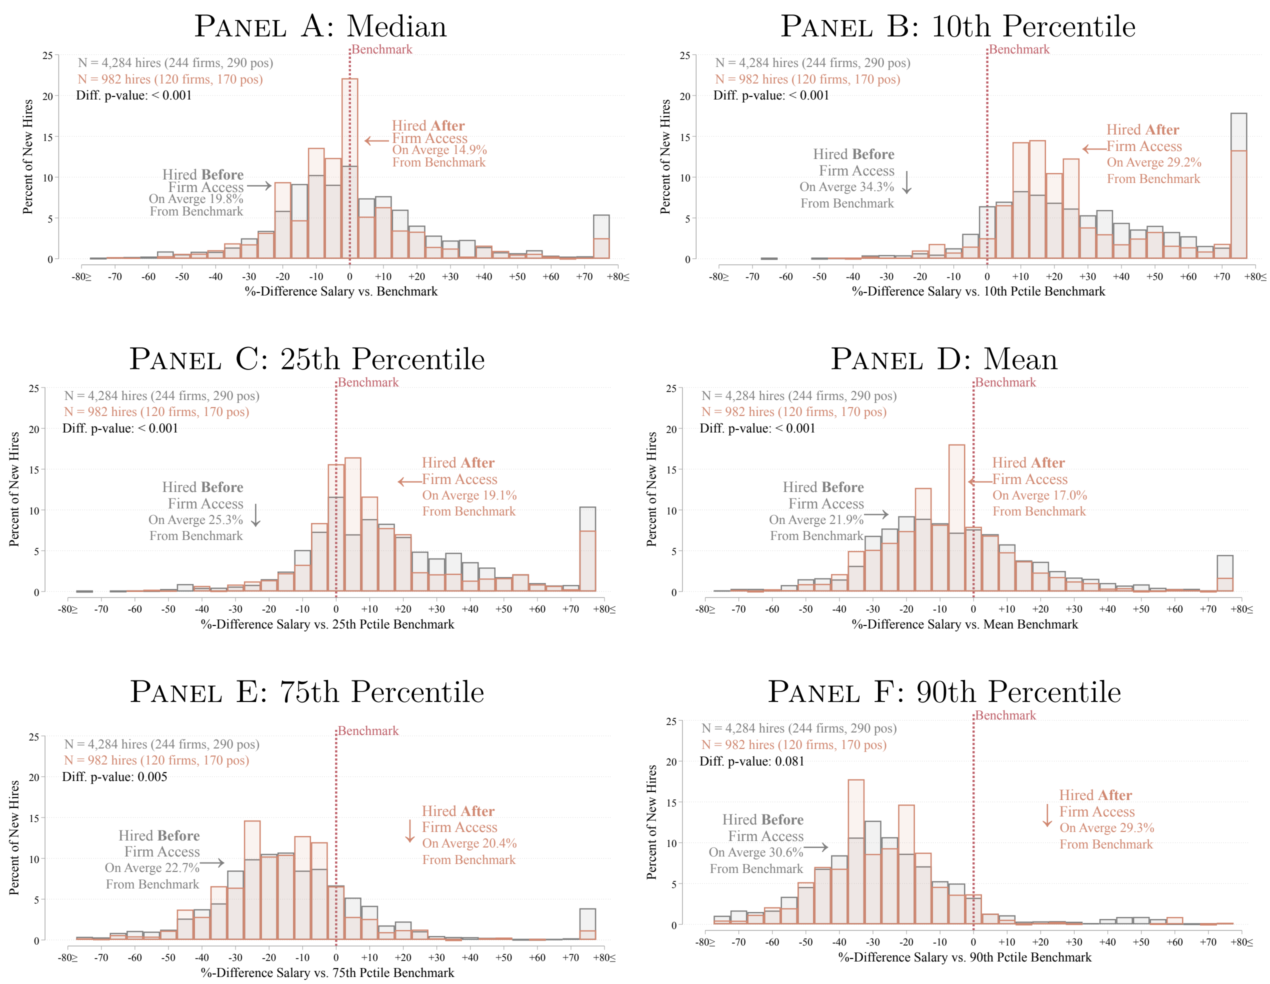

Notes: Panel A is a reproduction of Panel A of Figure 2. All other panels are identical, but using absolute dispersion from the specified percentile of the benchmark distribution rather than the median. For more details, see notes to Figure 2.

Figure D.3: Heterogeneity by Skill: Event-Study Analysis

PANEL A: Low-Skill: Searched vs. Non-Searched

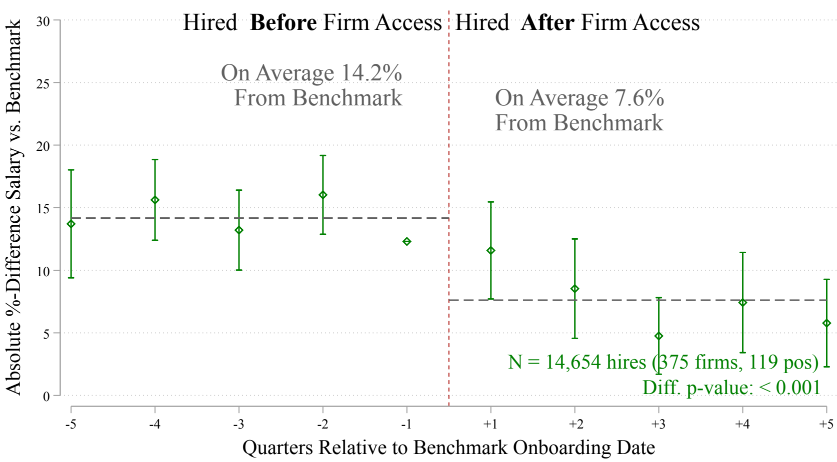

PANEL B: Low-Skill: Searched vs. Non-Searchable

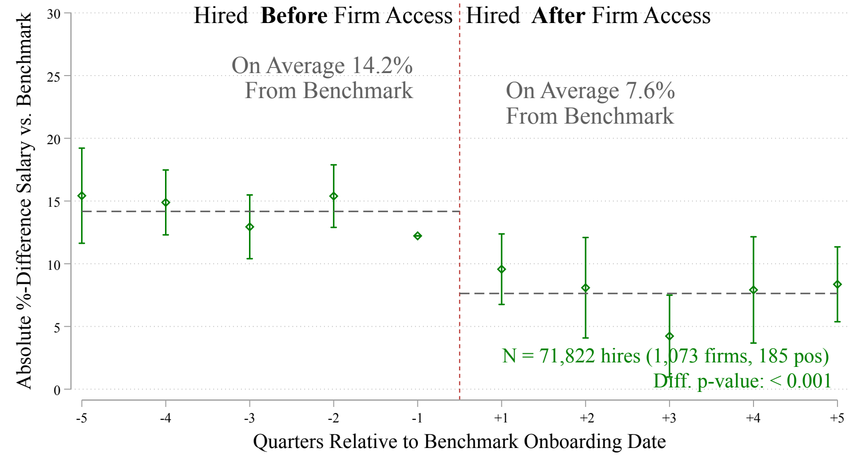

PANEL C: High-Skill: Searched vs. Non-Searched

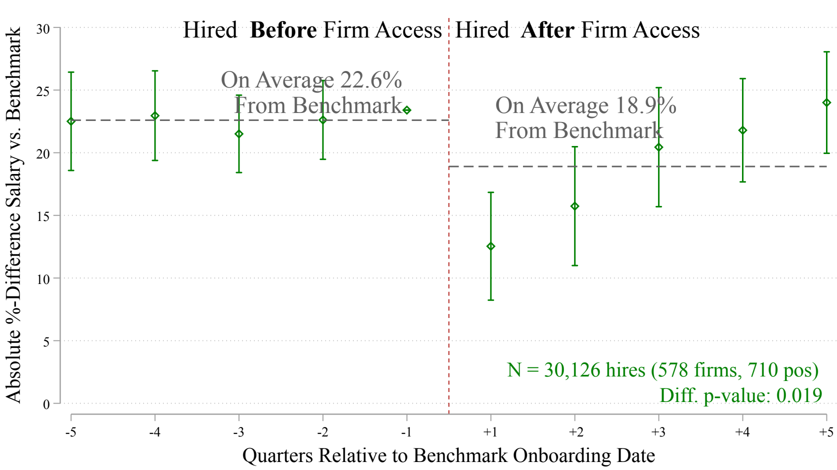

PANEL D: High-Skill: Searched vs. Non-Searchable

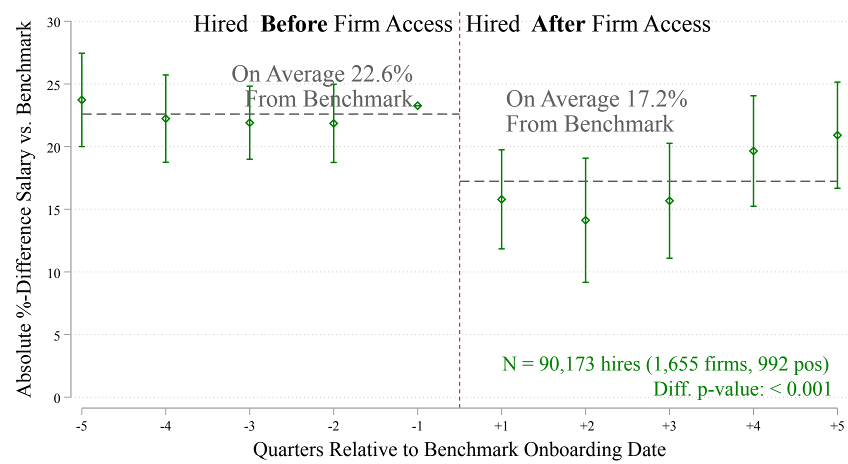

Notes: Panels A and C are a reproduction of Panel C from Figure 3, and Panels B and D are a reproduction of Panel D, but for the specified sub-samples. *Skill* is defined in Section 2.1. See the notes of Figure 3 for more details.

Figure D.4: Heterogeneity by Market Dispersion: The Effects on Pay Dispersion

Panel A: Low Market Dispersion: Searched vs. Non-Searched

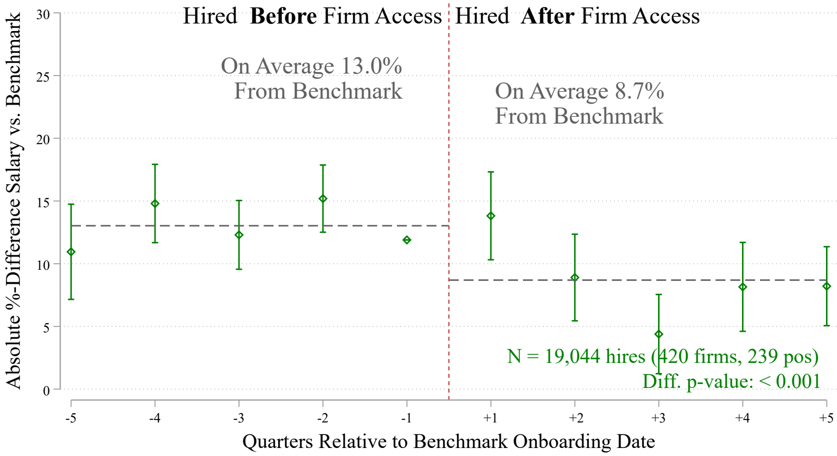

Panel B: Low Market Dispersion: Searched vs. Non-Searchable

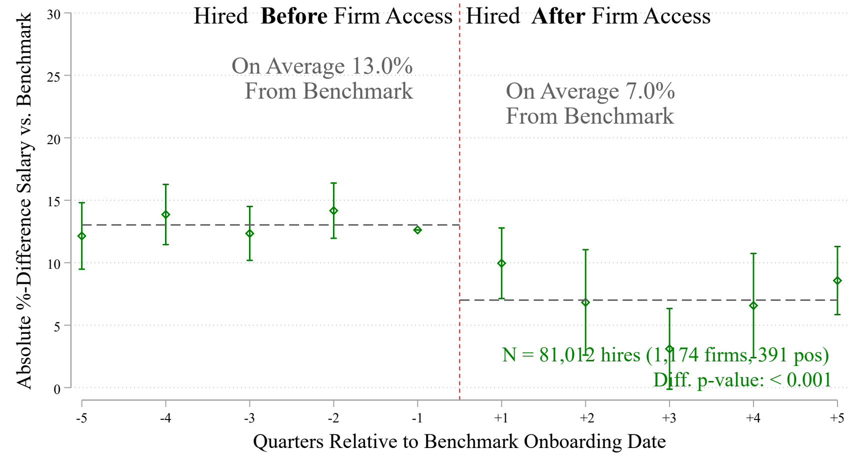

Panel C: High Market Dispersion: Searched vs. Non-Searched

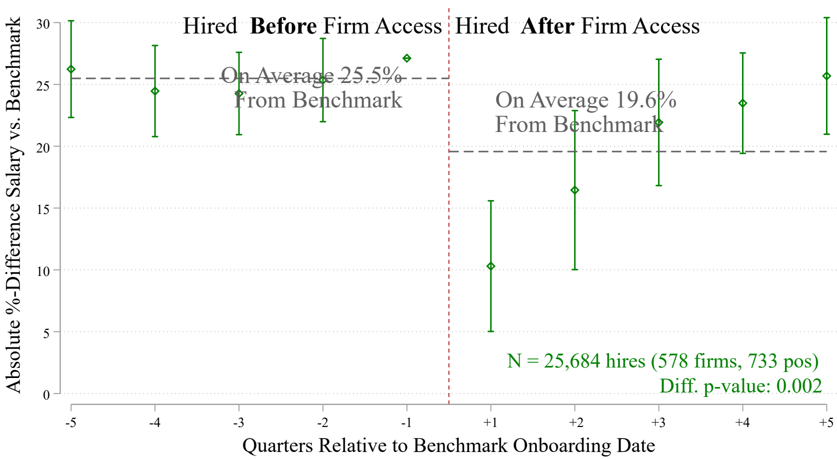

Panel D: High Market Dispersion: Searched vs. Non-Searchable

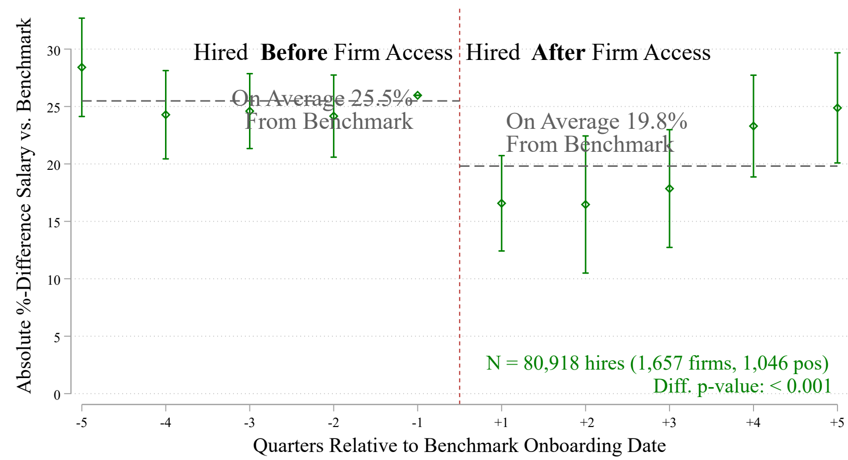

Notes: Panels A and C are a reproduction of Panel C from Figure 3, and Panels B and D are a reproduction of Panel D, but for the specified sub-samples split by the market dispersion. See the notes of Figure 3 for more details.

Figure D.5: Heterogeneity by Gender: Event-Study Analysis

Panel A: Male: Searched vs. Non-Searched

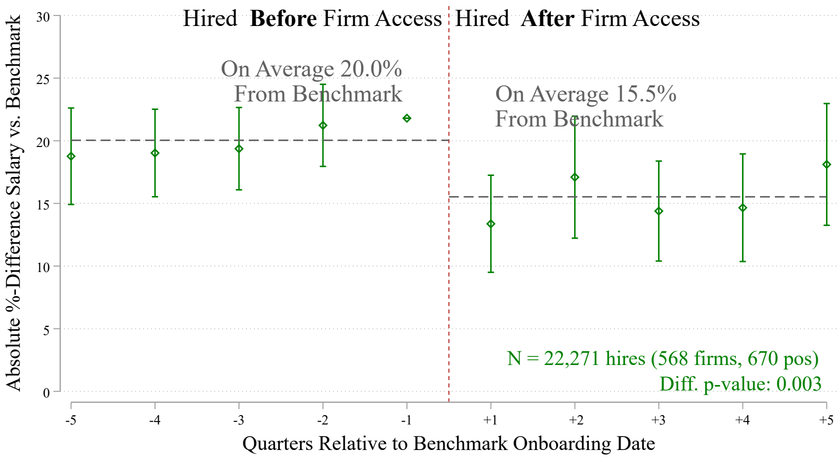

Panel B: Male: Searched vs. Non-Searchable

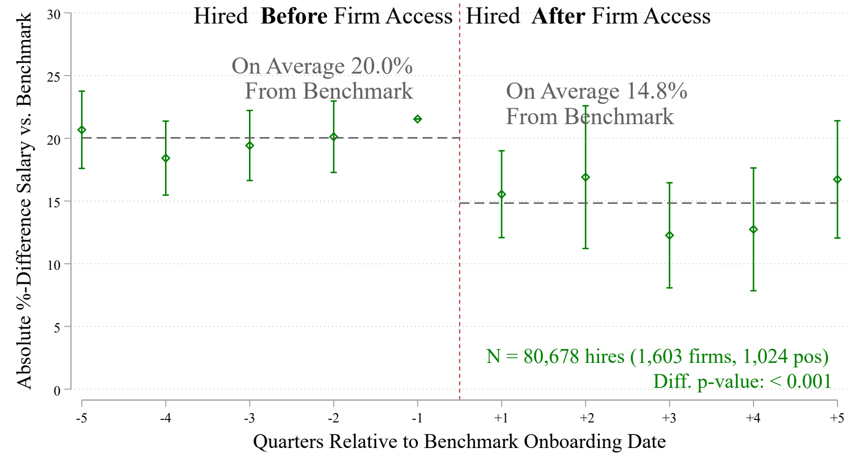

Panel C: Female: Searched vs. Non-Searched

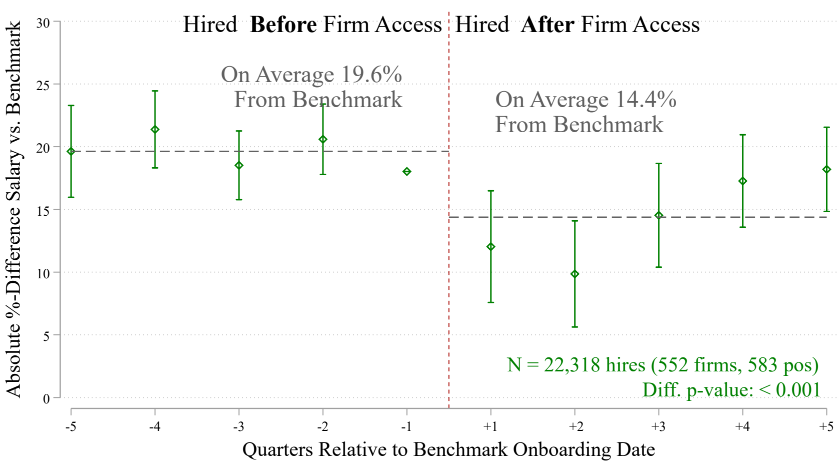

Panel D: Female: Searched vs. Non-Searchable

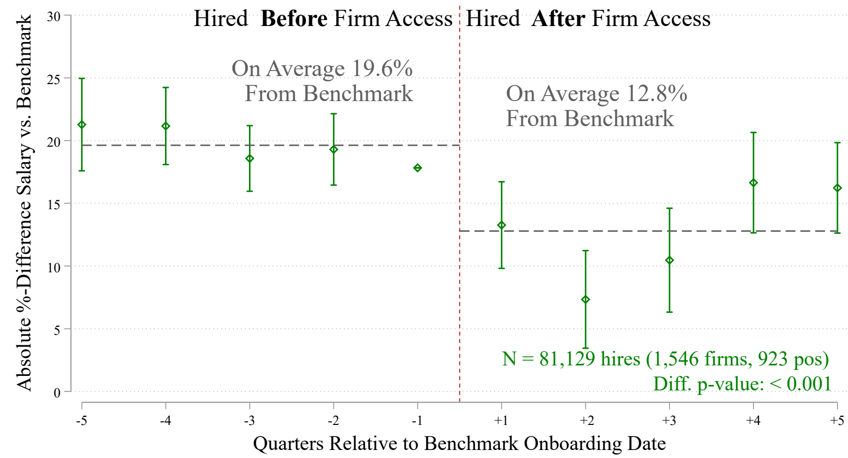

Notes: Panels A and C are a reproduction of Panel C from Figure 3, and Panels B and D are a reproduction of Panel D, but for the specified sub-samples split by gender. See the notes of Figure 3 for more details.

Figure D.6: Heterogeneity by imputed HHI: The Effects on Pay Dispersion (Low Skill)

Panel A: Low HHI: Searched vs. Non-Searched

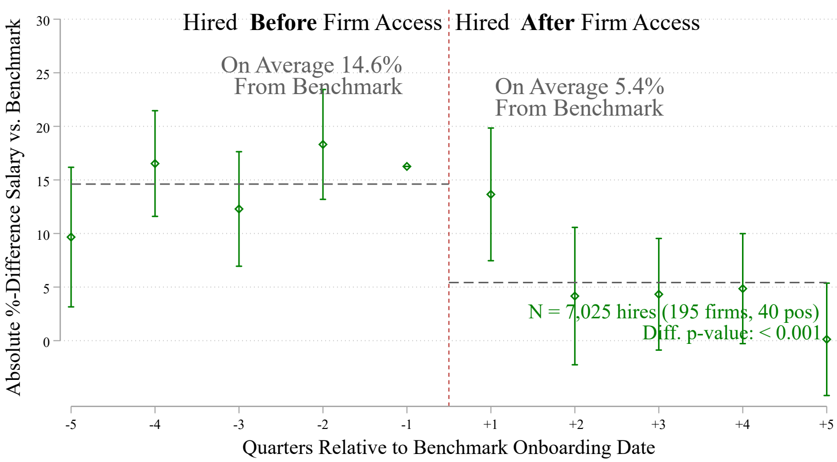

Panel B: Low HHI: Searched vs. Non-Searchable

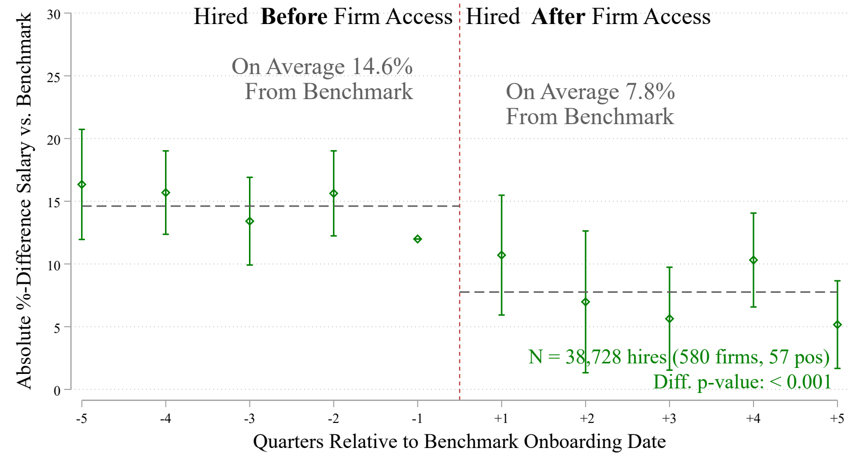

Panel C: High HHI: Searched vs. Non-Searched

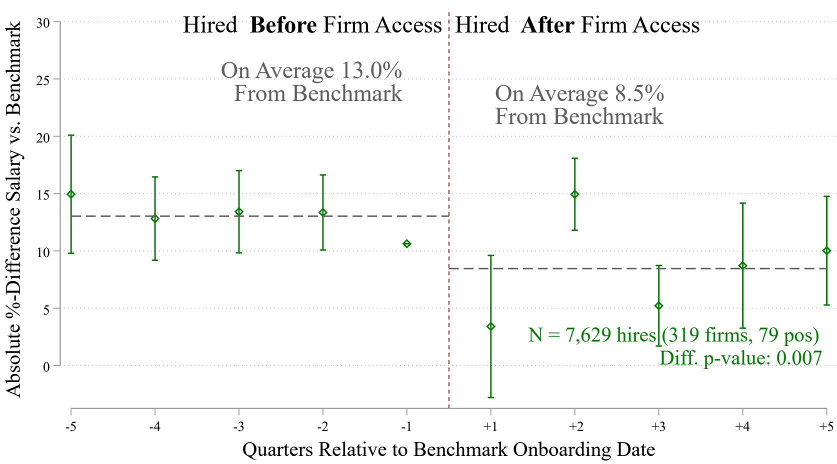

Panel D: High HHI: Searched vs. Non-Searchable

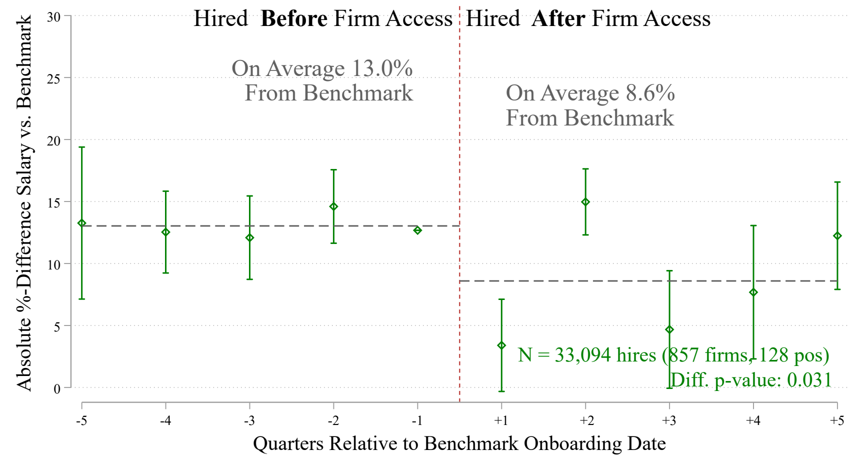

Notes: Panels A and C are a reproduction of Panel C from Figure 3, and Panels B and D are a reproduction of Panel D, but for the specified sub-samples split by the imputed HHI. See the notes of Figure 3 for more details.

Table D.1: The Effects of Benchmarking on Employee Composition

	(1) Female	(2) Hourly	(3) Age
Panel A: Post-treatment			
Searched vs. Non-Searched	-2.306	0.487	-0.749
	(2.424)	(1.722)	(0.740)
Searched vs. Non-Searchable	-1.987	2.708	-0.117
	(2.338)	(1.654)	(0.761)
Panel B: Pre-treatment			
Searched vs. Non-Searched	-1.051	1.605	-0.409
	(2.133)	(1.640)	(0.606)
Searched vs. Non-Searchable	0.584	3.241**	0.0461
	(2.009)	(1.651)	(0.529)
Mean Dep. Var. (Baseline)	54.870	76.482	34.598
Observations			
Searched	5,253	5,253	5,253
Non-Searched	39,527	39,527	39,527
Non-Searchable	156,734	156,734	156,734

Notes: Significant at *10%, **5%, ***1%. Robust standard errors in parentheses. All columns follow the specification of column (1) from Table 3, with the exception that here we exclude the additional controls. The dependent variables are a dummy equal to 100 if a new hire is *Female* and zero otherwise in column (1), a dummy equal to 100 if a new hire is an *Hourly* worker and zero otherwise in column (2), and a new hire's *Age* in column (3).

Table D.2: The Effects of Benchmarking on Salary Dispersion: Clustering at the Firm Level

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	$\|\%\Delta\|$	$\|log\Delta\|$	$\|\%\Delta\| > 10$	$\|\%\Delta\|$	$\|\%\Delta\|$	$\|\%\Delta\|$	$\|\%\Delta\|$	$\|\%\Delta\|$	$\|\%\Delta\|$	$\|\%\Delta\|$	$\|\%\Delta\|$	$\|\%\Delta\|$
Panel A: Post-treatment												
Searched vs. Non-Searched	-4.775***	-5.155***	-16.270***	-5.148***	-4.775***	-4.786***	-5.324**	-4.950*	-4.421***	-4.887***	-4.880**	-4.564***
	(1.530)	(1.740)	(5.311)	(1.763)	(0.906)	(1.604)	(2.262)	(2.631)	(1.505)	(1.567)	(2.283)	(1.554)
Searched vs. Non-Searchable	-6.149***	-7.118***	-13.861**	-6.836***	-6.149***	-6.128***	-7.494***	-7.450**	-5.714***	-6.163***	-5.044**	-5.934***
	(1.545)	(1.712)	(5.919)	(1.715)	(0.824)	(1.561)	(2.104)	(3.582)	(1.533)	(1.575)	(2.143)	(1.603)
Panel B: Pre-treatment												
Searched vs. Non-Searched	-0.346	-0.129	-5.872	-0.233	-0.346	-0.488	-1.646	-2.062	-0.714	-0.144	-2.205	-0.199
	(1.398)	(1.517)	(4.324)	(1.559)	(0.751)	(1.424)	(2.035)	(1.750)	(1.379)	(1.441)	(1.758)	(1.416)
Searched vs. Non-Searchable	-0.310	0.156	-4.221	-0.513	-0.310	-0.318	0.021	-1.029	0.241	-0.247	-0.754	-0.500
	(1.294)	(1.406)	(3.710)	(1.463)	(0.643)	(1.304)	(1.978)	(1.610)	(1.285)	(1.311)	(1.689)	(1.301)
Winsorizing at +/- 100%				✓								
No Clustering					✓							
No Additional Controls						✓						
No Position FE							✓					
Firm FE								✓				
Exclude High-Tip Jobs									✓			
Searched Positions Only										✓		
No Re-weighting											✓	
Ages 21-60												✓
Mean Dep. Var. (Baseline)	19.812	20.590	63.732	21.004	19.812	19.812	19.812	19.812	19.430	19.812	19.802	19.903
Observations												
Searched	5,253	5,253	5,253	5,253	5,253	5,253	5,266	5,262	5,105	5,253	5,331	4,611
Non-Searched	39,527	39,527	39,527	39,527	39,527	39,527	39,686	39,673	37,841	34,954	39,810	34,338
Non-Searchable	156,734	156,734	156,734	156,734	156,734	156,734	156,865	156,817	148,521	127,145	157,018	135,051

Notes: Significant at *10%, **5%, ***1%. Standard errors clustered at the firm level in parentheses. Each column corresponds to two regressions: one for Searched vs. Non-Searched new hires and one for Searched vs. Non-Searchable new hires. Post-treatment coefficients in Panel A refer to parameters α_1^k from equation (2), while pre-treatment coefficients in Panel B refer to parameters α_3^k from equation (2) (see Section 3.2 for details). All columns include year fixed effects. In columns (1) and (4)–(12) the dependent variable is the absolute percent difference between the annual base salary and median benchmark (Δ). The dependent variable in column (2) is the log of Δ and in column (3) is a dummy that equals 100 if $|\%\Delta|$ is greater than 10% and zero otherwise. We multiply $\%\Delta$ and $log(\Delta)$ by 100 so that the effects can be interpreted as percentage points. Δ is winsorized to $\pm$ 75 except in column (4) where it is winsorized to $\pm$ 100. All columns except (6) include additional controls (female dummy, high education dummy, hourly dummy, age, position tenure). Column (7) excludes position fixed effects. Column (8) includes firm fixed effects instead of position fixed effects. Column (9) excludes the three positions where gross pay most exceeds base pay: Waiter/Waitress, Chauffeur, and Bartender/Mixologist. Column (10) restricts the sample to only positions of Non-Searched or Non-Searchable new hires in positions that are searched and hired by firms in the data.

Table D.3: The Effects of Benchmarking on Salary Dispersion: Clustering at the Firm-Position Level

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	$\lvert\%\Delta\rvert$	$\lvert log\Delta\rvert$	$\lvert\%\Delta\rvert > 10$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$
Panel A: Post-treatment												
Searched vs. Non-Searched	-4.775***	-5.155***	-16.270***	-5.148***	-4.775***	-4.786***	-5.324**	-4.950**	-4.421***	-4.887***	-4.880**	-4.564***
	(1.551)	(1.763)	(5.439)	(1.806)	(0.906)	(1.610)	(2.227)	(2.453)	(1.535)	(1.579)	(2.240)	(1.560)
Searched vs. Non-Searchable	-6.149***	-7.118***	-13.861**	-6.836***	-6.149***	-6.128***	-7.494***	-7.450**	-5.714***	-6.163***	-5.044**	-5.934***
	(1.524)	(1.730)	(5.497)	(1.769)	(0.824)	(1.532)	(2.139)	(3.147)	(1.519)	(1.549)	(2.072)	(1.578)
Panel B: Pre-treatment												
Searched vs. Non-Searched	-0.346	-0.129	-5.872	-0.233	-0.346	-0.488	-1.646	-2.062	-0.714	-0.144	-2.205	-0.199
	(1.252)	(1.416)	(3.941)	(1.397)	(0.751)	(1.272)	(1.927)	(1.506)	(1.237)	(1.290)	(1.784)	(1.275)
Searched vs. Non-Searchable	-0.310	0.156	-4.221	-0.513	-0.310	-0.318	0.021	-1.029	0.241	-0.247	-0.754	-0.500
	(1.198)	(1.354)	(3.544)	(1.361)	(0.643)	(1.205)	(1.860)	(1.295)	(1.197)	(1.218)	(1.679)	(1.231)
Winsorizing at +/- 100%				✓								
No Clustering					✓							
No Additional Controls						✓						
No Position FE							✓					
Firm FE								✓				
Exclude High-Tip Jobs									✓			
Searched Positions Only										✓		
No Re-weighting											✓	
Ages 21-60												✓
Mean Dep. Var. (Baseline)	19.812	20.590	63.732	21.004	19.812	19.812	19.812	19.812	19.430	19.812	19.802	19.903
Observations												
Searched	5,253	5,253	5,253	5,253	5,253	5,253	5,266	5,262	5,105	5,253	5,331	4,611
Non-Searched	39,527	39,527	39,527	39,527	39,527	39,527	39,686	39,673	37,841	34,954	39,810	34,338
Non-Searchable	156,734	156,734	156,734	156,734	156,734	156,734	156,865	156,817	148,521	127,145	157,018	135,051

Notes: Significant at *10%, **5%, ***1%. Standard errors clustered at the firm-position level in parentheses. Each column corresponds to two regressions: one for Searched vs. Non-Searched new hires and one for Searched vs. Non-Searchable new hires. Post-treatment coefficients in Panel A refer to parameters α_1^k from equation (2), while pre-treatment coefficients in Panel B refer to parameters α_3^k from equation (2) (see Section 3.2 for details). All columns include year fixed effects. In columns (1) and (4)–(12) the dependent variable is the absolute percent difference between the annual base salary and median benchmark (Δ). The dependent variable in column (2) is the log of Δ and in column (3) is a dummy that equals 100 if $|\%\Delta|$ is greater than 10% and zero otherwise. We multiply $\%\Delta$ and $log(\Delta)$ by 100 so that the effects can be interpreted as percentage points. Δ is winsorized to $\pm$ 75 except in column (4) where it is winsorized to $\pm$ 100. All columns except (6) include additional controls (female dummy, high education dummy, hourly dummy, age, position tenure). Column (7) excludes position fixed effects. Column (8) includes firm fixed effects instead of position fixed effects. Column (9) excludes the three positions where gross pay most exceeds base pay: Waiter/Waitress, Chauffeur, and Bartender/Mixologist. Column (10) restricts the sample to only positions of Non-Searched or Non-Searchable new hires in positions that are searched and hired by firms in the data.

Table D.4: The Effects of Benchmarking on Salary Dispersion: Additional Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$	$\lvert\%\Delta\rvert$
Panel A: Post-treatment							
Searched vs. Non-Searched	-4.775***	-6.876***	-5.167***	-4.231***	-3.588***	-4.786***	-4.996***
	(1.143)	(1.870)	(1.113)	(1.104)	(1.218)	(1.094)	(1.153)
Searched vs. Non-Searchable	-6.149***	-8.163***	-6.018***	-5.701***	-6.288***	-6.582***	-6.420***
	(1.070)	(1.498)	(1.041)	(1.064)	(1.152)	(1.008)	(1.096)
Panel B: Pre-treatment							
Searched vs. Non-Searched	-0.346	-0.549	-0.669	0.152	-0.301	-0.369	-0.578
	(1.167)	(1.594)	(1.172)	(1.163)	(1.134)	(1.168)	(1.194)
Searched vs. Non-Searchable	-0.310	-0.154	-0.703	0.293	0.207	-0.315	-0.454
	(1.055)	(1.432)	(1.041)	(1.109)	(1.032)	(1.055)	(1.083)
No Filters		✓					
Filtered Benchmark $\geq$ 100			✓				
Include Match Outliers				✓			
Restricted Sample					✓		
After Aug-2020						✓	
Exclude HR Positions							✓
Mean Dep. Var. (Baseline)	19.812	24.290	20.163	20.478	19.812	19.812	19.790
Observations							
Searched	5,253	5,253	5,253	6,150	5,246	5,414	5,080
Non-Searched	39,527	39,527	39,527	50,943	15,958	46,435	38,156
Non-Searchable	156,734	156,734	156,734	196,768	83,348	160,595	151,778

Notes: Column (1) follows the specification of column (1) from Table 3. Column (2) uses the same specification as column (1), except using absolute dispersion from the unfiltered median benchmark, as opposed to using the state and sector filtered benchmark when available, as the outcome. Column (3) uses only filtered benchmarks computed using 100 or more employees, as opposed to the baseline threshold of 30 employees, and unfiltered benchmarks otherwise. Column (4) drops new hires who's organization specific job title has a low match score to the designated position title (scores less than the 20th percentile of scores in that quarter). Column (5) using the *Restricted Sample* uses only control observations after September 2019, the start of our search data. Column (6) adds data from Aug-2020 to July-2021 to the sample. Column (7) excludes new hires in HR positions. See Table 3 for more details.

E Effects on Average Salary and Retention: Additional Results and Robustness Checks

E.1 Main Robustness Checks

Regarding the effects of salary benchmarking on the average salary, the difference-in-differences estimates are presented in Table E.1. The post-treatment coefficients (α_1^k, from equation (1)) are presented in Panel A. Column (1) of Table E.1 corresponds to the baseline specification. The post-treatment coefficients are positive: 0.003 log points (p-value=0.745) when using Non-Searched as a control group and 0.019 log points (p-value=0.918) when using Non-Searchable as control group.

Columns (2) through (11) of Table E.1 are identical to column (1), except that they change a different feature of the baseline specification. In column (2) we use salary as the dependent variable: i.e., in \$s, without the log transformation. The results from column (2) are qualitatively consistent with the results from column (1): the post-treatment coefficients are modest (-\$724.51 and \$271.15 for the comparison to Non-Searched and Non-Searchable, respectively) and statistically insignificant (p-values of 0.829 and 0.871). The results are consistent in magnitude too. For example, the first post-treatment coefficient from column (1) suggests a -1.6% ($= \frac{-\$724.51}{\$44{,}146.85}$) increase in average salary relative to the baseline, while the corresponding coefficient from column (2) suggests an increase of 0.6% ($= \frac{\$271.15}{\$44{,}146.85}$).

The specification from column (3) of Table E.1 is identical to the baseline specification from column (1), except that the dependent variable is winsorized at 10% and 90% of the benchmark instead of the 2.5 and 97.5 percentile by position title. Column (4) is identical to column (1), except that the standard errors are not clustered. Column (5) is identical to column (1), except that it does not include any of the additional control variables. Column (6) excludes position fixed effects. Column (7) includes firm fixed effects instead of position fixed effects. Column (8) is identical to column (1), except that it excludes positions for which the base salary is not a major component of compensation: Waiter/Waitress, Chauffeur, and Bartender/Mixologist. Column (9) is identical to column (1), except that it restricts to the 329 positions that are searched at least once in the sample. Column (10) is identical to column (1), except that it does not re-weight by SOC groups. Last, column (11) is identical to column (1), except that it only includes new hires aged 21 through 60. In all these alternative specifications, the results are both qualitatively and quantitatively similar to those from column (1). Panel B of Table E.1 presents the corresponding “pre-treatment” coefficients. As expected, with few exceptions, these coefficients are close to zero, statistically insignificant and precisely estimated.

E.2 Additional Robustness Checks

In Table D.4, we show that the effects on salary dispersion are robust to a wide range of alternative specifications. In this Appendix, we show that the effects on salary levels are also robust to this same range of alternative specifications. The results are presented in Table E.2. Columns (1)–(5) measure the effect on the log of average base salary, with column (1) identical to column (1) of Table E.1. Columns (2)–(5) each change one feature of the baseline specification. In column (2), we include *Match Outliers* (positions with low match scores, indicating low match quality between the firm-specific job title and the position title) that are excluded from the main analysis. In column (3), we restrict the sample to only include Non-Searched and Non-Searchable positions that are after September 2019, the start of our search data. In column (4), we include new hires from August 2020 through July 2021. In column (5), we exclude HR positions, as HR professionals are the most common users of the benchmarking tool.

One potential concern is that firms may be reacting by changing beyond base salary, through bonuses, commissions or even hours worked. In addition to the base salary, our employee data includes the monthly gross wage: this is how much money the firm effectively pays to the employee each month, which reflects not only the base salary but also a myriad of other factors such as hours worked, tax withholdings, commission, bonuses and reimbursements. The last columns of Table E.2 measures the effects on average gross pay (instead of average base salary). In the first specification of column (6), we define the annual gross wage as the average monthly gross pay during the first three months working at the firm, then multiplied by 12 to transform it to an annual basis (i.e., so that it is comparable to the base salary outcome).[68] The base salary and the gross compensation are highly correlated, but not perfectly so (correlation coefficient of 0.848, p-value<0.001). One minor shortcoming with the gross pay data is that it is missing for 11.2% of the observations, for a variety of reasons.[69] In any case, as shown in column (9) of Table E.2, the results are also similar if we impute these missing values.

The results using the (log) gross pay outcome are presented in columns (6)–(9) of Table E.2. First of all, notice that the coefficients for gross pay are much less precisely estimated than the corresponding coefficients for base salary. For example, the standard errors of the post-treatment coefficients for gross pay (0.042 and 0.043, from column (6)) are 2 times as large as the corresponding coefficients for base salary (0.017 and 0.016, from column (1)).

[68] We compute the average starting on the 1st day of the month following the hire date, to make it more comparable across different employees. For employees who work fewer than three months at the firm, the average will be based on the one or two months they worked at the firm.

[69] For example, payroll data are not available for 7 firms, and for other firms it is missing for some employees for a variety of reasons such as failure of data entry from the manager.

This should be expected: relative to the base salary outcome, the gross pay outcome is more volatile because it includes a myriad of factors such as differences in tax withholdings, commissions and so on. The point coefficients for the gross pay outcome (0.002 and 0.037, from column (6)) are similar in magnitude to the corresponding coefficients for base salary (0.003 and 0.019, from column (1)), and statistically indistinguishable from each other.

E.3 Heterogeneity by Skill

Figure E.1 displays the event-study analysis of salary levels, reproducing Figure 5, but for the high-skill sub-sample. When compared to the Non-Searched control group, the Searched group have an average salary level that is 0.029 lower on average (p-value = 0.119) in the post period. Compared to the Non-Searchable control group, the post-period salary level is 0.016 log points lower (p-value = 0.288). These differences are economically modest and statistically insignificant, suggesting that the salary levels before and after gaining access to the salary benchmarking tool remain steady among the high-skill sub-sample. The difference-in-differences estimates under all the different specifications are presented in Table E.3 and Table E.4, respectively for low-skill and high-skill sub-samples. The results are largely robust across specifications.

E.4 Heterogeneity by Gender

The analysis of heterogeneity by the gender of the employee is presented in Figure E.2. Panel A and Panel B reproduce Panel C and Panel D of Figure 5 for the male subsample, and Panel C and Panel D use female subsample only. We find gender differences that are small in magnitude and statistically insignificant.

E.5 Heterogeneity by Market Dispersion

Figure E.3 displays the event-study analysis of salary levels split by above- and below-median market dispersion. Market dispersion is defined by the dispersion in salaries across all employees within the same position, state and industry as the new hire in our sample. Low dispersion markets are likely markets with more standardized or homogeneous labor, and therefore more compressed pay. Comparing Panels A and B (low-dispersion markets), with Panels C and D (high-dispersion markets) reveals that salary-levels rise modestly after accessing the salary benchmark tool primarily within low-dispersion markets. Among low-dispersion markets, the average salary level is 0.035 higher on average (p-value = 0.254) in the post period when comparing Searched vs. Non-Searched groups, and 0.032 higher (p-value = 0.056) when comparing Searched vs. Non-Searchable. Among high-dispersion markets, the respective change in average salary levels are -0.030 (p-value = 0.159) and -0.004 (p-value = 0.922), suggesting that the salary levels do not rise in high-dispersion markets.

E.6 Effects on Retention Levels

In Figure E.4, we present the event-study analysis for the retention outcome in the entire sample. Specifically, the dependent variable is whether an employee hired in a given month is still employed at the same firm one year later. The effects of benchmarking on retention are not statistically significant for the full sample. Panel C of Figure E.4 presents the difference between the Searched and Non-Searched groups. This analysis suggests that the retention rate rose from 52.0 pp to 56.1 pp after onboarding (p-value=0.163). Panel D presents the same difference for Searched and Non-Searchable positions. After onboarding, the retention rate rose from 52.0 pp to 53.4 pp (p-value=0.619). The corresponding difference-in differences estimates are presented in Table E.5, and the results are robust across specifications.

Figure E.5 presents the results for the high-skill sample (for reference, the results for the low-skill sample are presented in Figure 7 above). Consistent with the fact that we do not find significant effects on the average salaries in high-skill positions, we do not find any significant effects on the average retention rate either. More precisely, Panel C of Figure E.5 shows that, when using Non-Searched as control group, the average retention level increases by 3.8 pp (p-value = 0.304) after onboarding. Panel D shows that, when using Non-Searchable as the control group, the average retention drops by 0.2 pp (p-value = 0.976) after onboarding. These differences are economically small and statistically insignificant.[70]

The heterogeneity analysis by market dispersion is presented in Figure E.6. Commensurate with our findings on salary levels, retention rises in low-dispersion markets after onboarding, and remains largely unchanged in high-dispersion markets. Panels A and B present the low-dispersion markets. Compared to the Non-Searched control group, the Searched group experiences a 7.5 pp increase (p-value = 0.042) in the likelihood of being at the same employer a year after being hired, off a base of 41.8%, and a 6.9 pp increase (p-value = 0.019) when compared to the Non-Searchable control group. By contrast, Panels C and D present the high-dispersion markets. The estimated retention effect remains within 2 pp and statistically insignificant when restricting our sample to the high-dispersion markets.

[70]The corresponding difference-in-differences estimates are presented in Table E.6 (for low-skill positions) and Table E.7 (for high-skill positions).

Figure E.1: The Effects of Salary Benchmarking on Salary Levels: High-Skill Subsample

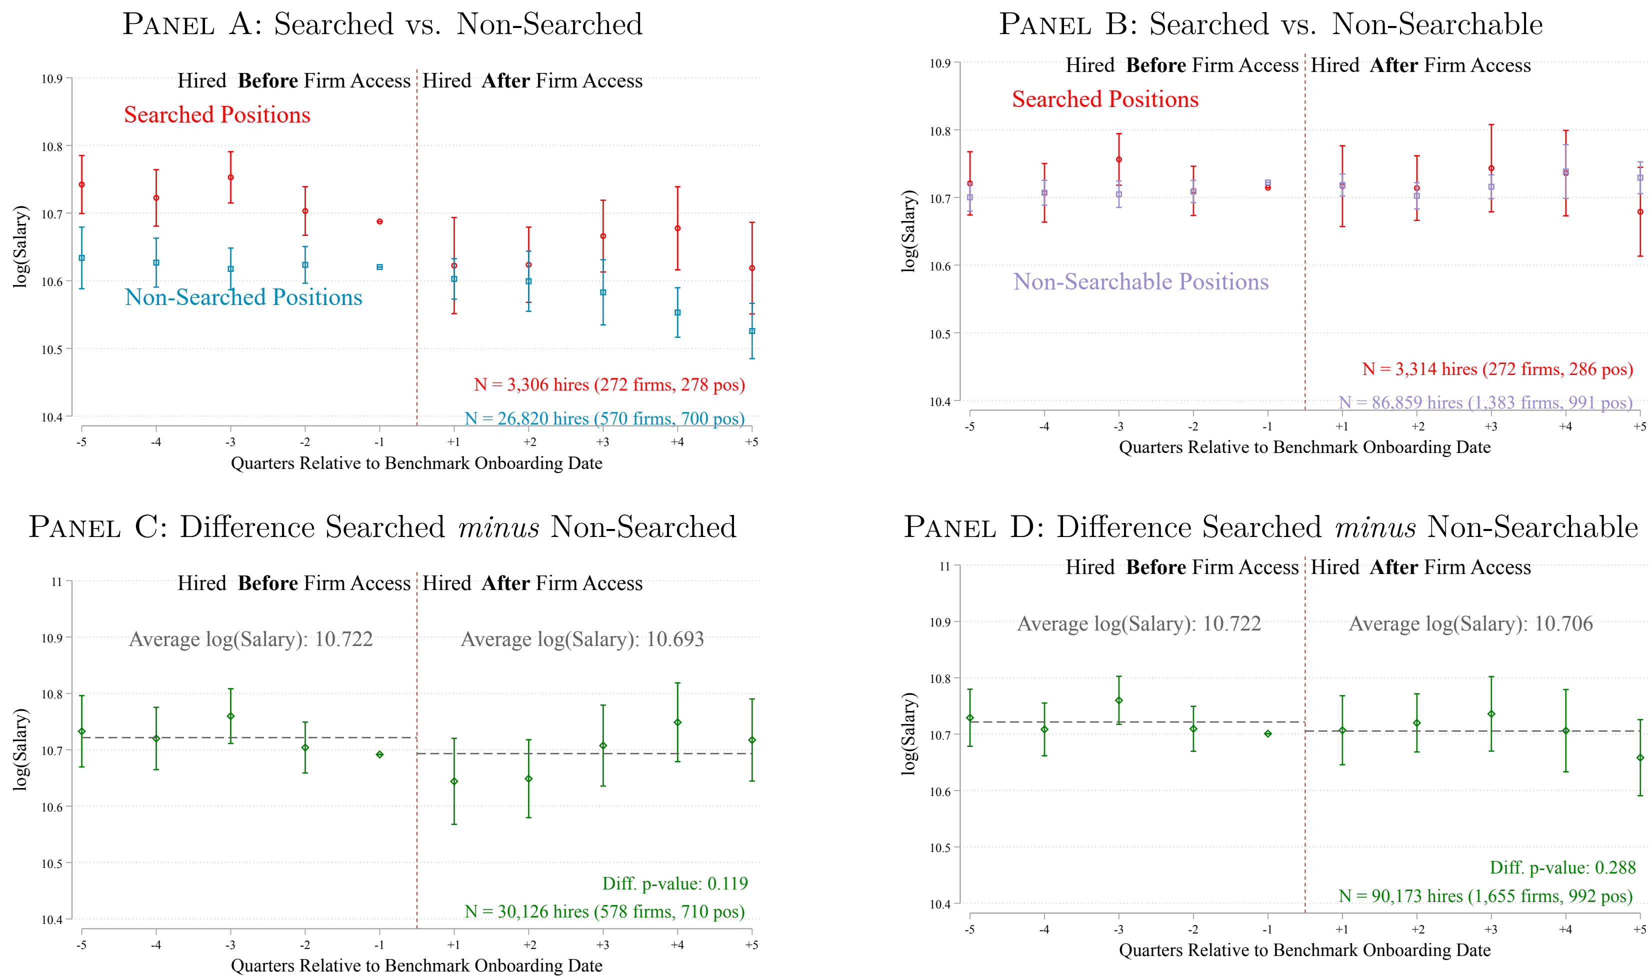

Notes:

Figure E.2: Heterogeneity by Gender: Event-Study Analysis

Panel A: Male: Searched vs. Non-Searched

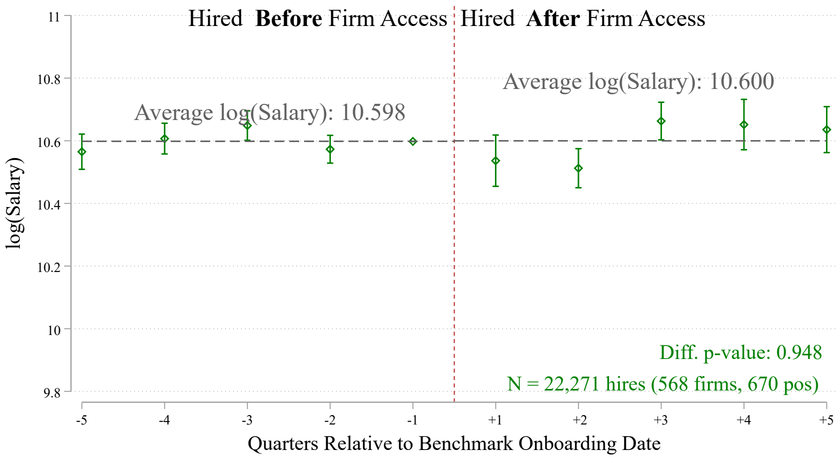

Panel B: Male: Searched vs. Non-Searchable

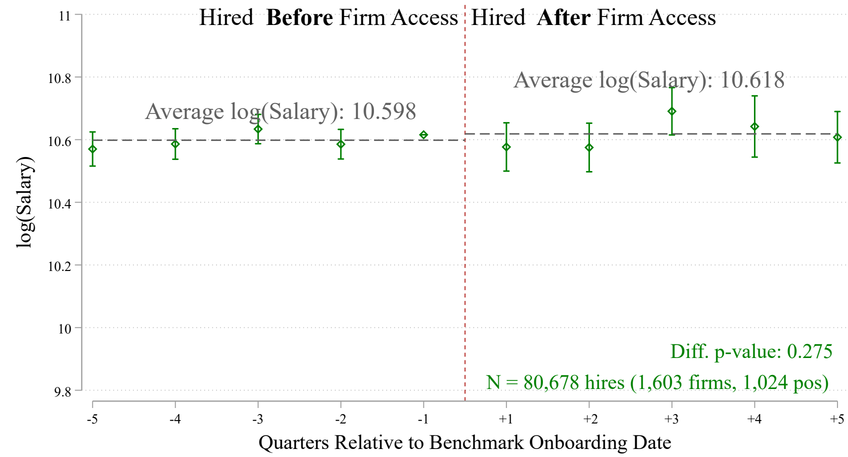

Panel C: Female: Searched vs. Non-Searched

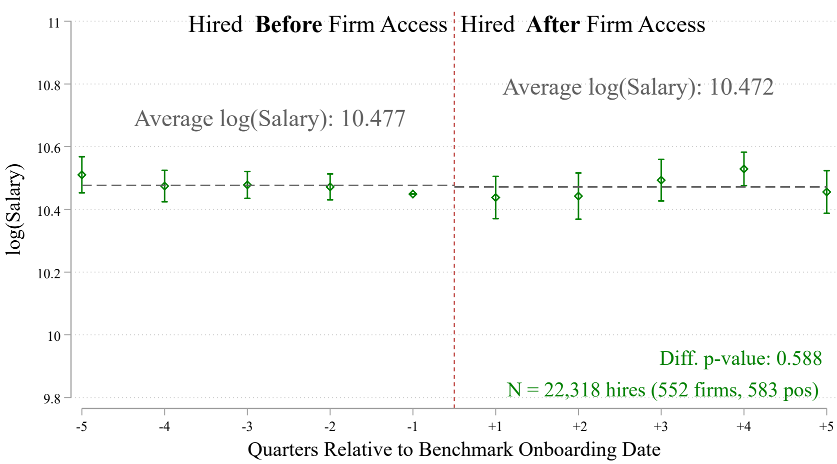

Panel D: Female: Searched vs. Non-Searchable

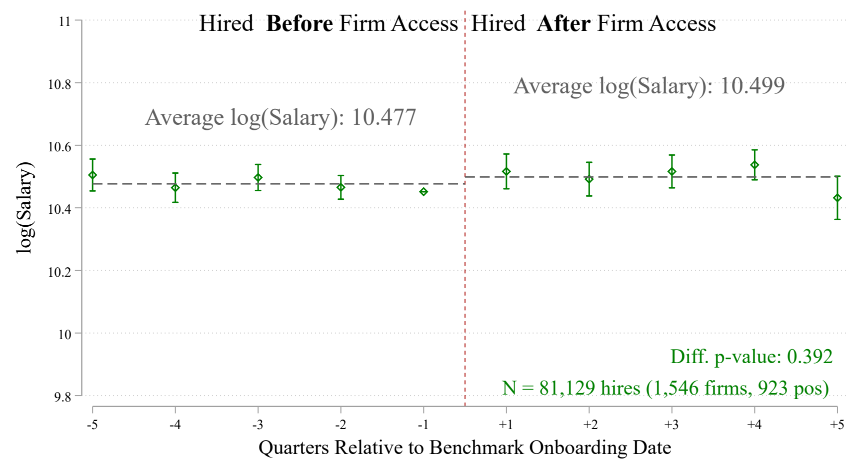

Notes: This is a reproduction of Figure 5 Panel C and D, split by the gender of employees. See the notes of Figure 5 for more details.

Figure E.3: Heterogeneity by Market Dispersion: The Effects on Salary Levels

Panel A: Low Market Dispersion: Searched vs. Non-Searched

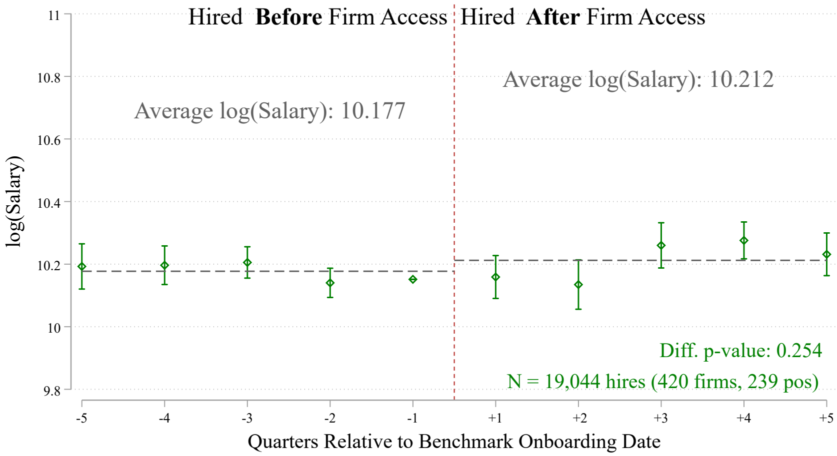

Panel B: Low Market Dispersion: Searched vs. Non-Searchable

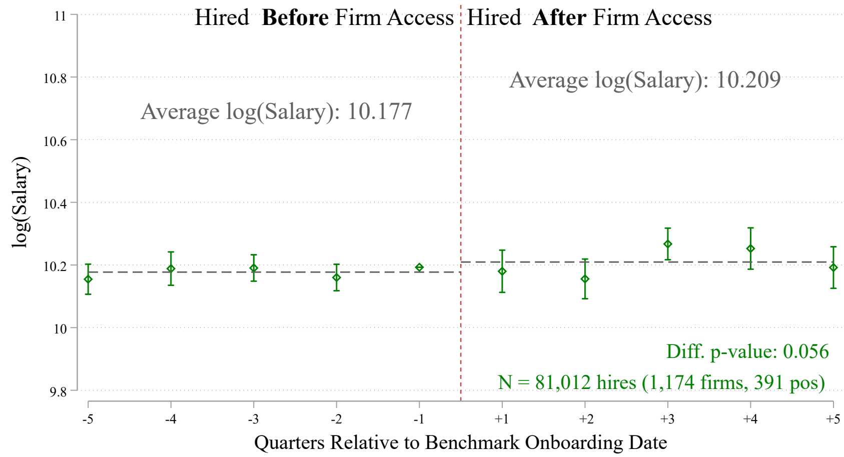

Panel C: High Market Dispersion: Searched vs. Non-Searched

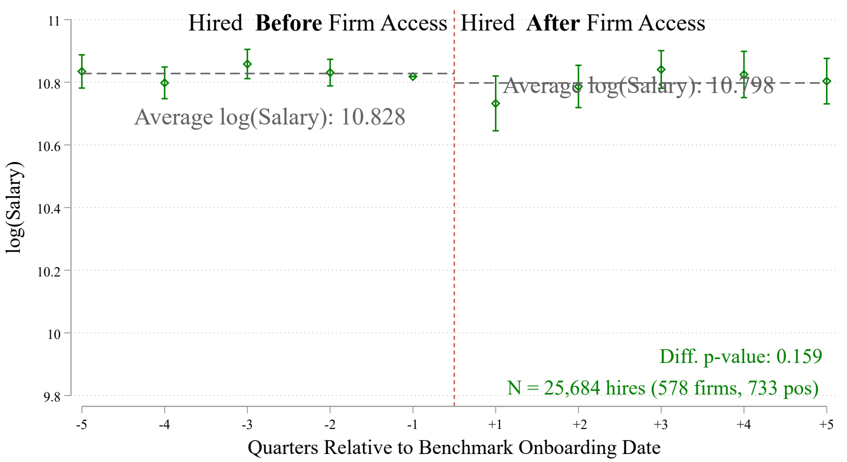

Panel D: High Market Dispersion: Searched vs. Non-Searchable

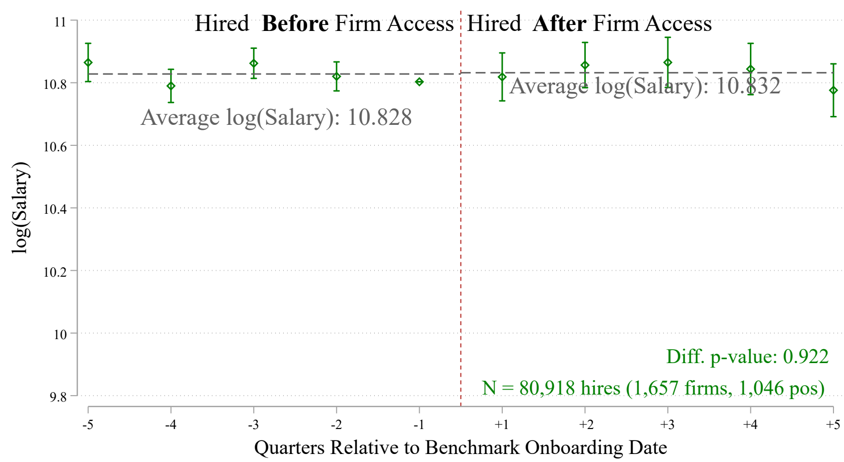

Notes: This is a reproduction of Figure 5 Panel C and D, split by the market dispersion. See the notes of Figure 5 for more details.

Figure E.4: Retention: Event-Study Analysis

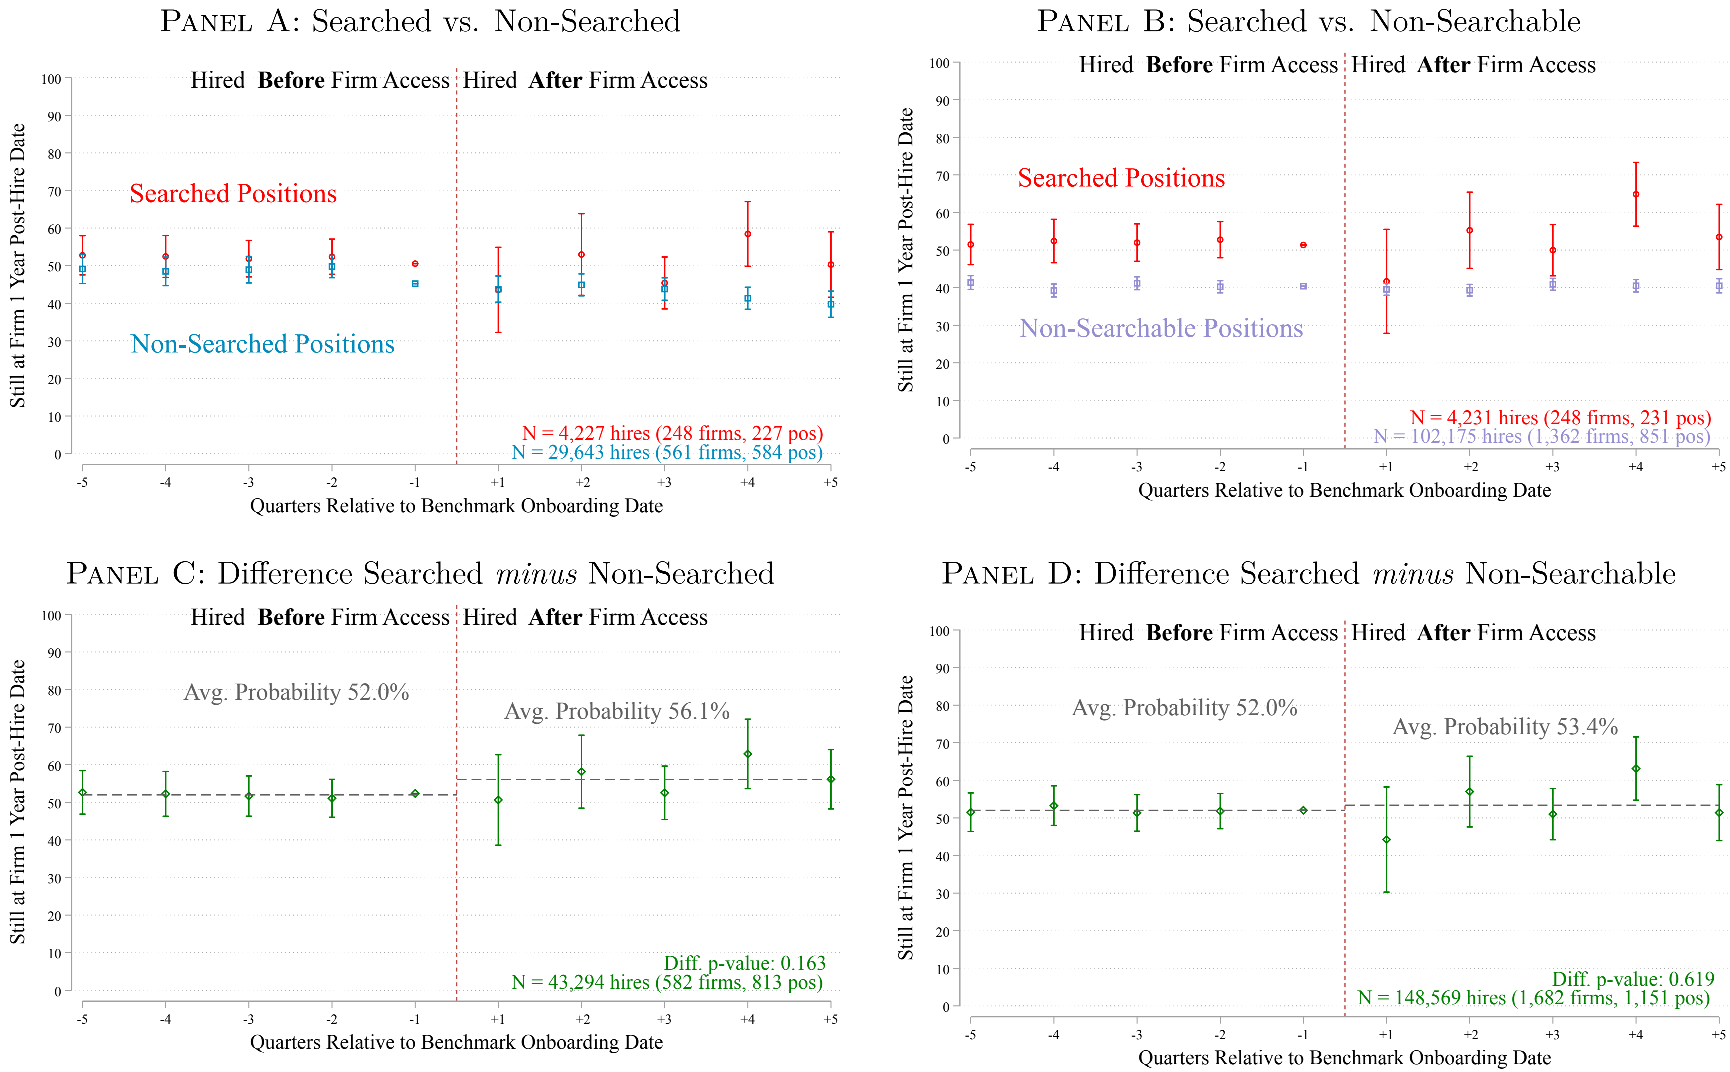

Notes: This is a reproduction of Figure 5, but with the outcome being a dummy equal to 100 if a new hire in a given month is still at the same firm 1 year later. Because our main sample ends in March 2020 and our data ends in July 2021, we observe this outcome for all new hires in our main sample. For more details, see notes to Figure 5.

Figure E.5: The Effects of Salary Benchmarking on Retention Rates: High-Skill Subsample

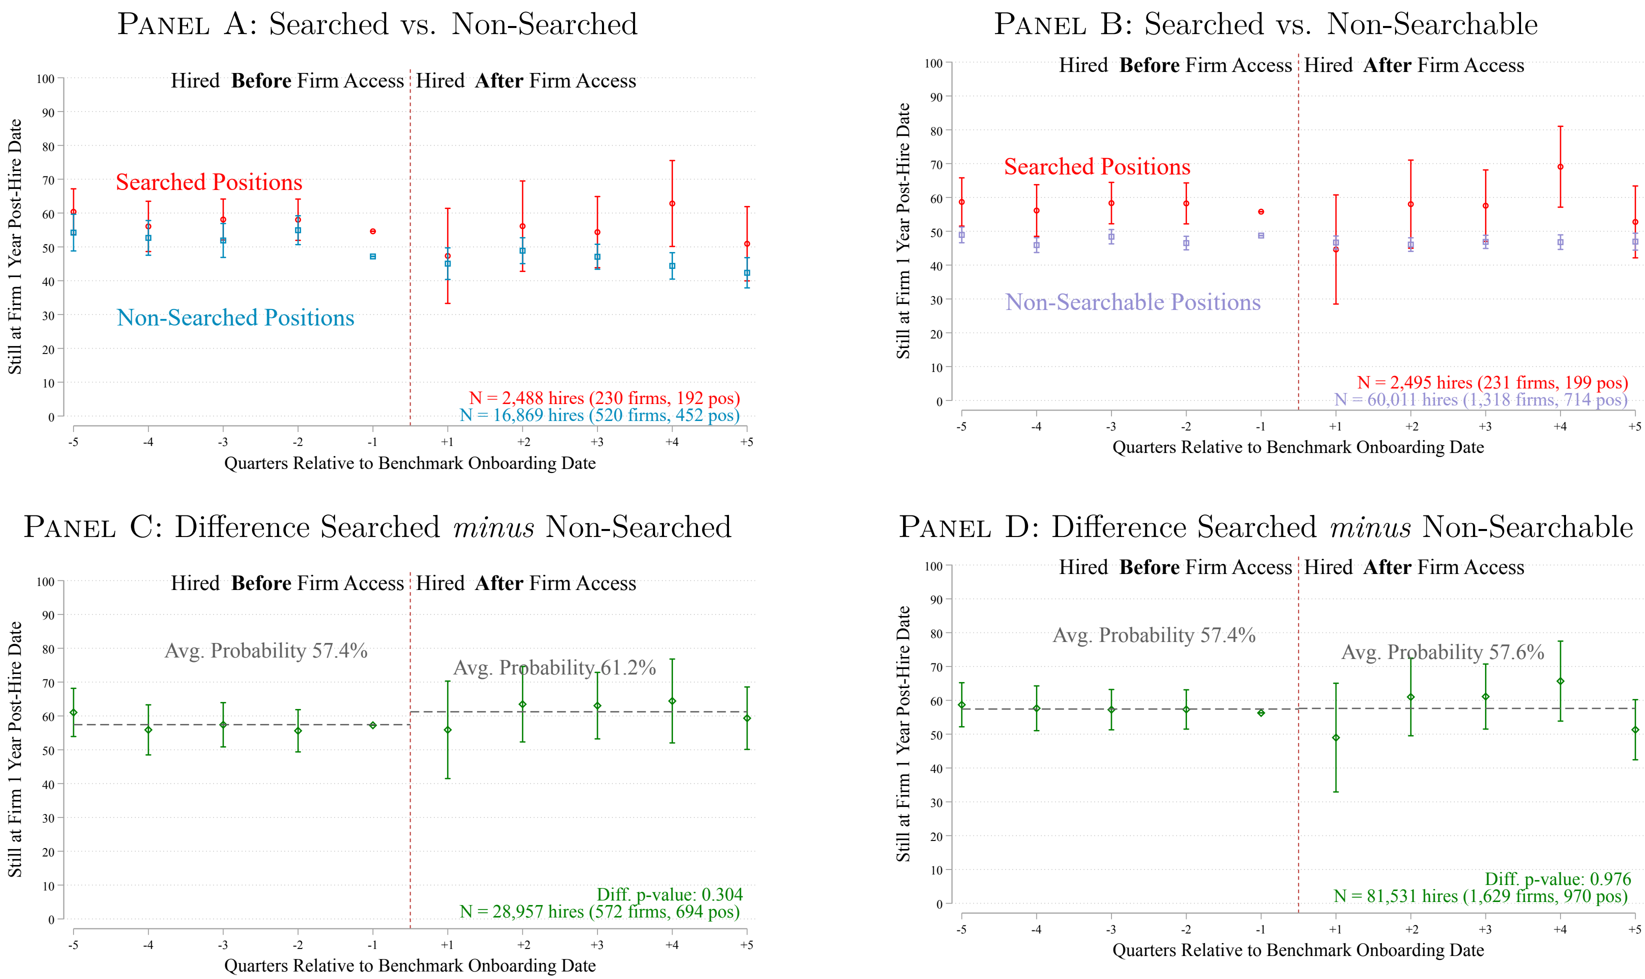

Notes:

Figure E.6: Heterogeneity by Market Dispersion: The Effects on Retention Rates

Panel A: Low Market Dispersion: Searched vs. Non-Searched

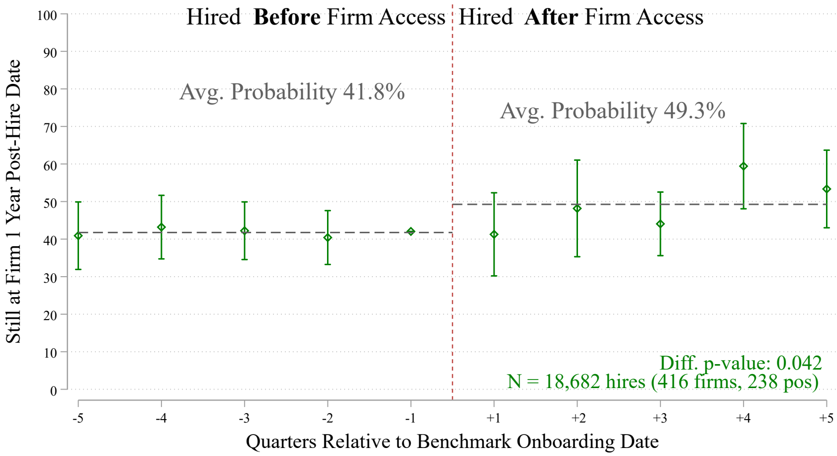

Panel B: Low Market Dispersion: Searched vs. Non-Searchable

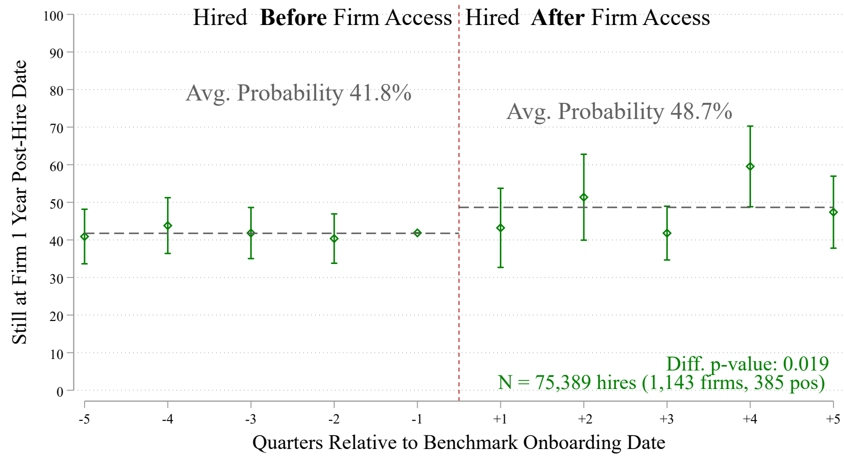

Panel C: High Market Dispersion: Searched vs. Non-Searched

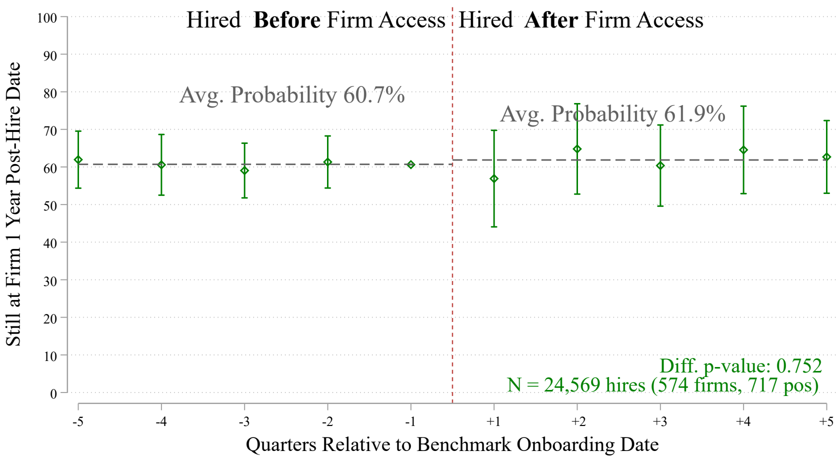

Panel D: High Market Dispersion: Searched vs. Non-Searchable

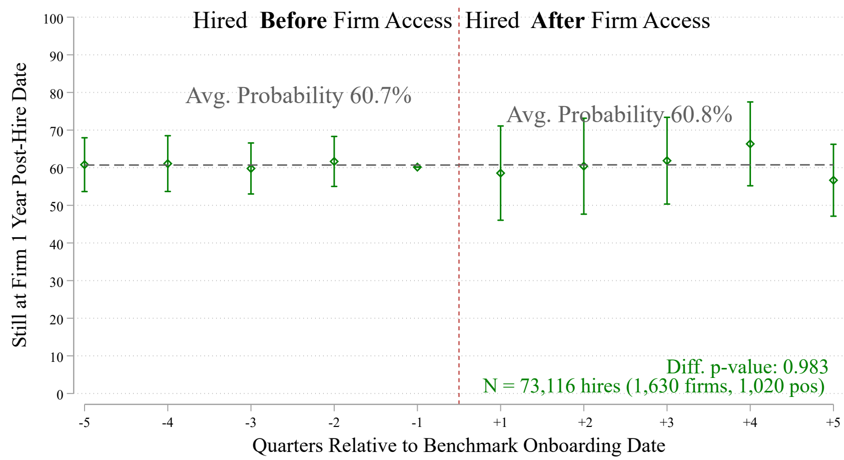

Notes: This is a reproduction of Figure E.4 Panel C and D, split by the market dispersion. See the notes of Figure E.4 for more details.

Table E.1: The Effects of Benchmarking on Salary Levels: Full Sample

	(1) log(*Salary*)	(2) Salary	(3) log(*Salary*)	(4) log(*Salary*)	(5) log(*Salary*)	(6) log(*Salary*)	(7) log(*Salary*)	(8) log(*Salary*)	(9) log(*Salary*)	(10) log(*Salary*)	(11) log(*Salary*)
Panel A: Post-treatment											
Searched vs. Non-Searched	0.003	-724.514	0.004	0.003	0.001	-0.001	-0.000	-0.000	0.005	0.017	-0.001
	(0.017)	(861.430)	(0.017)	(0.012)	(0.019)	(0.032)	(0.019)	(0.018)	(0.018)	(0.017)	(0.018)
Searched vs. Non-Searchable	0.019	271.150	0.021	0.019*	0.009	0.015	0.021	0.015	0.017	0.030*	0.012
	(0.016)	(926.594)	(0.016)	(0.011)	(0.018)	(0.025)	(0.023)	(0.016)	(0.016)	(0.018)	(0.017)
Panel B: Pre-treatment											
Searched vs. Non-Searched	-0.022	-1166.949	-0.021	-0.022**	-0.028	-0.030	0.006	-0.019	-0.019	-0.013	-0.022
	(0.018)	(870.587)	(0.018)	(0.011)	(0.019)	(0.038)	(0.016)	(0.018)	(0.018)	(0.019)	(0.017)
Searched vs. Non-Searchable	-0.005	-686.358	-0.005	-0.005	-0.017	-0.014	0.000	-0.011	-0.004	-0.004	-0.002
	(0.017)	(806.078)	(0.016)	(0.010)	(0.019)	(0.027)	(0.013)	(0.016)	(0.017)	(0.018)	(0.017)
Alternate Winsorization			✓								
No Clustering				✓							
No Additional Controls					✓						
No Position FE						✓					
Firm FE							✓				
Exclude High-Tip Jobs								✓			
Searched Positions Only									✓		
No Re-weighting										✓	
Ages 21-60											✓
Mean Dep. Var. (Baseline)	10.532	44146.850	10.523	10.532	10.532	10.532	10.532	10.547	10.532	10.506	10.565
Observations											
Searched	5,253	5,253	5,253	5,253	5,253	5,266	5,262	5,105	5,253	5,316	4,611
Non-Searched	39,527	39,527	39,527	39,527	39,527	39,686	39,673	37,841	34,954	39,645	34,338
Non-Searchable	156,734	156,734	156,734	156,734	156,734	156,865	156,817	148,521	127,145	156,883	135,051

Notes: Significant at *10%, **5%, ***1%. Standard errors clustered at the firm-position-month level in parentheses. Each column corresponds to two regressions: one for Searched vs. Non-Searched new hires and one for Searched vs. Non-Searchable new hires. Post-treatment coefficients in Panel A refer to parameters α_1^k from equation (2), while pre-treatment coefficients in Panel B refer to parameters α_3^k from equation (2) (see Section 3.2 for details). All columns include year fixed effects. In columns (1) and (3)–(11) the dependent variable is the log of annual base salary. The dependent variable in column (2) is the annual base salary (in $s). Log salary and salary are winsorized to the 2.5 and 97.5 percentiles of all salaries for their position. The exception is column (3) where wages are winsorized to ± 90% of the median benchmark. All columns except (5) include additional controls (female dummy, high education dummy, hourly dummy, age, position tenure). Column (6) excludes position fixed effects. Column (7) includes firm fixed effects instead of position fixed effects. Column (8) excludes the three positions where gross pay most exceeds base pay: Waiter/Waitress, Chauffeur, and Bartender/Mixologist. Column (9) restricts the sample to only titles of Non-Searched or Non-Searchable new hires in positions that are searched and hired by firms in the data.

Table E.2: The Effects of Benchmarking on Salary Levels: Additional Robustness Checks

	(1) log($Salary$)	(2) log($Salary$)	(3) log($Salary$)	(4) log($Salary$)	(5) log($Salary$)	(6) log($Gross$)	(7) log($Gross$)	(8) log($Gross$)	(9) log($Gross$)
Panel A: Post-treatment									
Searched vs. Non-Searched	0.003	0.003	-0.007	0.004	0.002	0.002	-0.001	-0.005	-0.004
	(0.017)	(0.016)	(0.019)	(0.017)	(0.018)	(0.042)	(0.043)	(0.041)	(0.040)
Searched vs. Non-Searchable	0.019	0.016	0.000	0.018	0.018	0.037	0.031	0.026	0.033
	(0.016)	(0.016)	(0.017)	(0.015)	(0.017)	(0.043)	(0.043)	(0.042)	(0.040)
Panel B: Pre-treatment									
Searched vs. Non-Searched	-0.022	-0.037**	-0.016	-0.022	-0.018	0.017	0.026	0.024	0.020
	(0.018)	(0.017)	(0.017)	(0.018)	(0.018)	(0.035)	(0.035)	(0.035)	(0.033)
Searched vs. Non-Searchable	-0.005	-0.009	-0.011	-0.005	-0.001	0.011	0.008	0.019	0.020
	(0.017)	(0.017)	(0.016)	(0.017)	(0.017)	(0.031)	(0.031)	(0.031)	(0.029)
Include Match Outliers		✓							
Restricted Sample			✓						
After Aug-2020				✓					
Exclude HR Positions					✓				
3 Month Window						✓			
2 Month Window							✓		
6 Month Window								✓	
Imputed									✓
Mean Dep. Var. (Baseline)	10.532	10.527	10.532	10.532	10.507	10.382	10.381	10.394	10.378
Observations									
Searched	5,253	6,150	5,246	5,414	5,080	4,869	4,864	4,875	5,253
Non-Searched	39,527	50,943	15,958	46,435	38,156	35,884	35,844	35,915	39,527
Non-Searchable	156,734	196,768	83,348	160,595	151,778	138,178	138,002	138,409	156,734

Notes: Columns (1)–(5) look at effects on the log of annual base salary. Column (1) is exactly column (1) from Table E.1. Column (2), *Include Match Outliers*, reproduces column (1), but including new hires whose organization specific job title has a low match score to the designated position title (scores less than the 20th percentile of the scores in that quarter). Column (3) using the *Restricted Sample* uses only control observations after September 2019, the start of our search data. Column (4) includes data from August 2020 through July 2021. Column (5) excludes new hires in HR positions. Columns (6)–(9) look at effects of the log of annual gross wages (as described in Section E.2). Column (6) uses a 3 month window to compute gross wages, while columns (7) and (8) use a 2 and 6 month window, respectively. Column (9) is equivalent to column (6) with missing values imputed. See Table E.1 for more details.

Table E.3: The Effects of Benchmarking on Salary Levels: Low-Skill Subsample

	(1) log(Salary)	(2) Salary	(3) log(Salary)	(4) log(Salary)	(5) log(Salary)	(6) log(Salary)	(7) log(Salary)	(8) log(Salary)	(9) log(Salary)	(10) log(Salary)	(11) log(Salary)
Panel A: Post-treatment											
Searched vs. Non-Searched	0.073***	1618.196***	0.058***	0.073***	0.070***	0.088***	0.034**	0.071***	0.078***	0.080***	0.066***
	(0.020)	(541.012)	(0.020)	(0.013)	(0.020)	(0.021)	(0.016)	(0.019)	(0.020)	(0.025)	(0.020)
Searched vs. Non-Searchable	0.087***	1862.423***	0.076***	0.087***	0.079***	0.090***	0.038**	0.087***	0.095***	0.089***	0.081***
	(0.021)	(509.221)	(0.021)	(0.013)	(0.021)	(0.021)	(0.015)	(0.019)	(0.022)	(0.025)	(0.022)
Panel B: Pre-treatment											
Searched vs. Non-Searched	0.002	17.476	-0.008	0.002	0.002	0.015	0.016	0.016	0.005	0.005	-0.001
	(0.023)	(593.343)	(0.025)	(0.012)	(0.025)	(0.025)	(0.012)	(0.024)	(0.024)	(0.024)	(0.022)
Searched vs. Non-Searchable	0.035	335.014	0.030	0.035***	0.030	0.037	0.020**	0.018	0.038	0.037	0.046
	(0.030)	(700.088)	(0.029)	(0.014)	(0.031)	(0.028)	(0.008)	(0.028)	(0.031)	(0.031)	(0.030)
Alternate Winsorization			✓								
No Clustering				✓							
No Additional Controls					✓						
No Position FE						✓					
Firm FE							✓				
Exclude High-Tip Jobs								✓			
Searched Positions Only									✓		
No Re-weighting										✓	
Ages 21-60											✓
Mean Dep. Var. (Baseline)	10.147	25929.204	10.147	10.147	10.147	10.147	10.147	10.150	10.147	10.145	10.168
Observations											
Searched	1,947	1,947	1,947	1,947	1,947	1,947	1,938	1,799	1,947	1,962	1,545
Non-Searched	12,707	12,707	12,707	12,707	12,707	12,724	12,671	11,021	11,015	12,715	10,211
Non-Searchable	69,875	69,875	69,875	69,875	69,875	69,890	69,755	61,662	55,625	69,903	56,380

Notes: Significant at *10%, **5%, ***1%. Standard errors clustered at the firm-position-month level in parentheses. Each column corresponds to two regressions: one for Searched vs. Non-Searched new hires and one for Searched vs. Non-Searchable new hires. Post-treatment coefficients in Panel A refer to parameters α_1^k from equation (2), while pre-treatment coefficients in Panel B refer to parameters α_3^k from equation (2) (see Section 3.2 for details). All columns include year fixed effects. In columns (1) and (3)–(11) the dependent variable is the log of annual base salary. The dependent variable in column (2) is the annual base salary (in \$s). Log salary and salary are winsorized to the 2.5 and 97.5 percentiles of all salaries for their position. The exception is column (3) where wages are winsorized to $\pm$ 90% of the median benchmark. All columns except (5) include additional controls (female dummy, high education dummy, hourly dummy, age, position tenure). Column (6) excludes position fixed effects. Column (7) includes firm fixed effects instead of position fixed effects. Column (8) excludes the three positions where gross pay most exceeds base pay: Waiter/Waitress, Chauffeur, and Bartender/Mixologist. Column (9) restricts the sample to only titles of Non-Searched or Non-Searchable new hires in positions that are searched and hired by firms in the data.

Table E.4: The Effects of Benchmarking on Salary Levels: High-Skill Subsample

	(1) log(Salary)	(2) Salary	(3) log(Salary)	(4) log(Salary)	(5) log(Salary)	(6) log(Salary)	(7) log(Salary)	(8) log(Salary)	(9) log(Salary)	(10) log(Salary)	(11) log(Salary)
Panel A: Post-treatment											
Searched vs. Non-Searched	-0.030	-1739.161	-0.021	-0.030*	-0.033	-0.045	-0.006	-0.030	-0.029	-0.025	-0.027
	(0.023)	(1236.904)	(0.022)	(0.017)	(0.026)	(0.048)	(0.026)	(0.023)	(0.024)	(0.021)	(0.023)
Searched vs. Non-Searchable	-0.022	-706.932	-0.011	-0.022	-0.032	-0.022	-0.019	-0.022	-0.028	-0.014	-0.026
	(0.020)	(1336.488)	(0.019)	(0.014)	(0.023)	(0.037)	(0.033)	(0.020)	(0.020)	(0.023)	(0.020)
Panel B: Pre-treatment											
Searched vs. Non-Searched	-0.039	-1935.302	-0.030	-0.039**	-0.048*	-0.072	-0.017	-0.039	-0.036	-0.025	-0.036
	(0.025)	(1312.840)	(0.024)	(0.016)	(0.027)	(0.058)	(0.021)	(0.025)	(0.025)	(0.026)	(0.024)
Searched vs. Non-Searchable	-0.026	-1380.412	-0.024	-0.026**	-0.042*	-0.042	-0.017	-0.026	-0.025	-0.019	-0.023
	(0.020)	(1174.559)	(0.020)	(0.013)	(0.024)	(0.044)	(0.019)	(0.020)	(0.020)	(0.023)	(0.020)
Alternate Winsorization			✓								
No Clustering				✓							
No Additional Controls					✓						
No Position FE						✓					
Firm FE							✓				
Exclude High-Tip Jobs								✓			
Searched Positions Only									✓		
No Re-weighting										✓	
Ages 21-60											✓
Mean Dep. Var. (Baseline)	10.722	53147.502	10.709	10.722	10.722	10.722	10.722	10.722	10.722	10.711	10.730
Observations											
Searched	3,306	3,306	3,306	3,306	3,306	3,319	3,316	3,306	3,306	3,354	3,066
Non-Searched	26,820	26,820	26,820	26,820	26,820	26,962	26,947	26,820	23,939	26,930	24,127
Non-Searchable	86,859	86,859	86,859	86,859	86,859	86,975	86,904	86,859	71,520	86,980	78,671

Notes: Significant at *10%, **5%, ***1%. Standard errors clustered at the firm-position-month level in parentheses. Each column corresponds to two regressions: one for Searched vs. Non-Searched new hires and one for Searched vs. Non-Searchable new hires. Post-treatment coefficients in Panel A refer to parameters α_1^k from equation (2), while pre-treatment coefficients in Panel B refer to parameters α_3^k from equation (2) (see Section 3.2 for details). All columns include year fixed effects. In columns (1) and (3)–(11) the dependent variable is the log of annual base salary. The dependent variable in column (2) is the annual base salary (in $s). Log salary and salary are winsorized to the 2.5 and 97.5 percentiles of all salaries for their position. The exception is column (3) where wages are winsorized to ± 90% of the median benchmark. All columns except (5) include additional controls (female dummy, high education dummy, hourly dummy, age, position tenure). Column (6) excludes position fixed effects. Column (7) includes firm fixed effects instead of position fixed effects. Column (8) excludes the three positions where gross pay most exceeds base pay: Waiter/Waitress, Chauffeur, and Bartender/Mixologist. Column (9) restricts the sample to only titles of Non-Searched or Non-Searchable new hires in positions that are searched and hired by firms in the data.

Table E.5: The Effects of Benchmarking on Retention: Full Sample

	Dep. Var.: Still at Firm 12 Months Later (=100)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	% retention	% retention	% retention	% retention	% retention	% retention	% retention	% retention	% retention
Panel A: Post-treatment									
Searched vs. Non-Searched	5.541*	5.541*	5.291*	4.401	0.001	4.849*	5.836**	1.072	6.061**
	(2.861)	(2.953)	(2.865)	(3.116)	(2.681)	(2.911)	(2.895)	(2.279)	(3.069)
Searched vs. Non-Searchable	2.447	2.447	1.792	1.945	1.203	1.853	2.397	2.086	3.212
	(2.649)	(2.760)	(2.628)	(3.066)	(2.648)	(2.684)	(2.660)	(2.154)	(2.684)
Panel B: Pre-treatment									
Searched vs. Non-Searched	0.446	0.446	0.068	1.762	1.924	0.877	0.396	2.248	0.561
	(2.534)	(2.300)	(2.541)	(2.718)	(2.381)	(2.580)	(2.584)	(2.614)	(2.641)
Searched vs. Non-Searchable	0.069	0.069	-0.512	1.125	1.918	0.147	-0.045	0.414	-0.457
	(2.305)	(2.025)	(2.326)	(2.617)	(2.213)	(2.342)	(2.327)	(2.485)	(2.388)
No Clustering		✓							
No Additional Controls			✓						
Position FE				✓					
Firm FE					✓				
Exclude High-Tip Jobs						✓			
Searched Positions Only							✓		
No Re-weighting								✓	
Ages 21-60									✓
Mean Dep. Var. (Baseline)									
Observations	5,111	5,111	5,111	5,124	5,121	4,974	5,111	5,189	4,478
Searched	38,189	38,189	38,189	38,355	38,339	36,575	33,787	38,472	33,091
Non-Searched	143,453	143,453	143,453	143,583	143,528	135,829	116,075	143,708	123,157

Notes: Significant at *10%, **5%, ***1%. Standard errors clustered at the firm-position-month level in parentheses. Each column corresponds to two regressions: one for Searched vs. Non-Searched new hires and one for Searched vs. Non-Searchable new hires.

Table E.6: The Effects of Benchmarking on Retention: Low-Skill Subsample

	Dep. Var.: Still at Firm 12 Months Later (=100)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	% retention	% retention	% retention	% retention	% retention	% retention	% retention	% retention	% retention
Panel A: Post-treatment									
Searched vs. Non-Searched	6.126	6.126	7.903*	5.123	-0.417	4.627	4.661	4.415	6.253
	(4.528)	(5.179)	(4.111)	(4.700)	(4.048)	(4.753)	(4.718)	(3.621)	(5.380)
Searched vs. Non-Searchable	4.732	4.732	4.988	4.323	5.271	3.622	4.730	3.668	4.443
	(3.822)	(4.989)	(3.743)	(3.914)	(3.509)	(3.975)	(3.852)	(2.966)	(4.671)
Panel B: Pre-treatment									
Searched vs. Non-Searched	1.982	1.982	3.250	0.493	1.221	3.042	1.643	3.295	1.777
	(4.345)	(3.833)	(4.478)	(4.447)	(4.166)	(4.588)	(4.462)	(4.410)	(4.617)
Searched vs. Non-Searchable	1.883	1.883	2.628	1.394	4.424	2.158	1.525	3.289	1.585
	(4.015)	(3.369)	(4.066)	(4.079)	(3.692)	(4.222)	(4.075)	(4.001)	(4.144)
No Clustering		✓							
No Additional Controls			✓						
Position FE				✓					
Firm FE					✓				
Exclude High-Tip Jobs						✓			
Searched Positions Only							✓		
No Re-weighting								✓	
Ages 21-60									✓
Restricted Sample	41.085	41.085	41.085	41.085	41.085	42.687	41.085	36.986	41.397
Mean Dep. Var. (Baseline)									
Observations	1,882	1,882	1,882	1,883	1,874	1,745	1,882	1,898	1,487
Searched	12,449	12,449	12,449	12,450	12,415	10,834	10,778	12,473	9,992
Non-Searched	65,152	65,152	65,152	65,156	65,033	57,526	51,747	65,202	52,397

Notes: Significant at *10%, **5%, ***1%. Standard errors clustered at the firm-position-month level in parentheses. Each column corresponds to two regressions: one for Searched vs. Non-Searched new hires and one for Searched vs. Non-Searchable new hires.

Table E.7: The Effects of Benchmarking on Retention: High-Skill Subsample

	Dep. Var.: Still at Firm 12 Months Later (=100)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	% retention	% retention	% retention	% retention	% retention	% retention	% retention	% retention	% retention
Panel A: Post-treatment									
Searched vs. Non-Searched	4.611	4.611	4.453	3.624	1.822	4.611	5.512	1.323	5.215
	(3.694)	(3.406)	(3.704)	(3.918)	(3.467)	(3.694)	(3.724)	(2.882)	(3.846)
Searched vs. Non-Searchable	0.918	0.918	0.236	1.961	-0.561	0.918	0.778	1.776	2.409
	(3.465)	(3.266)	(3.459)	(3.921)	(3.632)	(3.465)	(3.470)	(2.830)	(3.316)
Panel B: Pre-treatment									
Searched vs. Non-Searched	-0.318	-0.318	-0.779	0.749	0.636	-0.318	-0.055	-0.501	-0.564
	(3.037)	(2.848)	(3.041)	(3.288)	(2.943)	(3.037)	(3.086)	(3.158)	(3.145)
Searched vs. Non-Searchable	-1.549	-1.549	-2.271	-0.041	-0.520	-1.549	-1.685	-2.293	-2.213
	(2.795)	(2.507)	(2.815)	(3.573)	(2.852)	(2.795)	(2.823)	(3.370)	(2.896)
No Clustering		✓							
No Additional Controls			✓						
Position FE				✓					
Firm FE					✓				
Exclude High-Tip Jobs						✓			
Searched Positions Only							✓		
No Re-weighting								✓	
Ages 21-60									✓
Mean Dep. Var. (Baseline)									
Observations	3,228	3,228	3,228	3,241	3,239	3,228	3,228	3,291	2,991
Searched	25,739	25,739	25,739	25,888	25,868	25,739	23,006	25,999	23,099
Non-Searched	78,297	78,297	78,297	78,409	78,331	78,297	64,325	78,506	70,760

Notes: Significant at *10%, **5%, ***1%. Standard errors clustered at the firm-position-month level in parentheses. Each column corresponds to two regressions: one for Searched vs. Non-Searched new hires and one for Searched vs. Non-Searchable new hires.

F Additional Results: Sample of Existing Employees

Due to its simplicity, our main theoretical and empirical analysis focuses on new hires. For the sake of completeness, we provide some additional results for a sample of existing employees. For this sample, however, we must keep in mind some limitations with the data. The main limitation is that we cannot be certain which employee the information pertains to precisely. As an example, suppose the firm looks up the benchmark for bank tellers, of which there are 100 existing employees. The data challenge is that we only observe whether the firm looked up the benchmark for "bank teller", but not which of their 100 bank tellers the data was relevant for. It may be that the benchmark was looked up to adjust the salary ranges for all 100 tellers. Or maybe the firm only needs the information to respond to the outside offer of one particular teller. As a result, if we assume that the information should affect all 100 existing employees, the estimates could suffer from massive attenuation bias. For these reasons, the results shown below must be taken with a grain of salt.[71]

The outcome of interest when thinking of existing employees is not the salary level, but the changes in salaries. Firms review the salaries of their existing employees infrequently and, even when doing so, they are subject to strong downward wage rigidities (Kaur, 2019; Grigsby et al., 2021). These rigidities must be taken into account when interpreting the results: if hiring managers find it harder to lower the wages of existing employers than to raise them, that could mechanically lead to an increase in average salaries.

We constructed an annual panel of existing employees and calculated the percent change in their salary from January to December of each year. Figure F.1 presents the distribution of annual salary changes for all position types in the pre-onboarding period. Typically, employees do not experience a salary revision. When they do experience a revision, it tends to be positive and small, with a notable right tail of significant revisions that might accompany promotion. And, consistent with downward wage rigidities, negative salary revisions are extremely rare, occurring only in 1% of the employee-year observations.

Figure F.2 extends our event-study framework to study the effect of salary benchmarks on how firms change salaries for existing employees. Panels A and B show the estimates for Searched, Non-Searched, and Non-Searchable existing employees, while panels C and D show the difference-in-differences estimates. The post-treatment coefficients are positive and statistically significant: 1.407 (p-value$<$0.001) using Non-Searched positions as control group and 0.911 (p-value$=$ 0.005) using Non-Searchable positions as control. The modest gains in

[71] This source of attenuation bias still exists for the analysis of new hires, but is much less severe. By construction, the firm has to actively set a salary for every one of their new employees. As a result, if a firm looks up the information for a position in which they are hiring, it is likely that they will use that information in setting the salary of that new hire. To the extent that typically one or a few employees are hired at the same time, there is much less scope for misattribution.

salaries among existing employees are consistent with the gains in average salaries observed in the analysis of new hires. However, as discussed above, this result for existing employees must be interpreted in light of nominal wage rigidities: i.e. after looking up the benchmark some employers may desire to cut salaries of some employees but are unable to do so because of nominal wage rigidities.

Figure F.1: Analysis of Existing Employees: Annual Percent Change in Salary

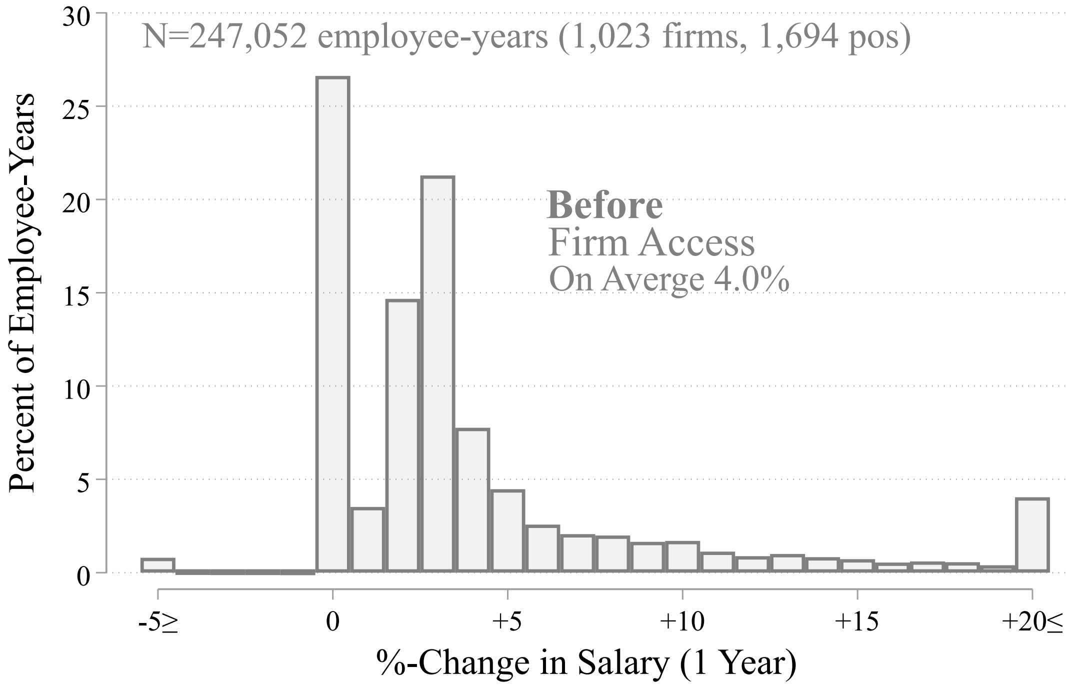

Notes: Distribution of the annual percent change in salary for existing employees of all position types before onboarding. Winsorized at -5 and +20.

Figure F.2: Analysis of Existing Employees: Event-Study Analysis of the Effects on Annual Percent Change in Salary

Panel A: Searched vs. Non-Searched

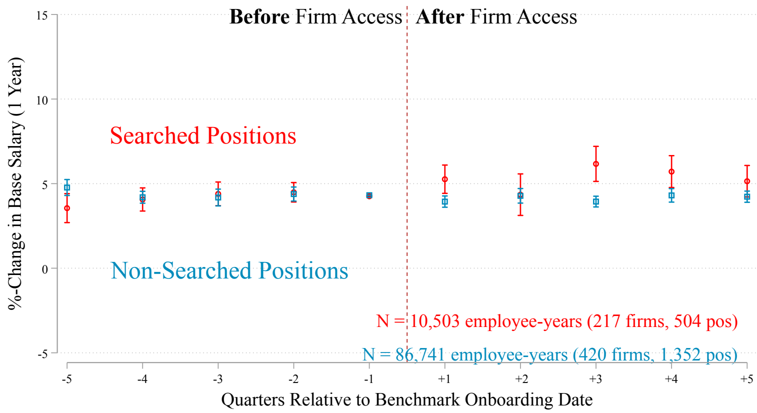

Panel B: Searched vs. Non-Searchable

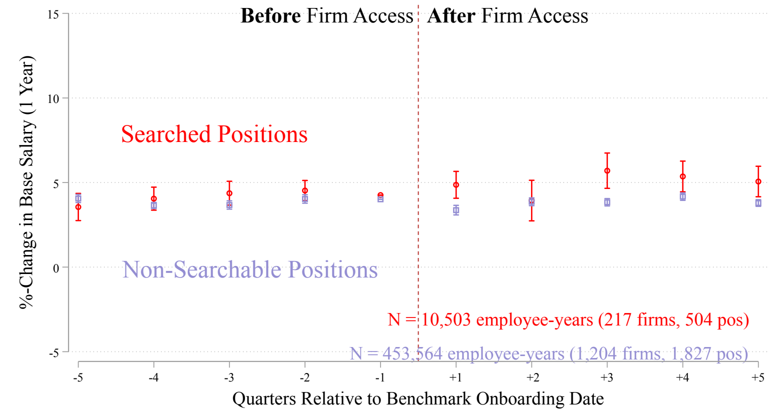

Panel C: Difference Searched *minus* Non-Searched

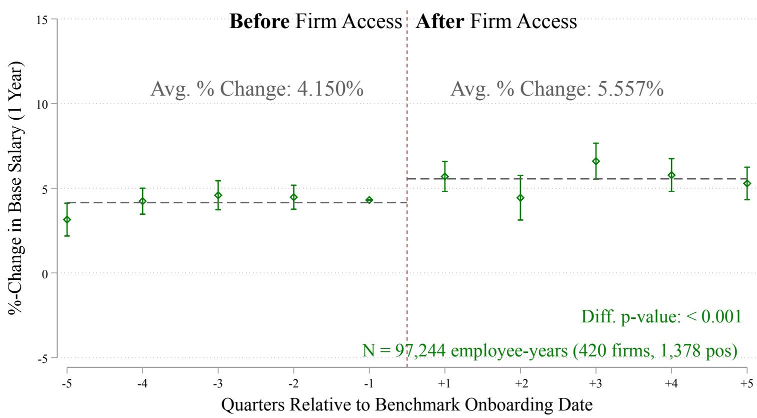

Panel D: Difference Searched *minus* Non-Searchable

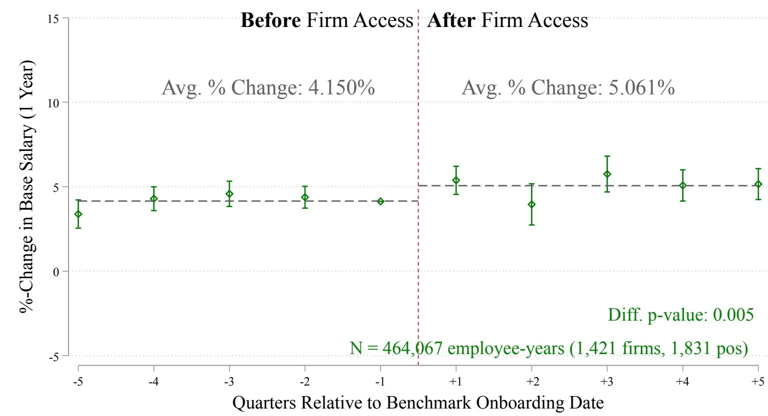

Notes: Each panel is a reproduction of the corresponding panel of Figure 3, except looking at the annual percent change in salary for existing employees and excluding new hires. For more details, see notes to Figure 3.

G Pass-through from Salaries to Benchmark

Recall, in our baseline model, each firm treats the benchmark as exogenous. However, in practice, some employers may be large enough to influence the market benchmarks. Their wages "pass-through" to the benchmark and hence affect the information that other firms use to make offers. Here we explore empirically how common these cases might be. Specifically, we investigate whether some firms hire a sufficiently large share of the employees in a labor market to shift the median of the benchmark when they adjust the salaries of their employees. In Figure G.1, we simulate how much the median of the salary benchmark for a position-industry-state would shift in the following quarter, when the benchmark is recalculated, if a firm decided to raise the salaries of all new hires (Panel A) and all existing employees (Panel B) by 10%. The result is very stark, shifting the salaries of all new hires by 10% would shift the benchmark median by 0.23% on average, and only 3% of firm-positions could shift the median by more than 1%. Even if a firm were to raise the salaries of all its employees by 10%, the median of the benchmark would only shift on average 0.59%. However, there are some firms that are large enough to have a meaningful impact on the market benchmarks.

Figure G.1: Salary Pass-through to Benchmark

Panel A: New Hires

Panel B: All Employees

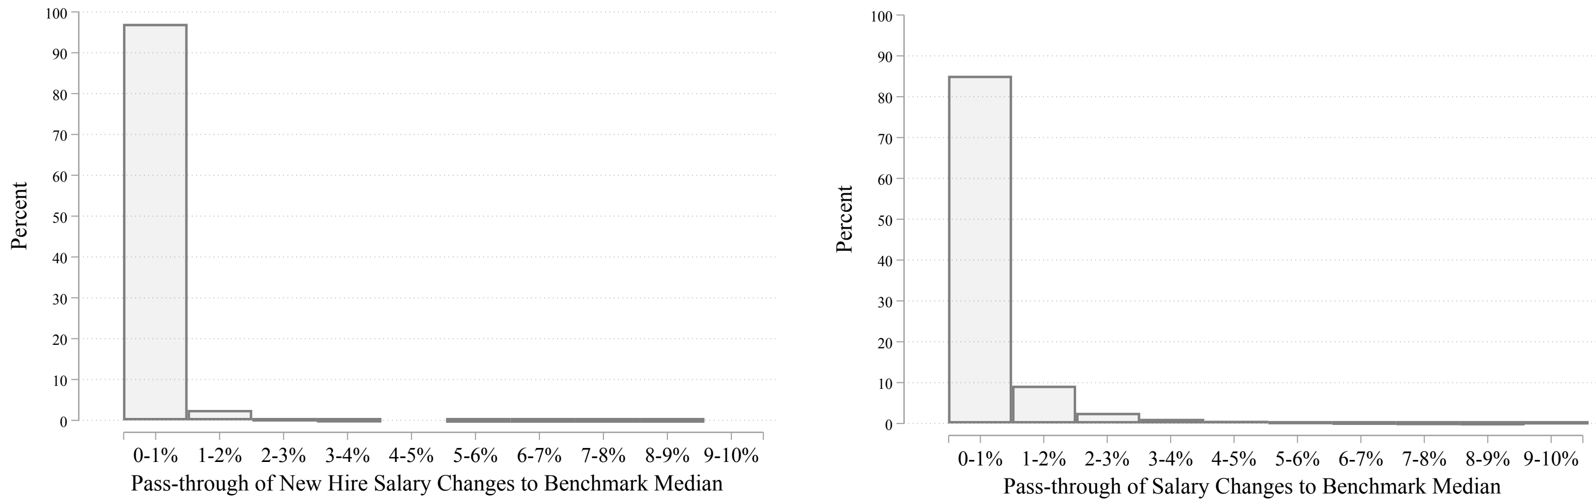

Notes: Distribution of the rate of pass-through for position-state-sector-firms with access to the benchmark. Panel A reflects the pass-through of raising all new hire salaries 10%, while Panel B shows the same but for all employees, not just new hires. Mean (median) pass-through for new hires is 0.23 (0.12) and for all employees is 0.59 (0.25).

H Wage Responses to Benchmark Shocks

We corroborate our main results on salary dispersion using an alternative identification strategy. Inspired by Derenoncourt et al. (2021), we leverage rare cases where large firms change wages in a particular position by 10% or more. Due to the granular mappings of position titles, this constitutes a sudden and localized shock to the benchmark information: the benchmark for one position may change significantly while holding constant the benchmarks for other adjacent positions. We then track the time it takes for other firms hiring in the same labor market to converge to the new benchmark as a function of whether they searched for that particular benchmark versus an adjacent benchmark. For example, if a large firm increases wages for employees in the position title say "Fulfillment Center Workers", the benchmark will be affected; however the benchmark for an adjacent position title, say "Warehouse Laborers", will not be affected even though the two positions are quite similar. We compare firms that then Searched (vs. Non-Searched and Non-Searchable) Fulfillment Center Workers versus firms that Searched (vs. Non-Searched and Non-Searchable) Warehouse Laborers, and compare the wage evolution with their respective benchmarks.

Figure H.1 shows how both the benchmark and salaries respond to the shock. We identify large firms as those with more than 95th percentile pass-through rates, as defined in Appendix G above. Panel A corresponds to the case where a large firm raises the salaries in a specific position-state-sector market by at least 10% in a single quarter. The dashed black line illustrates that the salary benchmarks associated with these wage changes also rises, closely tracking the changes made by the large firms. Panel B reproduces Panel A, but for adjacent markets. Adjacent markets are classified by taking the closest position title in the same SOC group by restricting to the same skill level categorization and selecting the title with the most similar average tenure and 2017Q1 benchmark. Panel B suggests that the large firm wage changes do not mechanically affect the benchmark of adjacent positions. The dashed black line representing the salary benchmark rises gradually over the subsequent 5 quarters rather than discretely around the time of the wage change. Even by the 5^{th} quarter, the adjacent benchmark has only risen modestly by comparison.

We next compare firms Searched (vs. Non-Searched and Non-Searchable) positions for benchmarks that were affected, with Searched (vs. Non-Searched and Non-Searchable) positions for the adjacent positions. In Panel A, we see that five quarters after the wage change, the median wages of employees in Searched positions have converged toward the affected benchmark, reaching an average 10% wage increase (equivalent to the benchmark) 5 quarters after the wage change. By contrast, Panel B displays wages in adjacent markets. In these adjacent markets, wages of Search positions rise less than 5% on average 5 quarters out, converging to the lower benchmark. The fact that we see divergent salary paths between Searched (vs. Non-Searched and Non-Searchable) in Panel A, but not in Panel B, suggesting that it is unlikely that the differential convergence is due to factors unrelated to the benchmark access.

Figure H.1: Wage Responses

Panel A: Markets with Wage Increases

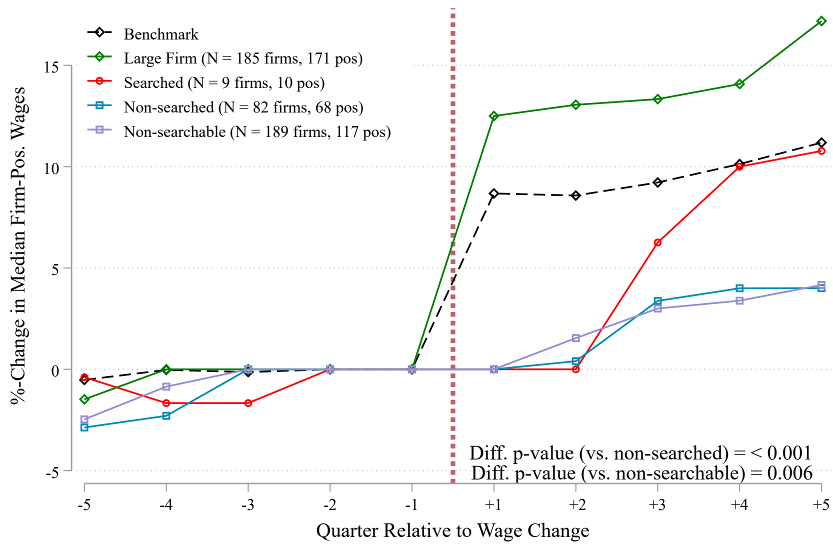

Panel B: Adjacent Markets

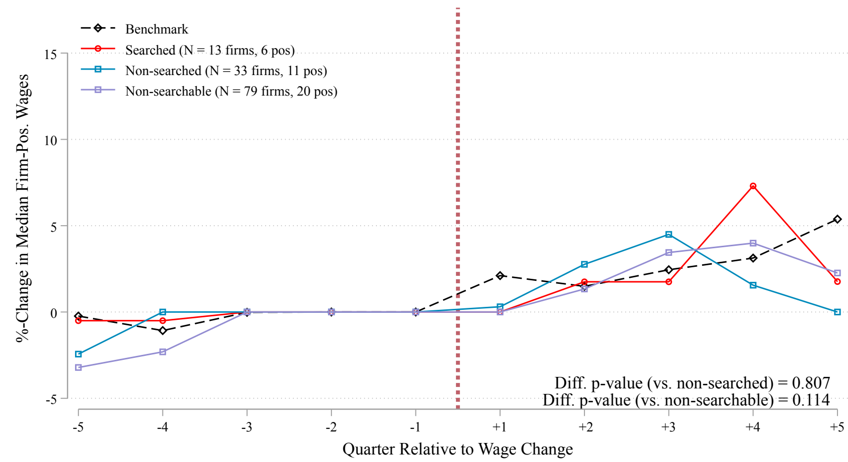

Notes: Panel A plots cases where we observed a firm (*Large Firm*) with greater than 95th percentile (24.5%) pass-through in a specific position-state-sector market raising their median wage in that market by at least 10%. We plot the median percent change in wages from the reference period and compare the quarterly benchmark and other position types in the same position-state-sector market, centered around the month of the wage change. P-values test the difference in medians in the percent change from the reference period in quarter 5 between Searched and control positions. Panel B shows those progressions in "Adjacent Markets", or similar position titles in the same state and sector. We identify similar position titles based on SOC group, average tenure, skill classification.

I Expert Prediction Survey

I.1 Survey Design

To assess whether the experimental results are surprising, we conduct a forecast survey with a sample of experts. A sample of the full survey instrument is attached as Appendix K. In this survey, which follows best practices (DellaVigna and Linos, 2022), we start by describing the benchmarking tool and the context. Then, we outline a hypothetical experiment where some firms are randomly given access to salary benchmarks and other firms are not – we opt for this simpler version because the full quasi-experimental design would have added too much complexity to the survey. We then elicit beliefs about the effects of access to the benchmarking tool on the distribution of salaries around the benchmark and the average salary level. We also included two questions about heterogeneous effects: one by education and one by gender. For each forecast, we elicit how confident the respondent feels about his or her answer, and we also include an open-ended question so that they can explain their choices.

I.2 Implementation

We collected responses from experts in two ways. First, we posted the survey on the Social Science Prediction Platform from May 6, 2022, to June 24, 2022. Second, on May 9, 2022, we emailed an invitation to the prediction survey directly to a list of 500 professors with publications related to the topic, and gave them 7 days to complete the survey. We excluded respondents who are not academics (8 respondents) or who had already seen our study (4 subjects). The final sample includes 68 experts. Of these, 11.8% responded to the survey through the Social Science Prediction Platform, and the remaining 88.2% responded through our direct invitation. This final sample is comprised of 90.7% professors, 2.9% PhD students, and 7.4% researchers. Around 91.2% of the respondents in the sample are economists, and 85.3% report having done research on labor economics.

I.3 Survey Results

The first result worth noting is that a majority of respondents did not feel confident about their own predictions. This is consistent with the fact that, prior to our study, there was little economics research on salary benchmarking and thus the experts do not have prior literature to base their predictions upon. Figure I.1 shows the distribution of certainty for each of the questions in the survey. If we pool all four predictions, the majority (64.3%) of responses were *Not Confident at All* or *Slightly Confident*, few respondents (4.0%) felt *Very*

Confident and no one ever responded *Extremely Confident*. There was more confidence in the responses to some questions than others; for example, 45.6% of respondents were *Very* or *Somewhat* confident in their response to the question about education heterogeneity.

To elicit beliefs about the effects on the distribution of salaries around the benchmark, we show six histograms (reproduced in Figure I.2) and ask the respondent to pick the histogram they think is most likely to represent the real data. One of those histograms looks like the results we see in the analysis (with compression from above and below the benchmark). The other histograms present alternative situations, such as no effect, compression only from above, compression only from below, and so on. Figure I.3 displays the frequency with which the experts chose each of the six histograms. A minority (30.9%) of respondents selected the histogram corresponding to our results, which showed compression from above and below.

To elicit beliefs about the effect of benchmarking on the average salary, we first asked respondents whether they expect positive, negative or no effects on this outcome. If they responded that they expect positive or negative effects, we then elicit the effect size in percent terms. The results for this prediction are presented in Figure I.4. A slight majority of respondents predicted a null effect or close to a null effect. For this question, the predictions are the most accurate: 58.8% of predictions fall within the 90% confidence intervals of our estimates from Table E.1, overlaid on Figure I.4. Many of the open-ended responses to this question echo the sentiment of one respondent who reasoned that it would be "equally likely that [employers] would revise their salary up or down given the information from the benchmarking."

We also elicited beliefs about heterogeneous effects. We asked respondents whether high or low education positions will be more strongly affected by benchmarking. Panel A of Figure I.5 shows that a majority (61.8%) of experts predicted that high-education positions would be most strongly affected. This goes against our findings, according to which low-skill positions are more strongly impacted. In the open-ended responses, respondents often noted that there should be less compression at baseline among high-education positions, which also goes against our findings. One common reason why experts believed that high-education positions would be more strongly affected was that for that type of positions the "information about the true distribution should be more valuable."

Last, we asked respondents whether salary benchmarking would increase the gender pay gap, decrease it or leave it unchanged. Panel B of Figure I.5 shows that a majority (66.2%) of experts responded that the wage gap would be reduced. Open-ended responses among those who predicted a reduction in the gender pay gap often mentioned that "bargaining becomes less important if an external source anchors the salary" or that "the employer may rely less on individual negotiations and biases, which often work against women." This prediction also goes against our finding, in that we do not find any significant effects of salary benchmarking on the gender pay gap. However, this comparison must be taken with a grain of salt, as we do not have enough statistical power to rule out small negative or positive effects.

Figure I.1: Experts' Confidence In Their Own Predictions

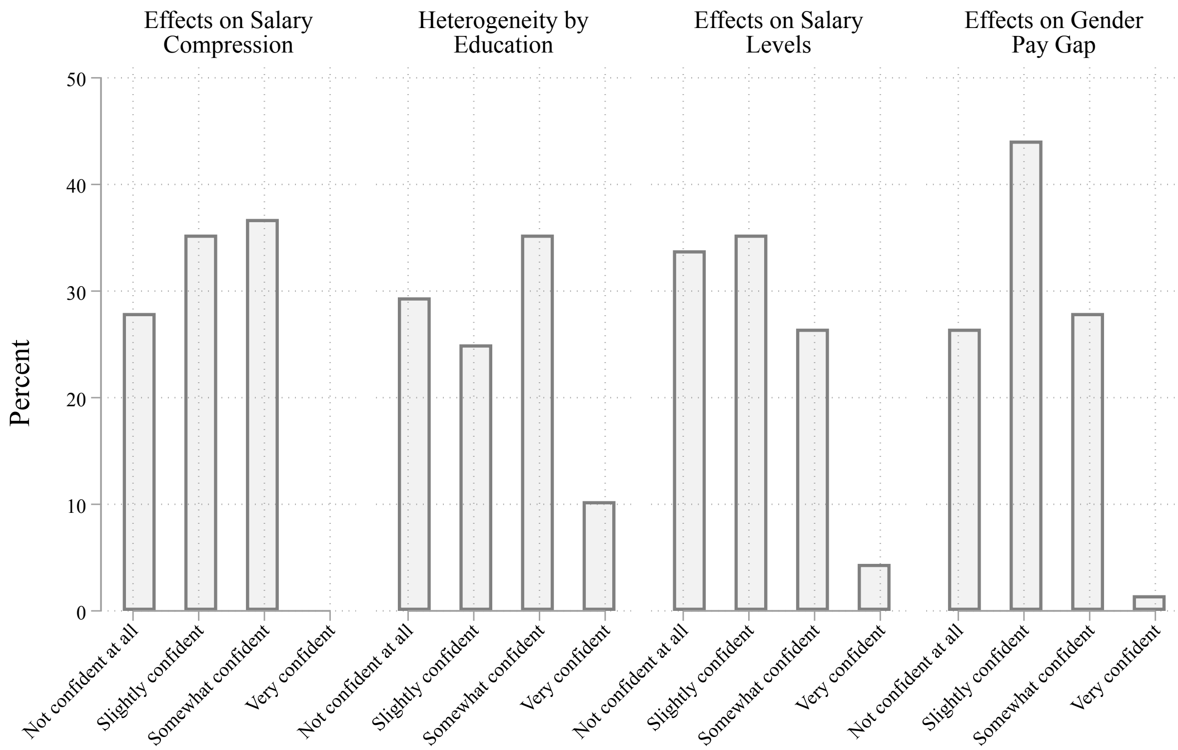

Notes: N = 68. Histogram of the response certainty for each question in the survey. Possible answers are *Not Confident at All*, *Slightly Confident*, *Somewhat Confident*, *Very Confident* and *Extremely Confident*.

Figure I.2: Expert Prediction Choice Set Regarding Salary Dispersion

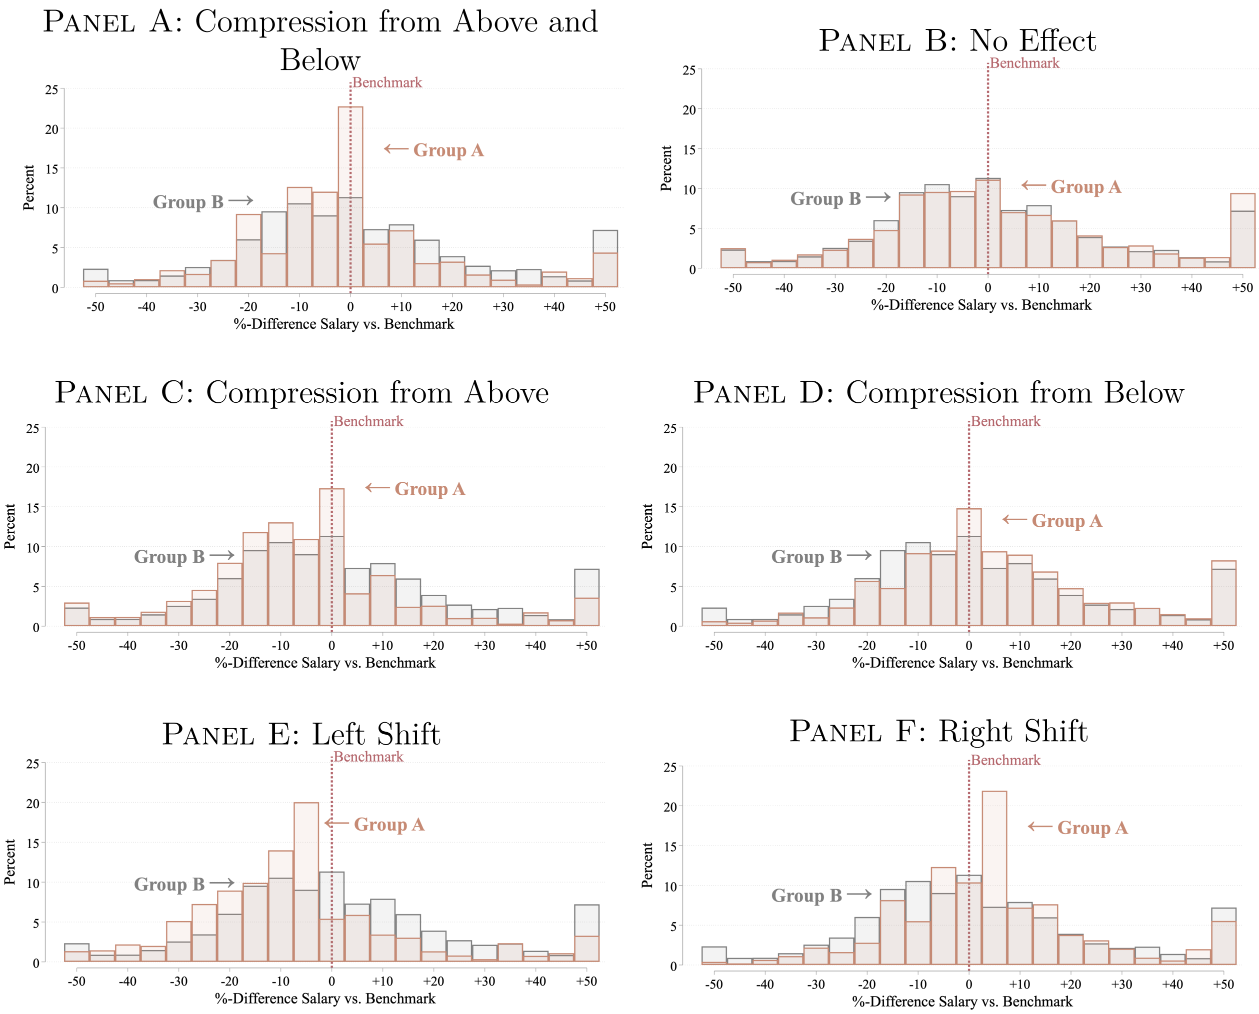

Notes: These are the images respondents could choose from when considering the compression effects of the benchmarking tool. Each figure is intended to show the effect described in the panel title. Panel A is an altered reproduction of Panel A of Figure 2.

Figure I.3: Expert Predictions Regarding Salary Dispersion

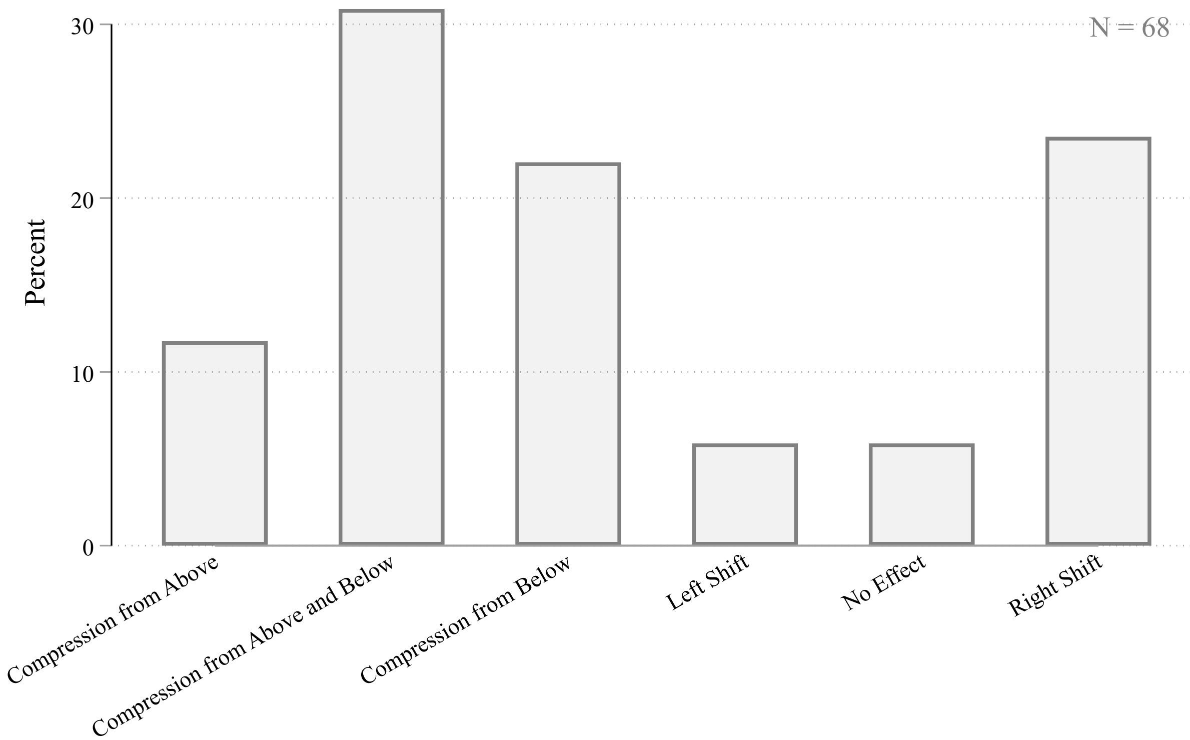

Notes: Histogram of the responses to survey question about the effects of benchmarking on compression. The six possible choices are displayed in Figure I.2.

Figure I.4: Expert Predictions Regarding the Average Salary

Notes: Histogram of the responses to survey questions about percent change in salary levels. Respondents were presented with a text box if they predicted salaries would go up or down. If they responded salary levels would stay the same, we include that here as 0% change. *Estimate 1* is the Searched vs. Non-Searched estimate of salary level effects from Table E.1 and *Estimate 2* is the Searched vs. Non-Searchable estimate. Displayed are the 90% confidence intervals.

Figure I.5: Expert Predictions Regarding Sources of Heterogeneity

Panel A: Education

Panel B: Gender

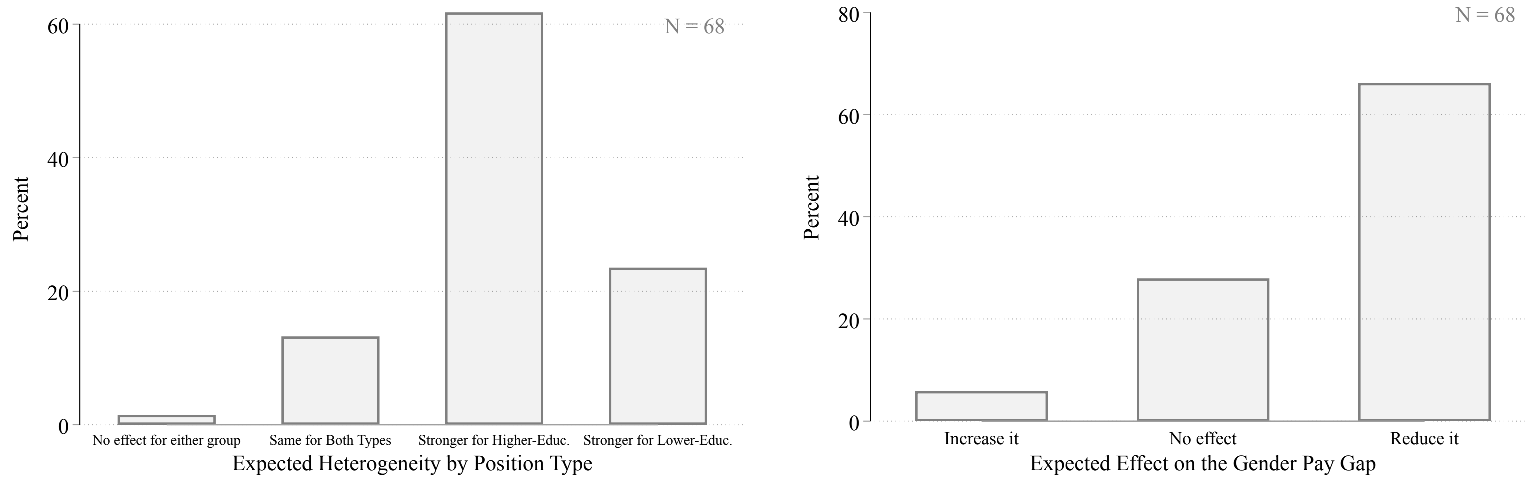

Notes: Histogram of the responses to survey questions about heterogeneity by education (Panel A) and gender (Panel B) in the effects of benchmarking.

J SHRM Survey

The following is a short summary of this study to help you decide whether or not to be a part of this research.

Here is some Key Information about the study:

- We are asking you to take part in a research study because you have hiring expertise.
- If you agree to be in this study you will be asked to complete a 10-minute online survey. In the survey, you will answer questions about compensation. At the start of the survey, you will be asked a screening question to determine your eligibility for this study.
- Your participation is completely voluntary. You can choose not to participate, or you can agree to participate and change your mind later and your decision will not be held against you. Your refusal to participate will not result in any consequences or any loss of benefits that you are otherwise entitled to receive. You can ask all the questions you want before you decide.
- If you have questions, concerns, or complaints, or think the research has hurt you, talk to Professor Cullen. She can be reached at 617-495-1876, or zcullen@hbs.edu.

○ Yes, I agree to take the survey

→

Do you participate in setting the salaries for employees?

- Yes
- No

→

How many years of experience do you have setting salaries?

→

How many employees does your company have (please consider all locations)?

- 1-49
- 50-99
- 100-999
- 1,000-4,999
- 5,000 or more

→

What main industry do you operate in? (start typing, then select a category that best describes your business.)

Please select a category from the list below to continue.

→

Are you working in the private sector or the public sector?

- Private sector
- Public sector

→

How would you describe your current role?

- Human Resources Professional
- Chief Human Resources Officer
- Executive (outside HR division)
- Manager (outside HR division)
- Recruiter (outside HR division)
- Other

→

Do you participate in salary settings for:

- New hires
- Current employees
- Both

→

Suppose you wanted to know the median salary that your company pays employees in a specific position. Would you be able to access that data?

- Yes, I can access it easily
- Yes, but it would take quite a bit of work
- No, I could not access that data even if I wanted to

→

When setting the compensation of their employees, some organizations use aggregate data on the market salaries for specific positions. This type of data is typically referred to as ***salary benchmarks***.

In your organization, do you use ***salary benchmarks***?

- Yes
- No

→

Which sources do you use to obtain ***salary benchmarks***? (Select all that apply.)

- ☐ Payroll data services
- ☐ Industry surveys
- ☐ Free online data sources
- ☐ Compensation consultants
- ☐ Paid online data sources
- ☐ Government data

→

What do you use the ***salary benchmark*** for? (Select all that apply.)

- [ ] To set salary ranges for specific job titles
- [ ] To plan ahead for headcount
- [ ] To determine salary in job advertisement
- [ ] To change salaries for current employees
- [ ] To set precise salaries for new hires
- [ ] In salary negotiations
- [ ] Other

→

How frequently do you use ***salary benchmarks*** to set salaries for new hires?

- For every hire
- A majority of hires
- Some of the hires
- A minority of hires
- Never

→

When do you use **salary benchmarks** *in relation* to new hires? (Select all that apply.)

- ☐ Before I publicize the position to include the expected salary in a job advertisement
- ☐ Right before I make an offer to the candidate
- ☐ After the candidate receives the offer, if the candidate wants to negotiate
- ☐ When the candidate presents an outside offer
- ☐ Other

→

Please explain briefly how you typically use ***salary benchmarks*** to set the salaries of new hires?

→

How frequently do you use ***salary benchmarking*** to change salaries for current employees?

- For all my employees
- For a majority of my employees
- For some of my employees
- For a minority of employees
- Never

→

When do you use ***salary benchmarks*** with current employees? (Select all the apply)

- [ ] When the employee goes through an annual review
- [ ] When the employee is up for promotion
- [ ] When the employee presents an outside offer
- [ ] When adjusting the salary ranges for positions
- [ ] Other

→

Can you please explain briefly how you typically use ***salary benchmarks*** to set the salaries of current employees?

→

Next, we'd like you to pick two different positions for which you are expecting to be hiring soon and tell us how you would set the salary for the new hire in each position.

Pick a position for which you are expecting to hire soon (start typing, then select a category - kindly allow a few seconds for the bold arrow to re-appear to continue.)

Please select a category from the list below to continue.

Pick a **second *different* position** for which you are expecting to hire soon (start typing, then select a category - kindly allow a few seconds for the bold arrow to re-appear to continue.)

Please select a category from the list below to continue.

→

Think about a future new hire in the role of Sales Engineers. What would be the **annual base salary** that you set for this person? (Please provide your best guess, and do not use commas.)

$

→

Suppose you look up the salary benchmark for this position using a highly accurate, up-to-date, compensation database and discover the median annual base salary is $93500. Upon reviewing that information, what salary would you pick?

$

→

Why did you use (or not use) the salary benchmark information in this compensation decision?

→

Think about a future new hire in the role of Sales Managers. What would be the **annual base salary** that you set for this person for a full-time position? (Please provide your best guess, and do not use commas.)

$

→

Suppose you look up the salary benchmark for this position using a highly accurate, up-to-date, compensation database and discover the median annual base salary is $92000. Upon reviewing that information, what salary would you pick?

$

→

Why did you use (or not use) the salary benchmark information in this compensation decision?

→

For which positions are ***salary benchmarks*** most useful?

- ◯ Most useful for lower-education positions
- ◯ Most useful for higher-education positions
- ◯ Equally useful for both groups
- ◯ Not useful for either group

→

Can you please explain briefly why salary benchmarks are most useful for higher-education positions?

→

Can you please explain briefly why salary benchmarks are most useful for lower-education positions?

→

Assume you are using a salary benchmark tool. The tool allows you to look at benchmark salaries, and to apply filters. When choosing filters there is a trade-off: applying filters can allow you to focus on a more relevant subgroup, but at the cost of smaller sample sizes and thus statistically imprecise benchmarks. Taking this into account, please select any filters from the set below that you would typically apply *after filtering for a particular position* (you can select more than one if you wish).

- ☐ Hourly vs Salaried
- ☐ State
- ☐ Industry
- ☐ Firm Size
- ☐ Revenue Size
- ☐ None of the above, the position-level filter is sufficient

→

Assume you are using a salary benchmark tool. It tells you the certain pieces of information about the salaries for that position. Please rank which information you typically care about the most (**drag and drop the options you care about, and then order them from most important to least important**). (1) Median salary (2) 10th Percentile (3) 25th Percentile (4) 50th Percentile (5) 75th Percentile (6) 90th Percentile (7) Average salary

Items	Order from most to least important (1=most important)
Median salary	
10th percentile	
25th percentile	
50th percentile	
75th percentile	
90th percentile	
Average salary	

→

For which position is the salary benchmark most useful for you? (start typing, then select a category - kindly allow a few seconds for the bold arrow to re-appear to continue.)
Please select a category from the list below to continue.

→

Think of your two closest competitors who also hire Retail Salespersons. For anonymity reasons, we'll refer to your competitors as firm A and firm B.

What is the maximum annual salary you think firm A would be willing to pay to hire in the full-time role of Retail Salespersons? (Please do not use commas)

$ [ ]

→

Let's say you find out that firm A would be willing to pay a maximum salary of $57500? After reviewing that information, what is the maximum annual salary you think firm B would be willing to pay to hire in the role of Retail Salespersons?

$

→

Suppose you lowered the salaries of your new hires by 10% and your competitors learned this information, what do you expect them to do with their new hires?

- Nothing, salaries of competitors would stay the same
- Salaries of competitors would fall between 1-5%
- Salaries of competitors would fall between 5-10%
- Salaries of competitors would fall 10%
- Salaries of competitors would fall by more than 10%

→

Suppose you raised the salaries of your new hires by 10% and your competitors learned this information, what do you expect them to do with their new hires?

- Nothing, salaries of competitors would stay the same
- Salaries of competitors would rise between 1-5%
- Salaries of competitors would rise between 5-10%
- Salaries of competitors would rise 10%
- Salaries of competitors would rise by more than 10%

→

If you adjusted the salaries of your new hires by 10%, do you think it would affect the salient numbers of the most commonly used salary benchmarks?

- No, the popular salary benchmarks would stay the same or shift by a negligible amount
- Yes, somewhat (eg. the median would shift by 1-2%)
- Yes, a significant amount (eg. the median would shift by >2%)

→

Have you ever used Glassdoor as your salary benchmark source?

- Yes
- No

Have you ever used [redacted] Data Cloud Compensation Explorer as your salary benchmarking source?

- Yes
- No

→

What share of your competitors do you think use [redacted] Data Cloud Compensation Explorer as a salary benchmarking source?

- The vast majority of my competitors
- Some of my competitors
- Very few of my competitors

→

Please share any feedback you have for us on the survey!

→

K Expert Prediction Survey

We are conducting an empirical study on labor markets. Due to your research record, you have been identified as an expert on the topic. We would love to elicit your expectations about the results of our analysis.

Please read the consent form below and click "I Agree" when you are ready to start the survey.

This survey involves no more than minimal risk to participants (i.e., the level of risk encountered in daily life). Participation typically takes between 5 and 10 minutes and is strictly confidential. Many individuals find participation in this survey enjoyable, and no adverse reactions have been reported thus far. Participation is voluntary, and participants may withdraw from the survey at any time.

Yes, I would like to take the survey

→

We would like to begin by providing some background information about this study.

When setting salaries, U.S. legislation prohibits employers from directly sharing compensation information with each other. However, employers are still allowed to use aggregated compensation data (e.g., median salary by position) provided by third parties. The practice of using aggregated market data is known as salary benchmarking.

We study the effects of salary benchmarking on the pay-setting of new hires. More precisely, we collaborated with a company that offers an advanced salary benchmarking tool that allows employers to look up market salaries in specific positions.

Employers in the Sample

Our sample covers a total of 1,982 firms from the United States, 583 of which gain access to the salary benchmarking tool. The average firm in our sample has 517 employees, ranging from 3 to 19,370 employees. Our firms cover all the main sectors in the U.S. economy, with the most common sectors being Manufacturing (21% of firms) and Finance and Insurance (14% of firms).

Employees in the Sample

The average employee in our sample earns $41,441 per year in base salary. On average, base salary accounts for the vast majority (92.7%) of the total compensation.

There are over three hundred unique positions that are looked up in the salary benchmarking tool. Some of the most commonly searched positions are Bank Teller, Customer Service Representative, Patient Care Coordinator and Software Developer.

→

Benchmarking Tool

To give you a bit more context, find below a screenshot of the benchmarking tool:

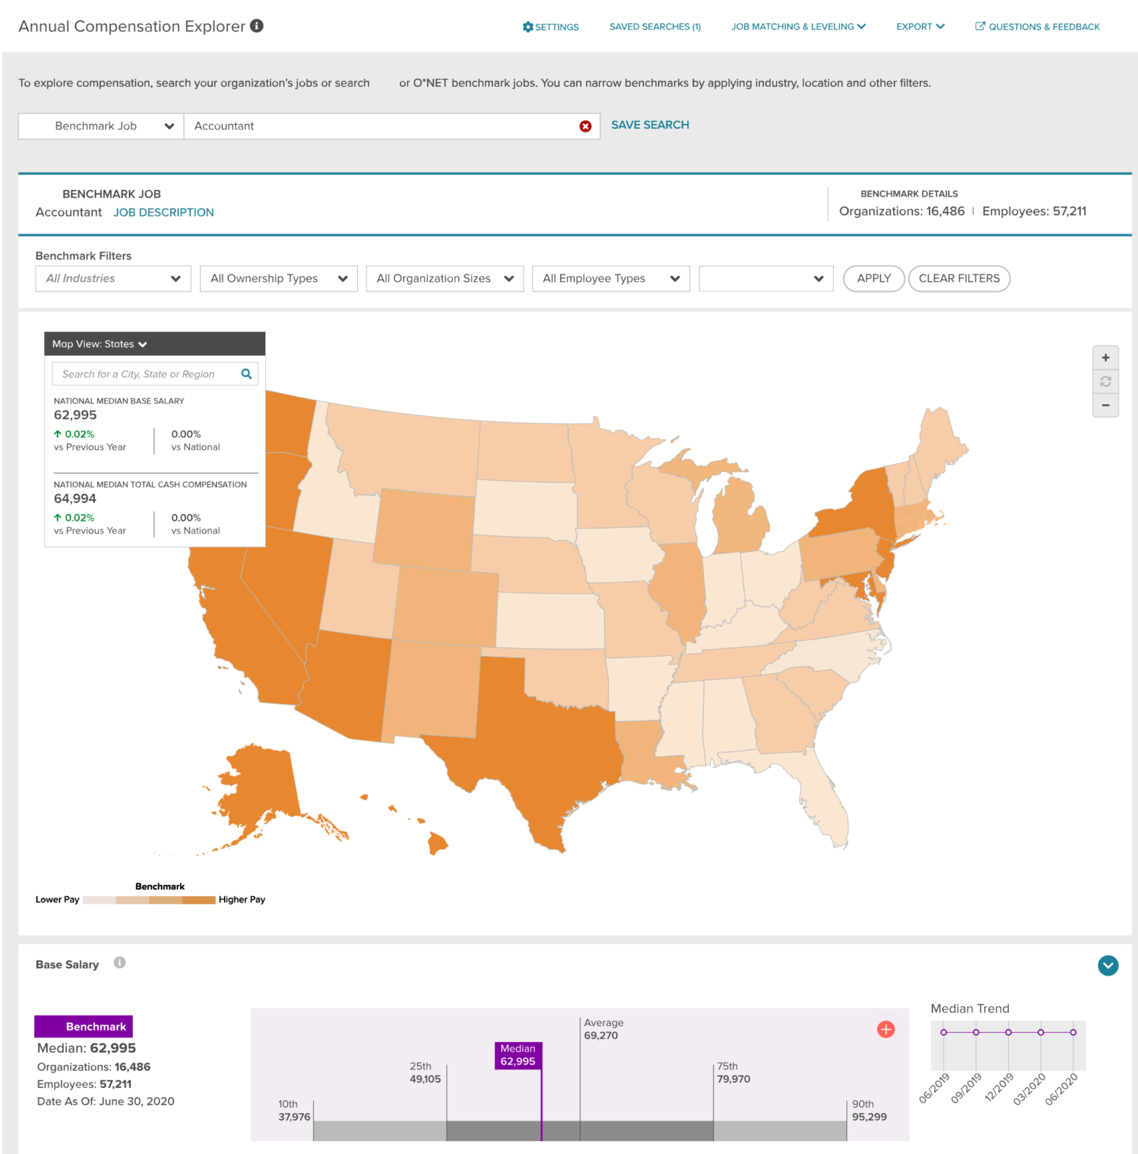

Employers can look up a any position (e.g., in the above screenshot, "Accountant"). Employers can apply filters to see the aggregate statistics within a specific state or industry, among other user-friendly features. The search results display the median annual base salary for the position, along with other key statistics (e.g., the 25th and 75th percentiles).

The benchmarks shown to the employers are of the highest quality. They are calculated using accurate payroll records from hundreds of thousands of firms and tens of millions of employees. As a result, the benchmarks are precisely estimated: e.g., in the above screenshot, the distribution of Accountants' salaries is based on 57,211 unique employees working at 16,486 unique firms. Moreover, the monthly frequency of the payroll records provides the most up-to-date benchmarks.

→

Are you familiar with the results from our research (e.g. have you seen it presented in a seminar)?

Yes

No

→

First, we want to elicit your forecasts about the effects of salary benchmarking on the average salary of new hires.

Consider the following thought experiment. Two employers (A and B) who just hired a new employee (e.g., a bank teller). The two employers are otherwise identical, except that Employer A was randomly chosen to gain access to a salary benchmarking tool, while Employer B was not chosen to receive access to the tool. As a result, employer A looked up the market pay before setting the salary of the new employee, while Employer B did not have access to that information at the time of setting the salary of the new employee.

Relative to Employer B (without salary benchmarking), do you think the average salary set by Employer A (with salary benchmarking) will be higher, lower, or about the same?

About the same

Lower

Higher

→

How confident are you regarding the previous forecast?

- Not confident at all
- Slightly confident
- Somewhat confident
- Very confident
- Extremely confident

Can you please explain briefly why you expect salaries to be lower, on average, for Employer A?

→

How much lower do you expect the average salary of Employer A to be (in percent terms)? Please enter a number between 0% and 100%.

0 %

→

In the previous question, we asked you about the effects of salary benchmarking on the **average** salary.

Next, we will ask you to forecast the effects of salary benchmarking on the **distribution** of salaries.

Consider two groups of employers:

Group A corresponds to employers ***with*** salary benchmarking: i.e., those who have access to the benchmark tool and look up the benchmarks before hiring a new employee.

Group B corresponds to employers ***without*** salary benchmarking: i.e., those who do not have access to the benchmark tool and thus cannot look up the benchmarks before hiring a new employee.

Consider the salaries of new hires relative to their corresponding benchmark (the median market salary for the position). For example, this is what the distribution of salaries looks like in Group B (without salary benchmarking):

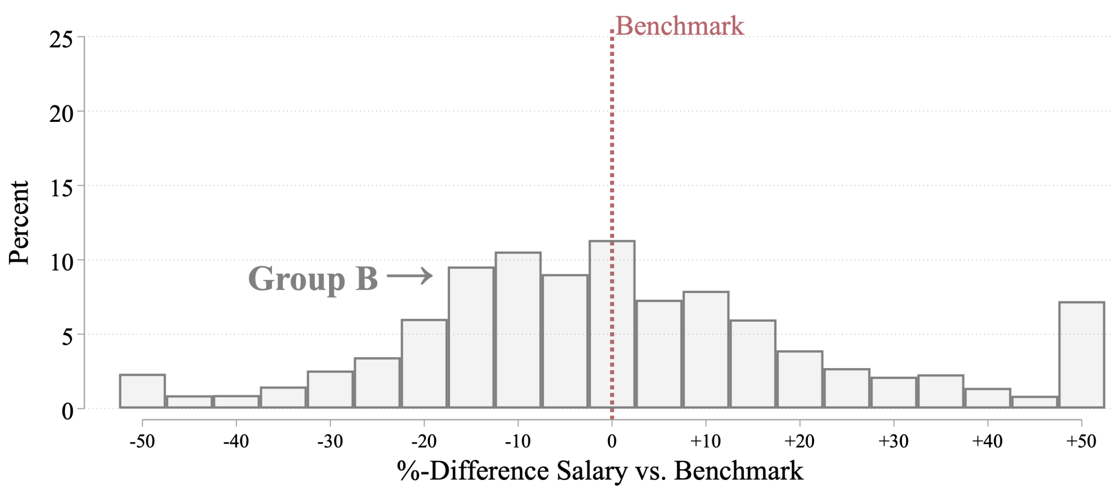

The middle bar corresponds to salaries that are close (i.e., within 2.5%) of the benchmark. The bars to the left of the middle bar correspond to new hires who are paid below the market benchmark, while the bars to the right of the middle bar correspond to new hires who are paid above the market benchmark.

Now that you have seen what the distribution of salaries looks like in Group B (without salary benchmarking), we want you to predict what the distribution would look like for Group A (with salary benchmarking).

Find below six histograms. In each of them, the gray bars denote the salaries in Group B (without salary benchmarking), while the red bars correspond to salaries of Group A (with salary benchmarking).

In your opinion, which of the following histograms best describes the effects of salary benchmarking?

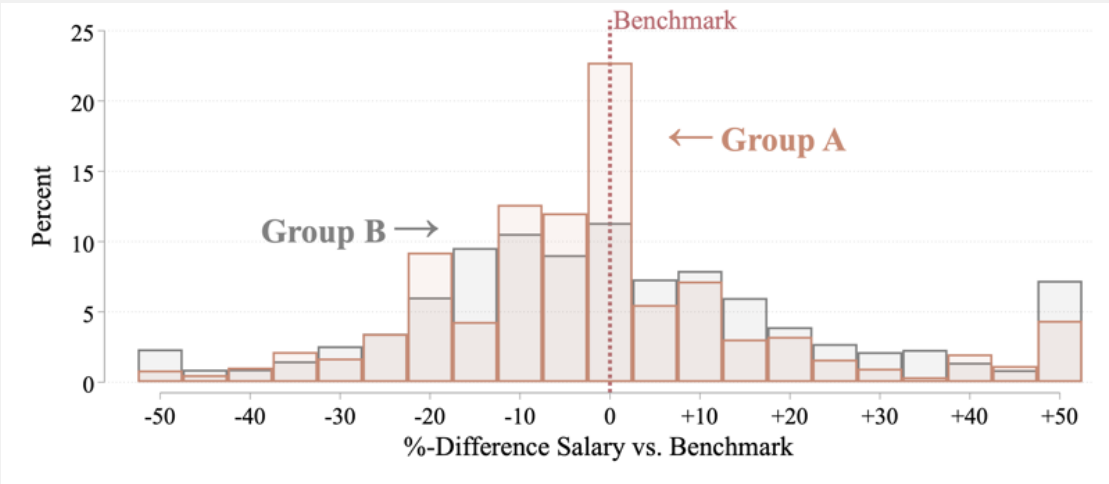

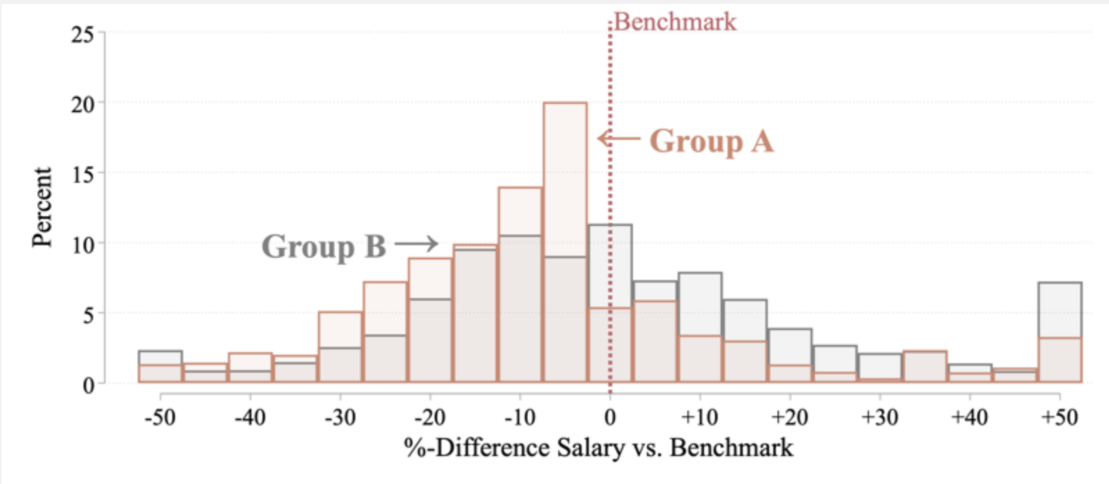

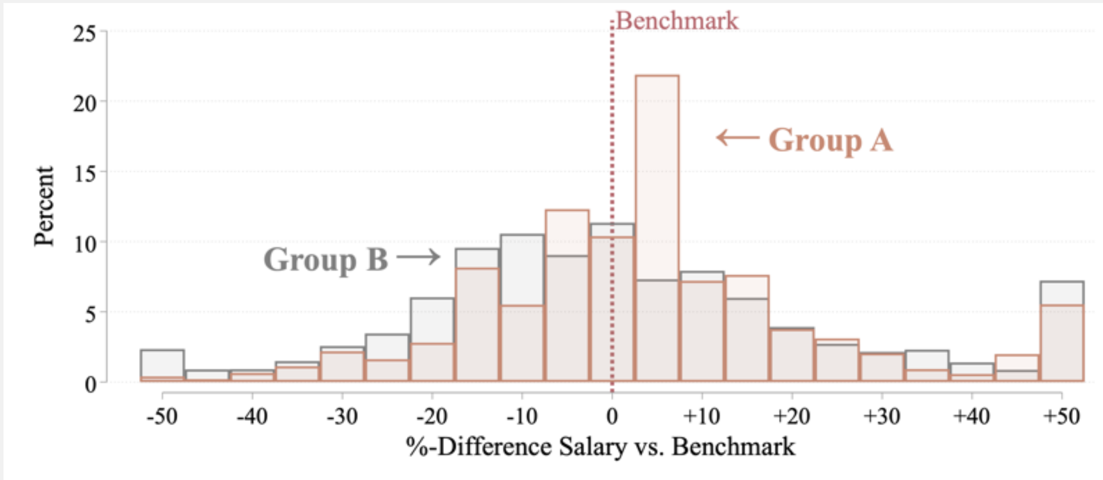

Benchmark
Group A
Group B
Percent
25
20
15
10
5
0
-50 -40 -30 -20 -10 0 +10 +20 +30 +40 +50
%-Difference Salary vs. Benchmark

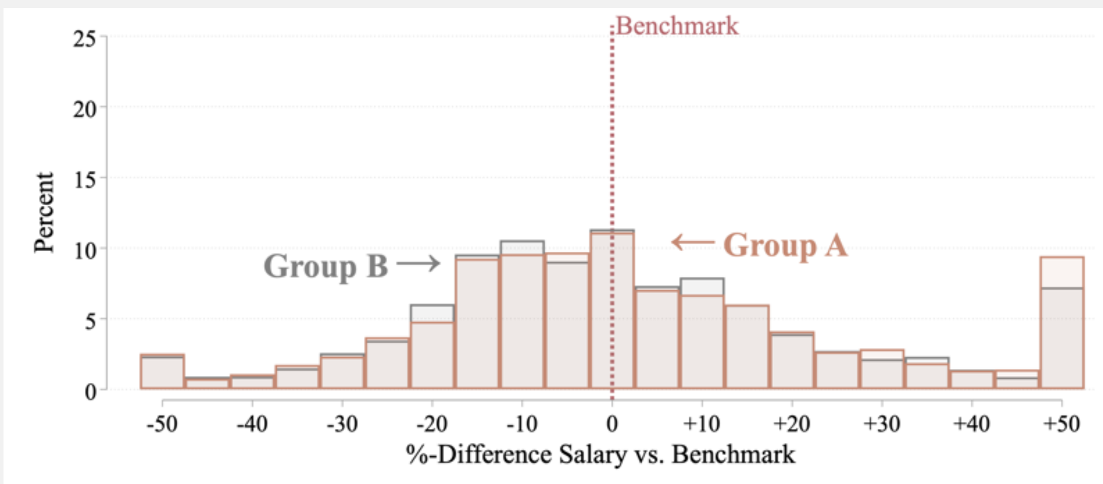

Benchmark
Group A
Group B
Percent
25
20
15
10
5
0
-50 -40 -30 -20 -10 0 +10 +20 +30 +40 +50
%-Difference Salary vs. Benchmark

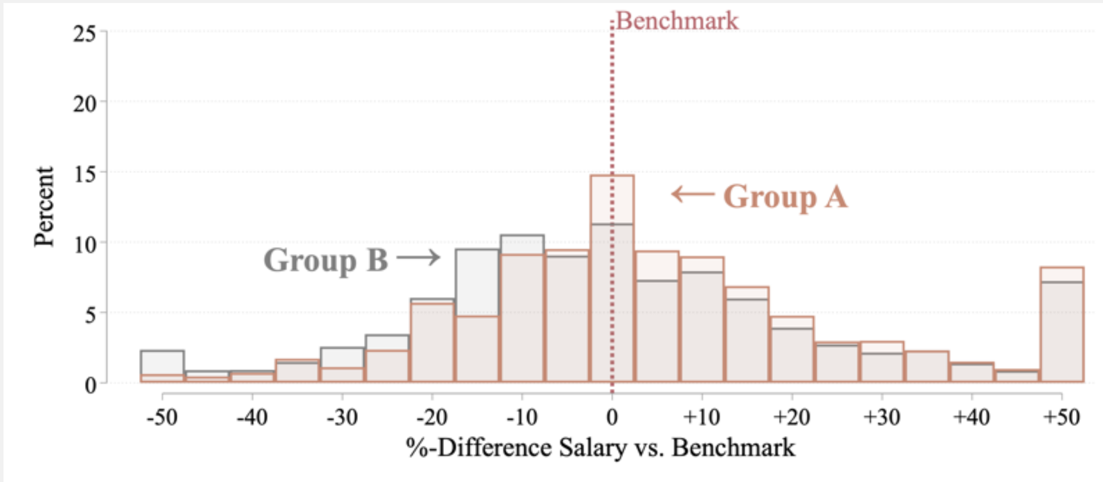
Benchmark
Group A
Group B
Percent
%-Difference Salary vs. Benchmark

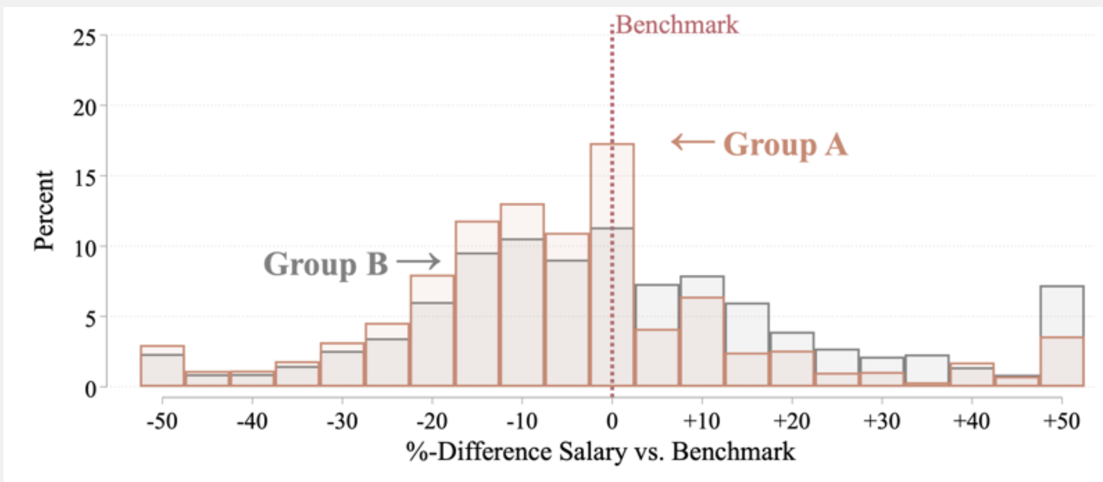
Benchmark
Group A
Group B
Percent
%-Difference Salary vs. Benchmark

→

How confident are you regarding the previous forecast?

- Not confident at all
- Slightly confident
- Somewhat confident
- Very confident
- Extremely confident

Can you please explain briefly why you think the histogram you selected best represents the effects of salary benchmarking?

→

Now, we want you to forecast which type of positions (if any) will be most strongly affected by salary benchmarking.

Consider lower-education vs. higher-education positions. The lower-education positions require little training and no more than a high school degree. The higher-education positions require more training and a College degree or more. Some common examples of lower-education positions are Bank Teller, Receptionist and Delivery Driver. Some common examples of higher-education positions are Legal Associate Specialist, Registered Nurse and Software Developer.

Do you expect the effects of salary benchmarking to differ between lower-education and higher-education positions?

- No effect for either group
- Stronger for lower-education positions
- Stronger for higher-education positions
- Equally strong for both groups

→

How confident are you regarding the previous forecast?

- Not confident at all
- Slightly confident
- Somewhat confident
- Very confident
- Extremely confident

Can you please explain briefly why you think the effects will be stronger for lower-education positions?

→

Do you expect salary benchmarking to affect the gender pay gap?

- No, it will not affect the gender pay gap
- Yes, it will reduce the gender pay gap
- Yes, it will increase the gender pay gap

How confident are you regarding the previous forecast?

- Not confident at all
- Slightly confident
- Somewhat confident
- Very confident
- Extremely confident

→

Can you please explain briefly why you think salary benchmarking will not affect the gender pay gap?

→

This is the last section of the survey. We would appreciate if you could share some information about yourself.

Are you currently one of the following: graduate student (either Master level or PhD level), faculty, post-doc or non-academic researcher?

- Yes
- No

→

Which of the following describes your current position?

- Professor (Associate or Full)
- Assistant Professor
- Post-doc
- Researcher
- PhD Student
- Master Student

Please select your discipline

- Economics
- Management
- Political Science
- Psychology
- Sociology
- Other

Do you have research experience in the following fields? Please select all that apply:

- Labor Economics
- Personnel Economics
- Public Economics
- Behavioral Economics
- Organizational Economics
- None of the above

→

This is the end of the survey. We thank you for taking the time to provide your forecasts!

If you have any comments for us, please leave them below:

→