

NBER WORKING PAPER SERIES

WHAT'S MY EMPLOYEE WORTH? THE EFFECTS OF SALARY BENCHMARKING

Zoe B. Cullen
Shengwu Li
Ricardo Perez-Truglia

Working Paper 30570
<http://www.nber.org/papers/w30570>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2022, Revised March 2024

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, Università della Svizzera Italiana, Norwegian School of Economics, UC-Berkeley (PF), UC-Berkeley (Labor), Amazon, University of Delaware, University of Copenhagen, University of Cologne, Goethe University, 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Zoe B. Cullen, Shengwu Li, and Ricardo Perez-Truglia. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What's My Employee Worth? The Effects of Salary Benchmarking
Zoe B. Cullen, Shengwu Li, and Ricardo Perez-Truglia
NBER Working Paper No. 30570
October 2022, Revised March 2024
JEL No. D83,J31,J38,M52

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.

Zoe B. Cullen
Rock Center 210
Harvard Business School
60 N. Harvard
Boston, MA 02163
and NBER
zcullen@hbs.edu

Ricardo Perez-Truglia
Haas School of Business
University of California, Berkeley
545 Student Services Building #1900
Berkeley, CA 94720-1900
and NBER
ricardotruglia@berkeley.edu

Shengwu Li
Harvard University
Department of Economics
Littauer Center
1805 Cambridge Street
Cambridge, MA 02138
shengwu_li@fas.harvard.edu

A data appendix is available at <http://www.nber.org/data-appendix/w30570>

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,

¹This 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. ophthalmic 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

²Notably, 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

³In 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).

⁴For 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-

⁵For 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.

⁶The 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

⁷The 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).

⁸Among 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

⁹We 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).

¹⁰Starting 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.

¹¹This 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.

¹²By 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.

¹³For 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

¹⁴The 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%.

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

¹⁶In 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.

¹⁷Given 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

¹⁸More 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.

¹⁹For more details, see Appendix B.2.

²⁰Due 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.

²¹We 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.

²²We 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.

²³More 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.

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

²⁵As 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.

²⁶As 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.

²⁷We 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

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

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

³⁰For 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.

³¹More 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

³²For 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.

³³More 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 ($\pm 2.5\%$) 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

³⁴The 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: $100 \cdot \left| \frac{\omega_{i,j,t} - \bar{\omega}_{i,t}}{\bar{\omega}_{i,t}} \right|$. 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:

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

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

³⁵In 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.

³⁶In 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:

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

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:

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

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.

³⁷One 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.

³⁸In 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 ($p\text{-value} < 0.001$), 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-elicite 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

³⁹In 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 proprietary 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 right 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 interchangeable. 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

⁴⁰For 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.

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

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

⁴³Among 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.

⁴⁴The 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

⁴⁵To 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.

⁴⁶A 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

⁴⁷More 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.

⁴⁸In 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.

⁴⁹One 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.

⁵⁰It 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).

⁵¹The 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

⁵²For 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 $F_s : [\underline{v}, \bar{v}] \rightarrow [0, 1]$. We assume that $0 \leq \underline{v}$, that $\bar{v} < \infty$, 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 $\eta : [\underline{v}, \bar{v}] \times \text{supp}(Q) \rightarrow \mathbb{R}_{\geq 0}$ 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

⁵³In 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).

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

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

⁵⁶Affiliation 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

$$m(v, c | q)m(v', c' | q) \leq m((v, c) \vee (v', c') | q)m((v, c) \wedge (v', c') | q). \quad (4)$$

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 $P(C \leq v | V_i = v, Q = q) = H(v | v, q)$. Let us define $\tau_q \equiv \inf \{v : H(v | v, q) > 0\}$. We assume that for all q and all $v > \tau_q$, there exists $\epsilon > 0$ such that $H(v | v + \epsilon, q) > 0$. 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:*

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

for

$$L(\alpha | v, q) \equiv \exp\left(-\int_{\alpha}^v \frac{h(\beta | \beta, q)}{H(\beta | \beta, q)} d\beta\right). \quad (6)$$

⁵⁷Jäger 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.

⁵⁸We 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*

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

and

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

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 $w^*(v, q) \leq v$. \square

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

⁵⁹For 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 $\alpha_i = \beta_i > 0$.

⁶⁰While 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.

⁶¹That is, (V_i, C) are affiliated conditional on Q and B , and there exists a distribution $H(c | v, q, b) \equiv P(C \leq c | V_i = v, Q = q, B = b)$ 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 $\bar{V}_i \equiv V_i + A_i$, equal to the revenue from filling the position plus the amenity. Moreover, assume that \bar{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 \bar{C} defined analogously. Then the function w^* can be reinterpreted as the total remuneration offered in equilibrium, as a function of \bar{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 $\eta : [\underline{v}, \bar{v}] \times \text{supp}(Q) \times \text{supp}(C) \rightarrow \mathbb{R}_{\geq 0}$ 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

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

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} ,

⁶²Observe that the payoff to a firm with value V_i and amenity A_i of hiring at total remuneration $w + A_i$ is $V_i - w = V_i + A_i - w - A_i = \bar{V}_i - (w + A_i)$. In particular, firms with the same \bar{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*

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

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

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

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 $E_{V_i, C} [\max\{V_i - C, 0\} \mid V_i = v, Q = q]$.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

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

⁶⁴This 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. Dworzak, 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*.

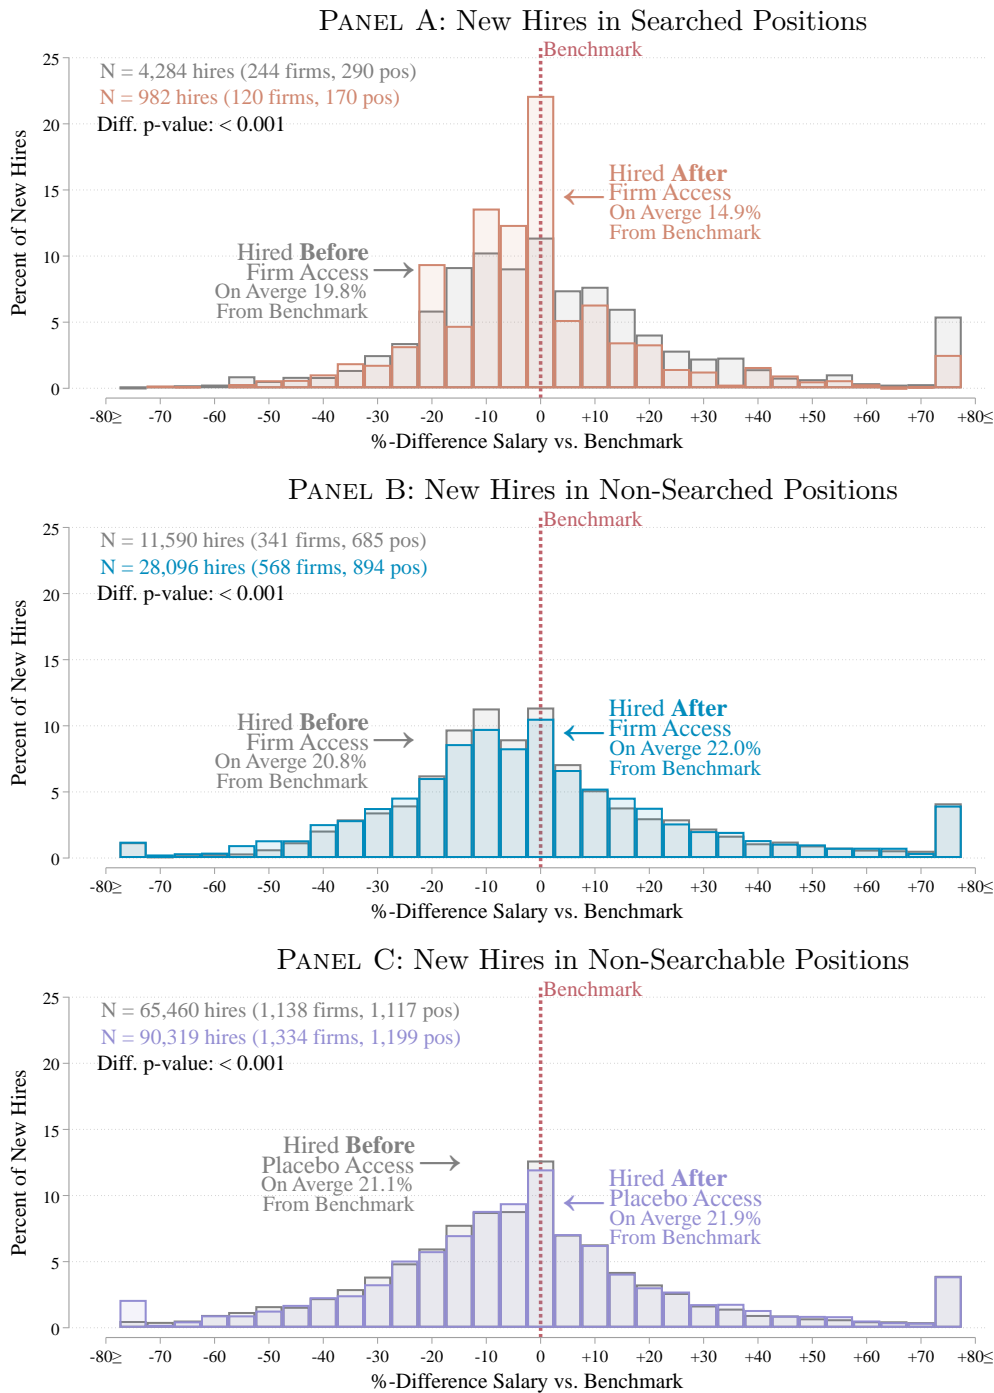
- Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of economic theory* 4(2), 103–124.
- Mas, A. (2016). Does Disclosure affect CEO Pay Setting? Evidence from the Passage of the 1934 Securities and Exchange Act. *Working Paper*.
- Mas, A. (2017). Does Transparency Lead to Pay Compression? *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



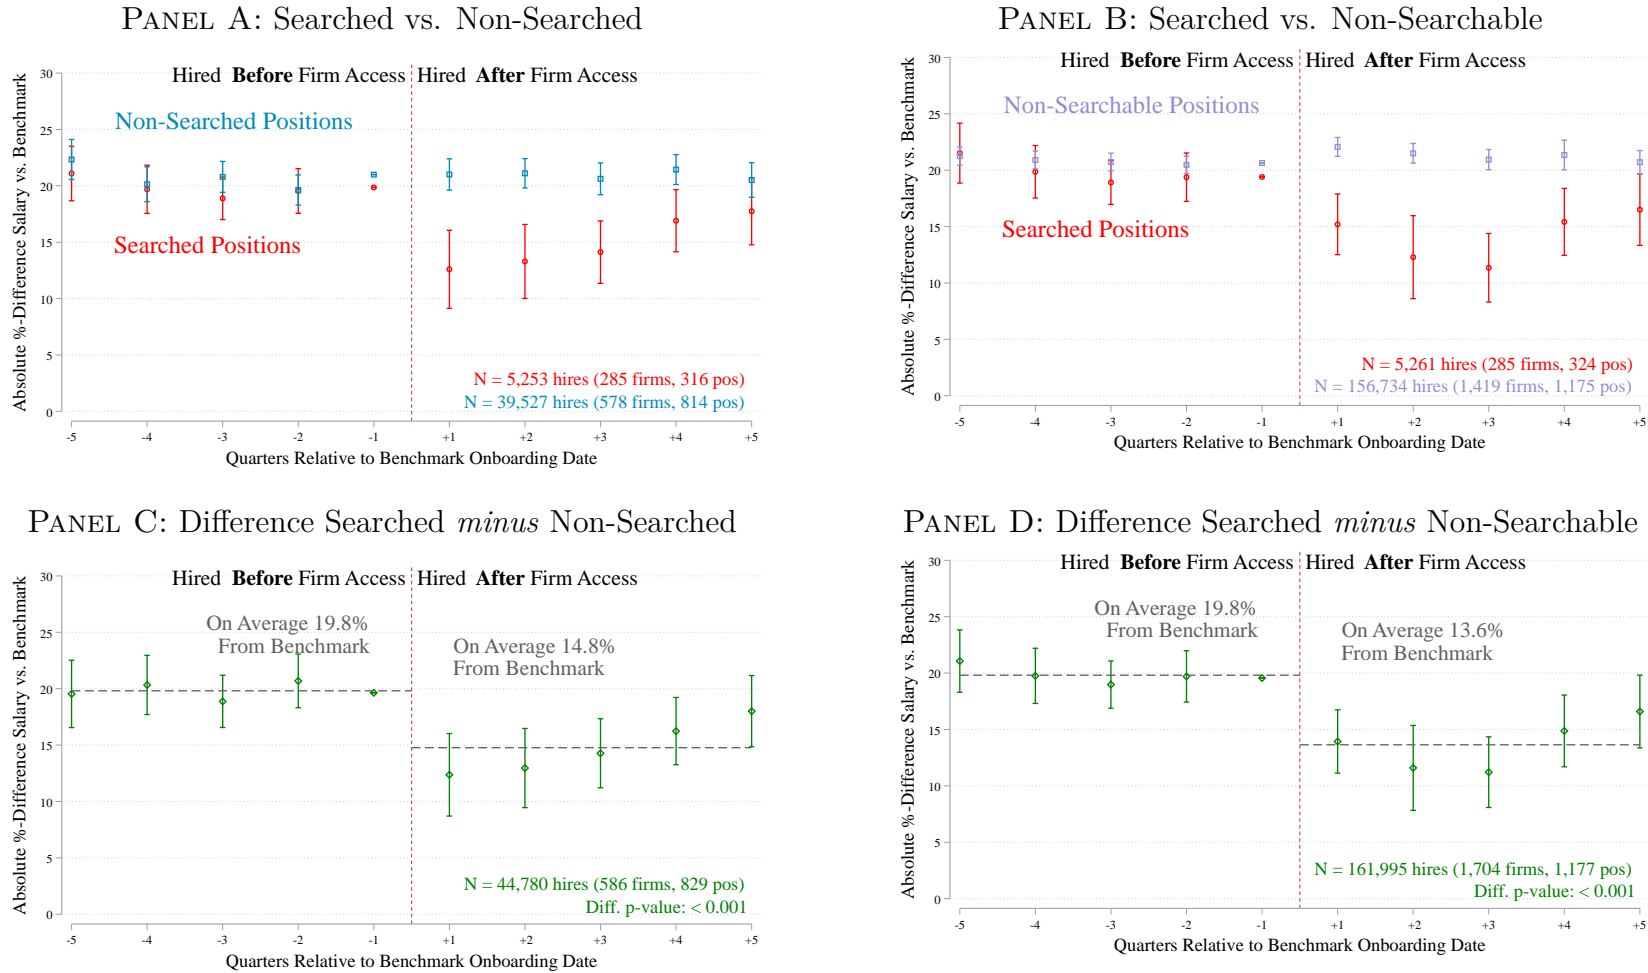
Notes: This is a screenshot of the pay benchmarking tool. It has been slightly altered to conceal the identity of the firm. This is the top of the screen. If you scroll down, you can see panels similar to the bottom panel titled *Base Salary* but for *Bonus*, *Overtime*, and *Total Compensation*.

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 $\pm 75\%$). 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).

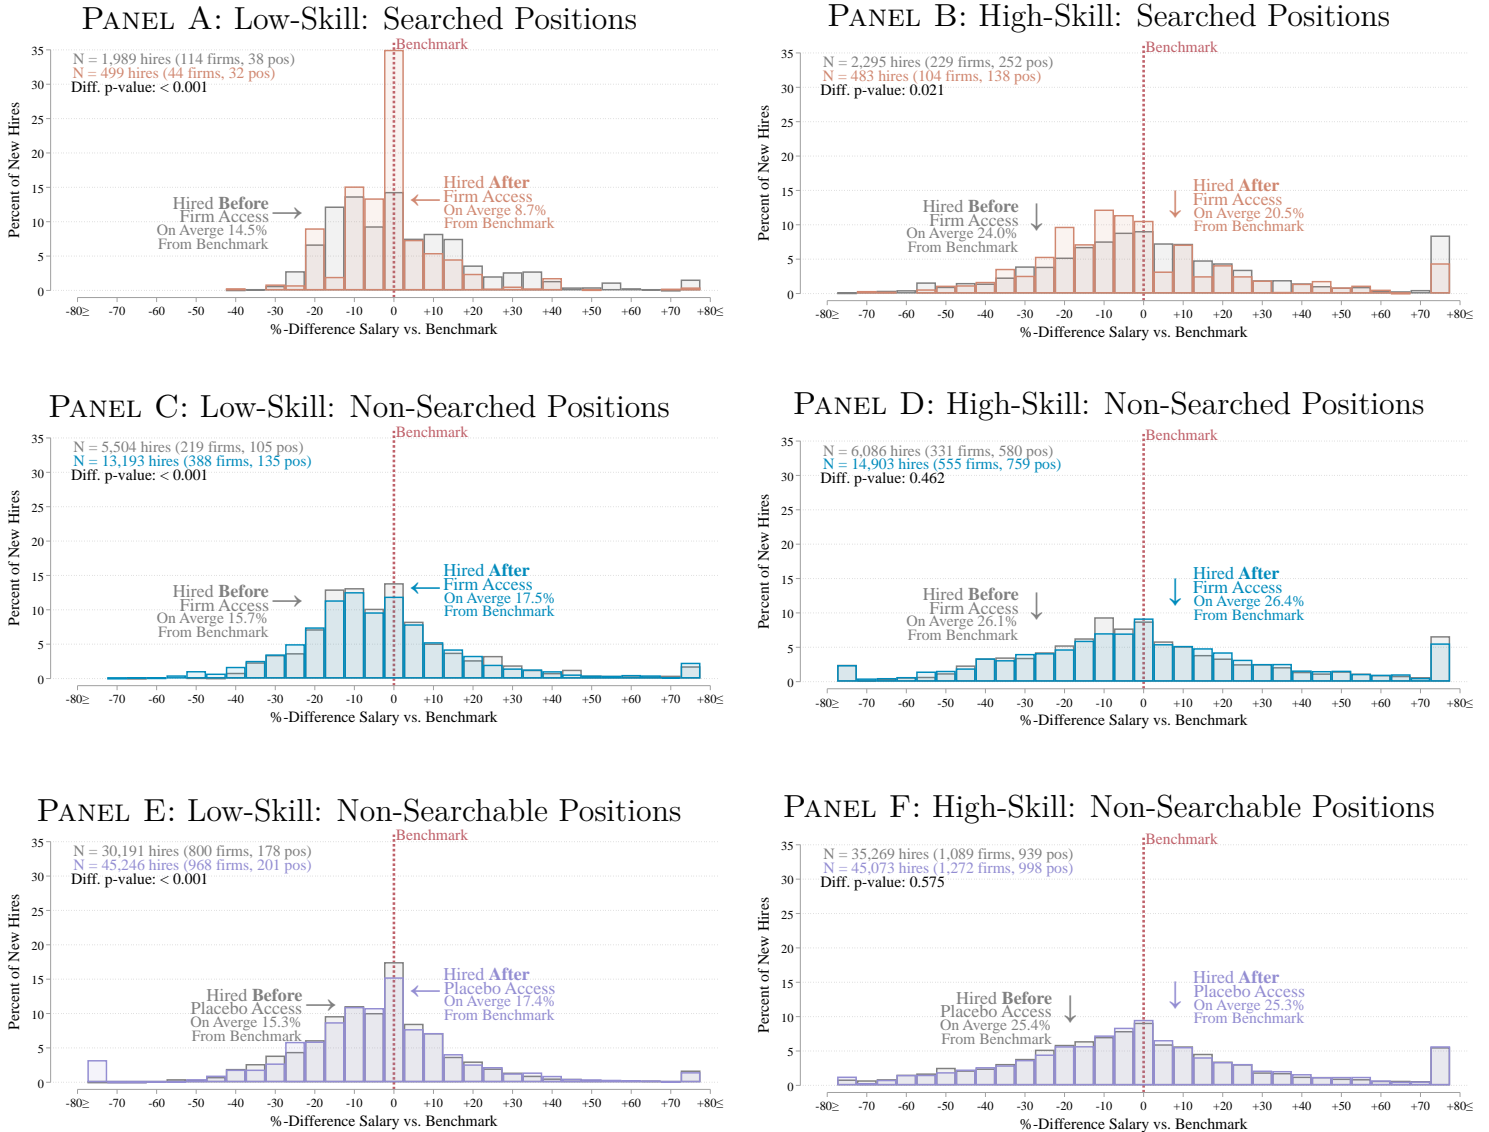
Figure 3: Event-Study Analysis: Effects on Pay Dispersion Around the Benchmark



45


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 $\alpha_{1,s}^k \forall s \in S$ 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




47


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 $\alpha_{1,s}^k \forall s \in S$ 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 (28.1)	509.8 (33.2)	483.2 (52.1)	525.7 (50.3)	444.5 (88.4)	482.1 (50.9)	483.9 (79.1)
Turnover Rate (%) [†]	2.424 (0.061)	2.438 (0.070)	2.382 (0.126)	2.392 (0.159)	2.374 (0.192)	2.680 (0.279)	2.192 (0.101)
Business Services Sector (%)	17.27 (0.99)	16.73 (1.13)	18.94 (2.07)	14.62 (2.71)	22.87 (3.07)	18.57 (3.30)	19.18 (2.67)
Hospitality Sector (%)	2.62 (0.42)	2.83 (0.50)	1.95 (0.73)	2.34 (1.16)	1.60 (0.92)	2.14 (1.23)	1.83 (0.91)
Retail & Wholesale Trade Sector (%)	12.04 (0.85)	11.97 (0.98)	12.26 (1.73)	16.37 (2.84)	8.51 (2.04)	12.14 (2.77)	12.33 (2.23)
Health Care Sector (%)	8.47 (0.73)	7.95 (0.82)	10.03 (1.59)	11.70 (2.46)	8.51 (2.04)	10.00 (2.54)	10.05 (2.04)
Banking Sector (%)	7.16 (0.68)	7.13 (0.78)	7.24 (1.37)	7.02 (1.96)	7.45 (1.92)	6.43 (2.08)	7.76 (1.81)
Other Sector (%)	52.44 (1.31)	53.38 (1.51)	49.58 (2.64)	47.95 (3.83)	51.06 (3.66)	50.71 (4.24)	48.86 (3.39)
Average Employee Characteristics							
Salary (annual) [†]	46,945 (794)	46,439 (956)	48,488 (1,356)	45,232 (1,632)	51,449 (2,103)	48,445 (2,366)	48,515 (1,634)
External Benchmark (annual) [†]	47,643 (652)	47,008 (752)	49,579 (1,307)	46,491 (1,650)	52,389 (1,977)	48,744 (1,931)	50,114 (1,754)
Abs. %-Diff. Salary vs. Benchmark [†]	22.16 (0.38)	22.46 (0.45)	21.26 (0.68)	19.41 (0.84)	22.95 (1.04)	21.28 (1.10)	21.25 (0.87)
Age	34.40 (0.18)	34.30 (0.22)	34.72 (0.32)	34.36 (0.42)	35.04 (0.48)	35.09 (0.55)	34.48 (0.39)
Share Female (%)	45.29 (1.29)	46.39 (1.48)	41.92 (2.57)	44.74 (3.78)	39.36 (3.51)	40.00 (4.09)	43.15 (3.32)
Share High Education (%)	56.92 (1.28)	55.30 (1.49)	61.84 (2.53)	57.89 (3.74)	65.43 (3.42)	62.86 (4.00)	61.19 (3.27)
Share Hourly (%)	71.89 (1.17)	73.08 (1.33)	68.25 (2.44)	71.35 (3.47)	65.43 (3.44)	71.79 (3.77)	65.98 (3.20)
Base Salary as Share of Total Comp. (%)	95.70 (0.20)	95.92 (0.25)	95.03 (0.33)	95.18 (0.45)	94.89 (0.47)	95.10 (0.48)	94.98 (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) All	(2) Searched	(3) Non-Searched	(4) Non-Searchable
Salary (annual \$) [†]	41,359 (146)	39,064 (462)	42,013 (390)	41,405 (166)
External Benchmark (annual \$) [†]	41,412 (113)	38,649 (409)	41,092 (295)	41,672 (128)
Abs. %-Diff. Salary vs. Benchmark [†]	20.36 (0.08)	17.36 (0.28)	21.03 (0.21)	20.45 (0.09)
Age	34.77 (0.05)	34.53 (0.22)	34.54 (0.13)	34.83 (0.06)
Share Female (%)	50.63 (0.20)	60.14 (0.83)	51.01 (0.53)	49.87 (0.23)
Share High Education (%)	42.21 (0.20)	34.49 (0.80)	42.28 (0.52)	42.76 (0.23)
Share Hourly (%)	81.11 (0.16)	82.94 (0.64)	80.13 (0.42)	81.16 (0.18)
Base Salary as Share of Total Comp. (%)	93.07 (0.06)	93.47 (0.19)	91.58 (0.14)	93.32 (0.06)
Occupation Groups				
Office and Administrative Support (%)	19.84 (0.16)	32.44 (0.79)	28.97 (0.48)	17.23 (0.17)
Building and Grounds Cleaning (%)	4.77 (0.09)	5.22 (0.38)	2.58 (0.17)	5.14 (0.10)
Management (%)	8.04 (0.11)	8.10 (0.46)	9.21 (0.31)	7.81 (0.12)
Production (%)	6.59 (0.10)	6.48 (0.42)	6.35 (0.26)	6.64 (0.11)
Transportation and Material Moving (%)	9.30 (0.12)	6.62 (0.42)	9.72 (0.31)	9.42 (0.13)
Other (%)	51.47 (0.20)	41.14 (0.83)	43.16 (0.52)	53.75 (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)
	$ \% \Delta $	$ \log \Delta $	$ \% \Delta > 10$	$ \% \Delta $	$ \% \Delta $	$ \% \Delta $	$ \% \Delta $	$ \% \Delta $	$ \% \Delta $	$ \% \Delta $	$ \% \Delta $	$ \% \Delta $
Panel A: Post-treatment												
Searched vs. Non-Searched	-4.775*** (1.143)	-5.155*** (1.266)	-16.270*** (3.626)	-5.148*** (1.338)	-4.775*** (0.906)	-4.786*** (1.198)	-5.324*** (1.282)	-4.950*** (1.286)	-4.421*** (1.153)	-4.887*** (1.165)	-4.880*** (1.276)	-4.564*** (1.178)
Searched vs. Non-Searchable	-6.149*** (1.070)	-7.118*** (1.211)	-13.861*** (3.681)	-6.836*** (1.220)	-6.149*** (0.824)	-6.128*** (1.076)	-7.494*** (1.233)	-7.450*** (1.576)	-5.714*** (1.078)	-6.163*** (1.087)	-5.044*** (1.231)	-5.934*** (1.127)
Panel B: Pre-treatment												
Searched vs. Non-Searched	-0.346 (1.167)	-0.129 (1.313)	-5.872 (3.690)	-0.233 (1.289)	-0.346 (0.751)	-0.488 (1.185)	-1.646 (1.514)	-2.062* (1.200)	-0.714 (1.133)	-0.144 (1.199)	-2.205 (1.528)	-0.199 (1.174)
Searched vs. Non-Searchable	-0.310 (1.055)	0.156 (1.175)	-4.221 (3.246)	-0.513 (1.184)	-0.310 (0.643)	-0.318 (1.057)	0.021 (1.375)	-1.029 (1.116)	0.241 (1.046)	-0.247 (1.069)	-0.754 (1.342)	-0.500 (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 ± 75 except in column (4) where it is winsorized to ± 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.