



Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Job Seekers' Beliefs and Preferences over Wages and Benefits: Experimental Evidence from Online Job Boards

Andreas Beerli, Stefano Fiorin, Andreas Gulyas, Daniel Kopp,
Mahsa Khoshnama and Michael Siegenthaler

Working Paper n. 734

This Version: October 2, 2026

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
<http://www.igier.unibocconi.it>

The opinions expressed in the working papers are those of the authors alone, and not those of the Institute, which takes non institutional policy position, nor those of CEPR, NBER or Università Bocconi.

Job Seekers' Beliefs and Preferences over Wages and Benefits: Experimental Evidence from Online Job Boards*

Andreas Beerli[†] Stefano Fiorin[‡] Andreas Gulyas[¶] Daniel Kopp^{||}
Mahsa Khoshnama[§] Michael Siegenthaler^{**}

October 2, 2026

Abstract

We estimate how strongly job seekers value wages and job benefits, how well informed they are about compensation differences across employers, and how much they update their beliefs when shown such information. In a field experiment, we randomly embed wage and benefit information from an employer-review platform on job ads for 265,000 users of a Swiss job-board network. The revealed-preference estimates suggest that a 1% higher posted wage increases clicks to view and engage with the ad by 0.4%. A conceptual framework decomposes this elasticity into preferences, prior knowledge, and belief updating, which we identify through a belief-elicitation experiment with users from the same job boards. We find that job seekers underestimate wage differences across jobs by half, know little about benefit differences, and update beliefs only partially. Accounting for these information frictions, the implied own-wage elasticity of applications is 2.5–2.9. Among the twelve benefits shown across the experiment, job seekers value four significantly: a company car (10% of wages), work from home (8%), firm-sponsored childcare (6%), and a parking spot (4%). The average position offers benefits worth 11% of wages. Because higher-paying jobs offer more benefits, inequality in total compensation exceeds wage inequality.

Keywords: Amenities, Job Benefits, Job Search, Wage Elasticity of Applications, Wage Transparency, Willingness to Pay

JEL Classification: J32, J33, J64

*Andreas Gulyas gratefully acknowledges funding by the German Research Foundation (DFG) through CRC TR 224 (Project A3). Michael Siegenthaler gratefully acknowledges financial support through SNSF grant no. 407740, entitled “What Workers Want.” This study is registered in the AEA RCT Registry (AEARCTR-0011023). It received ethical approval from EC Mannheim, Princeton, and ETH Zurich.

[†] KOF Swiss Economic Institute, ETH Zürich, IZA, and RFBerlin.

[‡] Bocconi University, IGIER, CEPR, and IZA.

[¶] University of Vienna and IZA.

^{||} KOF Swiss Economic Institute, ETH Zürich and IZA.

[§] Princeton University.

^{**} KOF Swiss Economic Institute, ETH Zürich, IZA, and RFBerlin.

1 Introduction

A fundamental question in labor economics is how strongly job seekers respond to wage and benefit differences across job ads. This response depends, in turn, on three channels: how much job seekers truly value wages and benefits in jobs, how much they already know about differences across jobs, and how much they update their beliefs in response to information in job ads. These channels are central to many core questions about labor markets, among them, the extent of employer wage-setting power, the consequences of pay transparency, workers’ willingness to pay (WTP) for benefits, and the contribution of non-wage benefits to labor-market inequality.

This paper combines a large-scale field experiment on one of the largest job-board networks in Switzerland with a complementary belief-elicitation survey experiment to separately identify these three channels. We randomly provide approximately 265,000 job seekers on over a hundred Swiss job boards with information on wages and the availability of non-wage benefits, and track their search behavior in detail. The wage and benefit information comes from employee reports on the market-leading employer-review platform. Since these reports cover a wide range of job attributes, we can simultaneously study twelve non-wage benefits, including company cars, firm-sponsored childcare, health measures, work from home, and flexible hours. Our novel experimental design recovers job seekers’ elasticity with respect to posted wages and their reaction to benefit information entirely from revealed-preference variation on a real job search platform. During our experiment, job seekers make over 12 million consequential decisions over nearly 200,000 real vacancies.

We develop a conceptual framework to guide the interpretation of the revealed-preference estimates from the field experiment. We model a job seeker’s decision to apply to a vacancy as depending on their preferences for wages and benefits and on their beliefs about the job’s compensation. When job seekers receive wage or benefit information, they form posterior beliefs as a weighted average of the information shown and their prior beliefs about the job’s wages and benefits. The framework highlights that the elasticity of applications with respect to posted wages or benefits does not directly identify a job seeker’s true underlying valuation. Even if job seekers value wages and benefits highly, they may not respond to wage and benefit information if they do not update their beliefs in response to the information or already know compensation differences across employers. We address this identification challenge by complementing the field experiment with a survey that presents job seekers with real ads tailored to their usual search scope. Some respondents are randomly given wage and benefit information for these ads, while others report their beliefs on the basis of the ad alone. This allows us to decompose the estimated effects of wage and benefit information into job seekers’ prior knowledge, belief updating, and structural preferences.

We run the field experiment on the 112 job platforms of Jobchannel, one of the largest private job-board networks in Switzerland. For three months in 2023, users of these platforms were randomly assigned either to a control group or to one of several treatment arms. The control group experienced the platform and search-results page as before. In the treatment arms, job seekers saw additional information on wages and benefits for job advertisements on the search-results page and on the full ad page. The information is sourced from kununu, the market-leading employer-review website in the German-speaking region. The wage information consists of the monthly full-time wage that kununu reviewers report for the specific job title and firm of the advertised vacancy. The benefit information indicates which benefits the posting firm offers, with a benefit classified as available if the share of reviewers confirming it exceeds a certain threshold.

On the job boards, we record users’ interactions in detail, including which vacancies were displayed on

their screen and all clicks on the website. Our interest lies in job seekers’ applications on the platform, which we proxy with two outcome measures. Our primary outcome is an ad view, defined as opening the vacancy page from the search-results list. Since the treatments may generate curiosity clicks that do not lead to further actions, our secondary outcome is a composite indicator for taking any consequential search decision, which comprises user clicks on the ‘apply now’ button on the website, navigating to the original job posting on the employer’s website, saving the ad to a watchlist, printing or sharing it. These ad actions capture real interest in applying after seeing the full vacancy information.

Our empirical analyses of the field experiment leverage two distinct sources of identifying variation. The first is *within-ad variation*: for the same job advertisement, some job seekers are randomly shown the average wage reported on kununu, while others are shown the median wage. An analogous design applies to benefits, where we experimentally varied the thresholds for confirming whether a benefit was available. The second source is *across-ad variation*, which we leverage within a continuous difference-in-differences design. Intuitively, we compare how the treatment changes users’ relative interest in higher- versus lower-paying vacancies—or in vacancies with versus without a given benefit—net of the corresponding differences in the control group. A key feature of our design is that we randomize information across job seekers rather than across vacancies. We can thus control for vacancy fixed effects, which absorb unobserved vacancy characteristics common across treatment arms, such as vacancy text and firm attractiveness.

The belief-elicitation survey experiment was conducted among approximately 1,500 job seekers whom we recruited from the Jobchannel platforms. The survey experiment mirrors the field-experimental design: respondents see eight real job ads from our field experiment personalized to their typical search region and occupation. Respondents in the control group receive the same result-list previews of job ads as the control group in the field experiment. The treatment groups receive previews with the same additional wage and benefit information as treated job seekers in the field experiment. After each ad, respondents report their beliefs about their expected monthly wages and benefits when getting the job. The control group’s answers reveal how well-informed job seekers are about wage and benefit differences across jobs in their typical search scope. Comparing these answers with the treatment group shows how much job seekers update their beliefs.

We document several important insights. First, posting wage information has a significant and robust effect on job seekers’ search behavior. Exploiting the across-ad variation, we estimate an elasticity of ad views with respect to posted wages of approximately 0.4. Extrapolating this posted-wage elasticity implies that displaying wage information on all vacancies would shift ad views toward vacancies in the top quintile of the wage distribution by approximately 30% relative to those in the bottom quintile. The wage elasticities are very similar when using the composite indicator of further ad actions as the outcome, indicating that higher posted wages translate into consequential search actions. The estimates are robust to a wide range of robustness and specification checks. For instance, using both the survey and field experiment, we find no evidence that wage information shifts job seekers’ beliefs about their likelihood of getting a job offer, the availability of not-shown benefits, or the workplace atmosphere. These findings mitigate concerns that the posted-wage elasticities are attenuated by directed search behavior or that information about one job dimension contaminates inferences about another.

Second, exploiting within-ad variation—which compares job seekers in the mean versus median wage treatment groups—yields an elasticity more than twice as large as the across-ad elasticity. Our conceptual framework rationalizes this difference. While both elasticities are driven by job seekers’ underlying wage

preferences and how much they update their beliefs in response to posted wages, the across-ad elasticity is additionally attenuated by their prior knowledge of wage differences across jobs. The within-ad elasticity, which focuses on variation *within* vacancies, does not depend on how much job seekers know about wage differences *across* vacancies. Consequently, the ratio of the two elasticities recovers job seekers’ prior knowledge—how closely their wage priors track benchmark wages. The comparison suggests that, for a job with a 10 percent higher benchmark wage, job seekers without wage information expect 5.4 percent higher pay. Thus, job seekers perceive roughly half of the wage differences between jobs. The belief-elicitation experiment provides an independent estimate of very similar magnitude: respondents without wage information expect 4.6 percent higher pay for a job with a 10 percent higher benchmark wage.

Third, we use the belief-elicitation survey experiment to estimate belief updating and to back out the structural wage preference parameter governing job search. The survey reveals that job seekers only partially update their beliefs about their own pay to wage information. The wage information closes 29 percent of the gap between the control group’s priors and the benchmark wage, which is in the range of recent estimates of belief updating to wage benchmarks [Altmann et al., 2026, Andresen et al., 2026, Bhatiya et al., 2026, Christensen et al., 2026]. The survey also shows that on average, job seekers in the control group underestimate full-time monthly wages by approximately 5 percent. Consistent with an underestimation of wages absent information, showing an average wage also increases ad views in the field experiment. Combining the field-experimental elasticities with the survey-based estimates of belief updating and prior knowledge, we recover the structural wage preference parameter—the elasticity of ad views with respect to job seekers’ beliefs about their own wages. Using across-ad variation, we find an elasticity of 2.5, and a very similar elasticity of 2.9 using the within-ad estimate. These results imply that job seekers are several times more responsive to differences in beliefs about their own pay than the posted-wage elasticity would suggest.

Fourth, we estimate job seekers’ WTP for job benefits as the ratio of the underlying preference parameters for benefits and wages. We recover it by combining the two experiments: the field experiment delivers the ratio of benefit-to-wage responses in actual search behavior. The survey adjusts this ratio for differences in how much job seekers know about benefits versus wages and in how strongly they update their beliefs about each. The survey reveals that job seekers are considerably less informed about benefit availability than about wages, implying that a comparison of posted-wage and posted-benefit elasticities would overstate the valuation of benefits by roughly 75 percent. Among the twelve benefits studied, the most highly valued is the provision of a company car, with a WTP of 10.4 percent of the wage. The second-highest valued benefit is the opportunity to work from home, at 8.3 percent. We also find statistically significant WTP for firm-sponsored childcare (6.1 percent) and parking spots (4.4 percent). The WTP estimates for childcare and company cars align closely with objective cost benchmarks. We find small and statistically insignificant valuations for the other benefits. A notable outlier is the company doctor, which we estimate to be a disamenity, perhaps because it signals hazardous working conditions. Interestingly, we find limited evidence of heterogeneity in responses to benefit provision across occupations. The main exception is job seekers in occupations with high female employment shares, who respond more strongly to work-from-home opportunities.

Finally, we combine the kununu wage and benefit information, the vacancy data, and the WTP estimates to compute the total job value in the sample of ads with matched kununu wage information. We define total job value as the sum of wages and the monetized value of benefits. We find that the average job offers benefits valued at approximately 11 percent of wages. Since higher-paying firms also offer more valued

benefits, inequality in total job value is significantly larger than wage inequality alone. For example, the Gini coefficient of total job value exceeds that of wages by approximately 6 percent. Hence, accounting for non-wage benefits amplifies, rather than compresses, compensation inequality across jobs. These results speak to a growing literature on the relationship between amenities and wages across jobs and firms [e.g., Caldwell et al., 2026, Costas-Fernández et al., 2026, Maestas et al., 2023, Ouimet and Tate, 2023, Sockin, 2024].

This paper contributes to several strands of the literature. First, we contribute to a growing body of work studying how job seekers respond to wages in job advertisements. Observational studies of job-board data [e.g., Banfi and Villena-Roldan, 2019, Escudero et al., 2025, Marinescu and Wolthoff, 2020] and field experiments that randomize wages on real or fictitious vacancies [Balgova et al., 2025, Belot et al., 2022, Dal Bó et al., 2013, He et al., 2023] typically find posted-wage elasticities of applicant interest in the range of 0.5–1.0. A conceptual contribution of our paper is to show that these posted-wage elasticities may substantially underestimate job seekers’ own-wage elasticity of job applications when job seekers have prior information about wage differences or update their beliefs only incompletely. We show that the own-wage elasticity of 2.5–3 is at least twice as large as the elasticity with respect to posted wages. Christensen et al. [2026], who adopt a belief-elicitation approach similar to ours, find a similar own-wage elasticity.

We also contribute to the literature on the valuation of non-wage job amenities. A long tradition in labor economics studies compensating differentials and workers’ WTP for job attributes [see Mas, 2026, for a recent overview]. Our paper is distinguished by its revealed-preference approach, in which job seekers make consequential search decisions over actual vacancies. Because job seekers on the platform search across a broad range of occupations, our estimates are not confined to a single occupation or recruiting setting. The inferential power of the large-scale experiment also allows us to estimate willingness to pay for a comprehensive set of twelve benefits, including rarely studied ones such as company cars, firm-sponsored childcare, and health measures. Finally, we show that job seekers update beliefs only partially to benefit information, suggesting that WTP estimates based on responses to posted information that do not account for incomplete updating would overstate the valuation of benefits.

In addition, we contribute to a growing literature on imperfect information in labor markets. Survey evidence indicates that workers anchor beliefs about outside options to their current wage [Andresen et al., 2026, Jäger et al., 2024], and that beliefs about pay at other firms are informative but compressed relative to true pay differences [Caldwell et al., 2025]. We add to this literature in three ways. First, we show that job seekers have informative but imperfect priors about relative pay differences within their search scope, with two independent estimates suggesting that they perceive roughly half of the wage differences across jobs. Second, when provided with wage information, job seekers close only about 29 percent of the gap between their priors and the benchmark wage. Third, we provide the first evidence on workers’ beliefs about job benefits, showing that workers are considerably less informed about differences in benefit provision than about wage differences.

Our results also have implications for the literature on employer wage-setting power. In canonical models of monopsony, wage-setting power derives from workers’ limited responsiveness to wage differences, reflecting either search frictions or preferences over non-wage amenities and idiosyncratic job characteristics [Azar and Marinescu, 2024, Manning, 2021]. Our decomposition suggests that both preferences for certain benefits and information frictions matter. The own-wage elasticity of job applications of roughly three implies that

job-seeker preferences remain an important source of wage-setting power even under full information. At the same time, job seekers underestimate pay differences across jobs by roughly half. Thus, imperfect information constitutes a second source of wage-setting power, in line with recent work incorporating information frictions into models of labor market power [Cheremukhin and Restrepo-Echavarria, 2025, Christensen et al., 2026].

Last, we contribute to the literature on the effects of wage transparency [see Bennedsen et al., 2023, Cullen, 2024, for literature reviews]. Job seekers today have access to far more information about pay, both because of pay transparency laws and the advent of employer-review and pay-comparison platforms such as Glassdoor, kununu, and Indeed. The existing literature has primarily studied how transparency legislation affects wage levels and wage gaps [Baker et al., 2023, Bennedsen et al., 2022, Gulyas et al., 2023]. Less is known about how transparency shapes workers’ job search behavior, a first-order mechanism through which greater transparency is expected to operate. Our field experiment provides direct evidence on this channel. The across-ad wage elasticity quantifies the extent to which posting benchmark wages redirects job seekers’ interest toward higher-paying vacancies, an explicit goal of the pay transparency policies recently introduced by the European Union and a growing number of US states. The scale of our experiment also allows us to estimate posted-wage elasticities separately by occupation. The posted-wage elasticity is roughly twice as large in the lowest-paying occupation tercile as in the highest-paying tercile and higher in occupations with a high share of female workers. These patterns suggest greater scope for wage transparency to redirect job search in lower-paid and female-dominated occupations.

2 Setting and Experimental Design

We combine a large-scale field experiment implemented on one of the largest private job-board networks in Switzerland with a complementary online survey experiment. In the field experiment (Experiment 1), we randomize whether and which information about wages and non-wage benefits is displayed to users and trace their revealed preferences through their search behavior. The survey experiment (Experiment 2) mirrors this design in a controlled online setting. It recruits job seekers from the same platform and exposes them to a personalized set of real job ads under analogous information treatments. It then elicits their beliefs to document both their priors and how these beliefs update when information is provided.

2.1 Experiment 1: Field Experiment

2.1.1 Experimental Design

The field experiment is conducted in collaboration with Jobchannel, which operates one of the largest private job-board networks in Switzerland, with a self-estimated market share of approximately 20% of the web traffic in the online vacancy market. Its network comprises more than one hundred job platforms, including generalist boards and portals with a regional (e.g., jobs-zuerich) or sector- or occupation-specific focus (e.g., healthcare, IT, hospitality). Together, these platforms list several hundred thousand vacancies per month and display job ads in users’ search-result lists millions of times each month.

All platforms share an almost identical layout. Job seekers search for vacancies by entering a job title or selecting from predefined categories and choosing a preferred region. The platform then displays a vertical list of matching job ads, as illustrated in Panel (a) of Figure 1. Clicking on a vacancy in the result list opens

the corresponding vacancy page, which displays the full vacancy text provided by the employer.

Our field experiment adds information about wages and benefits to job ads for some users.¹ The information was sourced from kununu, the leading employer-review website in the German-speaking region. Current and former employees can submit reviews on kununu that include information about wages, job characteristics, workplace culture, and the availability of benefits at their firms (see Appendix Figure A1 for screenshots of kununu’s wage-review and benefit-review forms).² In the information treatment arms, the platform disclosed the source and construction of the displayed information through mouse-over tooltips for desktop users (see Panel (b) of Figure 1 for an illustration). Hovering over the wage value or benefit icons revealed that the information was based on kununu reviews and gave details on how the information was computed. Survey evidence described below suggests that roughly three out of four respondents understood that the information comes from an employer-review website. The rest believed that the information stems from the employer.³

Users were randomized into different informational regimes.

Control group. Users in the control group saw the standard version of Jobchannel’s result list, which displays the posting firm’s average kununu five-star employer rating but does not include any wage or benefit information. Panel (a) of Figure 1 illustrates a typical search-results list shown to users randomized into the control group, similar to the status quo.

Information groups. Users in the information groups saw an additional gray bar for each job ad in the search-results list, displaying either wage information alone or wage information together with benefit availability (see Panel (b) of Figure 1 and Appendix Figure A5). The same information also appeared on the corresponding vacancy page (see Appendix Figure A6 and, for mobile users, Appendix Figure A7). The information bar appeared only when at least one type of kununu information (either wage or benefit information) was available for the ad.

- (a) **Wage-only information groups.** Users in the wage-only information groups saw wage information at the firm–job-title level. The displayed wage was constructed from reviews submitted by employees on kununu and was shown only when at least three such submissions were available for the firm and job title corresponding to the ad. As a result, wage information was available for approximately 10% of vacancies during the experiment.

Depending on the treatment arm, the displayed value corresponded to either the *average* or the *median* reported wage. Importantly, the two wage-information treatments are visually and

¹Unlike designs that require posting fictitious vacancies to generate experimental wage or benefit variation, our intervention involves only real job advertisements and displays additional information sourced from a third-party employer-review platform (kununu). Because we do not create non-existent job opportunities, we do not induce job seekers to spend effort evaluating or applying to fictitious openings. Some of the displayed information was, in principle, accessible independently of the experiment: employers could provide wage or benefit details in the vacancy text, and users in all arms, including the control group, could reach kununu reviews by clicking on the posting firm’s five-star rating. We thus do not restrict access to information for any users; at a minimum, the intervention made selected information salient at the search-results stage. Most respondents in the follow-up survey state that they would find wage and benefit information of this kind very useful (Appendix Figure A12).

²Up until March 2023, when the experiment started, kununu had collected 165,384 individual benefit reviews and 160,441 wage reviews covering close to 50k different employers in Switzerland. These microdata form the basis of our information treatments.

³Appendix Figure A10 shows that, among respondents shown wage information, 72% correctly attributed the information to an employer-review website, regardless of whether benefits were also shown or whether kununu was explicitly named as the source.

Figure 1: The Experimental Intervention: Search Results without and with Wage and Benefit Information



Notes: The figure provides an English-language illustration of a search-results list from the field experiment. Panel (a) shows the control-group version. The first line displays the job title as the main clickable element, and the second line shows the workplace location, the hiring firm, and, when available, the firm’s five-star rating from kununu. Panel (b) shows a version for the benefit and wage treatment groups, in which supplementary wage and benefit information from kununu is displayed below the standard job-ad content in a gray information bar. It displays wage information, benefit availability, or both. Black icons and text indicate that a benefit is classified as available at the posting firm, gray icons and text that it is not. For the third ad in Panel (b), the figure also illustrates the mouse-over tooltips shown to desktop users: hovering over the wage value or a benefit icon revealed that the information was based on kununu reviews and how it was computed.

linguistically identical: in both treatment arms, the mouse-over texts used the German term “mittlerer” (literally, “central” or “typical”). As a result, from the job seeker’s perspective, the only element that varies across arms is the displayed numerical wage. This design feature creates within-ad variation in the wage shown to job seekers: for the same job ad, some users saw the average kununu wage, while others saw the median wage. Appendix Figure A4 shows the distribution of gaps between log mean and log median wages. A broad set of job ads contributes to the within-ad wage variation: the mean differs from the median for 78% of ads, and the absolute gap exceeds 1% for 61% and 5% for 17% of ads.

- (b) **Wage-and-benefit information groups.** Users in the wage-and-benefit information groups saw wage information together with firm-level information about the availability of three benefits. Benefit availability was also constructed from employee reviews on kununu and was shown as soon as one review was available for the firm posting the ad. As a result, benefit information was available for 56% of vacancies during the experiment.⁴

⁴Benefit information had broader coverage than wage information because we used a lower display threshold for benefits and because kununu had started collecting benefit reviews several years before collecting wage reviews.

The set of three benefits varied across treatment arms. Across the experiment, the displayed benefits were drawn from a set of twelve benefits: flexible working hours, work from home, firm-sponsored childcare, parking spot, public transport connections, company car, canteen, food allowance, health measures, company doctor, coaching, and employee events.

For each benefit in a treatment arm, the information bar displayed whether the benefit was available at the posting firm or not. We classified a benefit as available if a sufficiently large share of kununu reviewers confirmed its availability. This threshold was 20% in most treatment arms. For three benefits—flexible working hours, work from home, and firm-sponsored childcare—some treatment arms instead used a threshold of 50%. When the threshold was met, the corresponding icon (a check) and text appeared in black. When it was not met, the icon (a dash) and text were shown in gray. Job seekers could therefore distinguish a benefit that is not offered from one about which no information is available. As with the average-versus-median wage variation, the two thresholds generate within-ad variation in the information shown to job seekers. For the same job ad, some users could be shown that a benefit was available at the firm, while others were shown that the benefit was not available.

The field experiment ran on Jobchannel’s job boards between March 6 and May 31, 2023. It was organized in three monthly waves, with eight approximately equally sized groups each month (always including a control group). Table A1 summarizes the full treatment-arm schedule. The March wave introduced the main treatment structure. Two wage-only arms displayed either average or median wages, while five wage-and-benefit arms displayed average wages together with three benefits grouped in thematic combinations, such as family-related, commuting-related, or health-related benefits (always based on the 20% threshold). The April wave focused on flexible working hours, work from home, and childcare. It retained the average-versus-median wage variation, varied the threshold used to classify benefit availability between 20% and 50%, and included an arm that reversed the order of the benefit icons and placed the wage information to the right. The May wave returned to the full set of twelve benefits and displayed average wages. It varied benefit combinations (always using the 20% threshold) so that each benefit appeared in multiple combinations across the experiment, rather than always being paired with the same two benefits.

Assignment to experimental conditions occurred only for users who accepted the platform’s statistics cookies. Jobchannel’s A/B testing framework operates only for these consenting users and is used for routine website optimization, including experimentation. As a result, both randomization and outcome measurement are restricted to users who accepted statistics cookies, who represent about 50% of all Jobchannel users during the experiment period. We define a user as a cookie identifier observed by the platform. Each month, randomization occurred at this user-cookie level, so that each user saw a consistent randomized information regime throughout their searches in a given month.⁵

⁵This implies that a single person may appear as multiple users when searching from different browsers or devices, or after deleting cookies. Conversely, multiple people using the same browser and device would be recorded as a single user. To gauge the magnitude of the first issue, we use the subsample of registered users, for whom the platform observes login identifiers. More than 75% of registered job seekers are associated with only one Google Analytics user ID, and the average registered job seeker has 1.66 such IDs.

2.1.2 Outcome Measures

A large part of the experimental literature on the valuation of wages and other job attributes uses stated-preference experiments in which respondents evaluate hypothetical job offers or report intended applications (e.g., Maestas et al., 2023, Wiswall and Zafar, 2018). Our field experiment instead provides revealed-preference evidence from users’ interactions with real vacancies on the platform. Our primary outcome is whether a user clicks on a job ad in the search-results list to open the corresponding vacancy page. We refer to this outcome as an *ad view*. An advantage of our setting is that the platform also records *ad impressions*: a job ad displayed in a user’s search-result list given her screen settings and browser. Therefore, we observe which vacancies appeared on a user’s screen and can clearly delineate the choice set from which job seekers select. Conditional on such an impression, the ad-view outcome captures whether the information shown in the search-results list changes users’ immediate interest in a vacancy.

In contrast to some field experiments that observe completed applications in more targeted recruiting settings (e.g., Abebe et al., 2021, Balgova et al., 2025, Dal Bó et al., 2013, He et al., 2021, Mas and Pallais, 2017), our experiment is embedded in a broad job-search platform covering a large cross-section of real vacancies. Applications to these vacancies do not occur exclusively through the platform itself. Many applications occur only after users navigate to the employer’s own website, and we cannot observe them directly.


Nevertheless, we also study a broader composite indicator for taking any *ad action* on the website. For each displayed ad, we observe whether the user applies through the platform, navigates to the original job posting on the employer’s website, saves the ad to a watchlist, prints, or shares it. The ad-action measure therefore captures search behaviors that indicate stronger interest than simply opening the vacancy page. Because most of these actions are taken after the job seeker has opened the ad, they also capture interest that persists once the job seeker has seen the full information. The ad-action outcome therefore helps address the concern that ad views merely reflect curiosity clicks on the result list that job seekers reverse upon viewing the full ad. In the empirical analysis below, we show that the estimated wage effects are very similar when using ad actions rather than ad views as the outcome. We also do not find evidence that job seekers reversed their initial interest after opening the ad. These results suggest that the information treatments shift downstream engagement, not only initial clicks.⁶

2.1.3 Job-Ad Sample and Coverage

Because the experiment studies how users respond to displayed wage and benefit information, the empirical analysis focuses on job ads for which we can construct at least some benefit information from kununu reviews. Panel (a) of Figure 2 assesses the coverage of job ads with matched benefit information relative to the broader Swiss vacancy market. It compares the number of job ads with benefit information that had at least one impression during the experiment and were active on March 31, 2023 with the number of job openings reported by the Swiss Federal Statistical Office on the same date. The comparison is constructed at the level of NACE sections. While the public sector is overrepresented in our matched sample relative to official job-opening statistics, the number of matched job ads is otherwise strongly correlated with official vacancies

⁶They are also consistent with recent large-scale evidence from Jalal [2026], which finds similar effects of pay disclosure on both job views and applications. In addition, Belot et al. [2022] find similar wage effects when using either views or saving for later application as measures of interest in artificial vacancies.

Figure 2: Coverage of the Experiment and Accuracy of Wage Information



Notes: Panel (a) plots the number of job ads with matched benefit information that had at least one impression during the experiment and were active on March 31, 2023, against the number of job openings in Switzerland reported by the Swiss Federal Statistical Office on the same date, by NACE section. Panel (b) compares the wages displayed in the experiment with wages according to official statistics, by two-digit ISCO occupation. The vertical axis shows the median kununu wage of matched job ads with wage information viewed during the experiment, computed across ads within the occupation. The horizontal axis shows the median monthly full-time-equivalent wage, averaged over the 2018, 2020, and 2022 waves of the Swiss Earnings Structure Survey. Markers are weighted by the number of matched ads with wage information.

across industries. This suggests that the industry distribution of matched ads broadly mirrors the overall vacancy market.

Table 1 summarizes the characteristics of job ads in the full sample and in the matched analysis sample. The analysis sample contains 198,079 matched job ads, compared with 352,534 ads in the full sample. Only 1% of job ads post a wage directly in the vacancy text. By construction, all matched ads have benefit review information, and 18.9% also have wage review information. The average review wage among matched ads with wage information is approximately CHF 6,600. Compared with the full sample, matched ads are somewhat more concentrated in manufacturing and trade and less represented in administrative and support services and hospitality. Their occupational distribution is nevertheless similar to that of the full sample, with professionals, technicians, clerical, and service occupations accounting for the largest shares. Matched ads are also posted by larger firms, reflecting the requirement that firms have at least one review in the underlying review data. On average, each matched ad receives about 64 in-sample impressions, 2.5 in-sample views, and 0.29 in-sample actions. With a total of 12.6 million ad impressions across nearly 200,000 matched ads, the sample size allows us to estimate treatment effects separately by detailed occupation and labor-market segment.

The wage analysis relies on a subset of roughly 37k matched ads with kununu wage information (Column 3 of Table 1). These ads are posted by larger firms and are more concentrated in some industries, especially trade, and less concentrated in the public sector. In terms of occupations, they overrepresent clerical support and service and sales occupations.

Table 1: Characteristics of Job Ads Shown in the Experiment

	Full Sample	Analysis Sample (Matched Sample)	With Wage Information
Has vacancy-posted wage (%)	0.56	0.97	1.37
Has review wage (%)	10.62	18.7	100
Avg review wage (mean)	6,614 (2,107)	6,614 (2,107)	6,614 (2,107)
Has benefit review information (%)	56.19	100	100
Reported WFH in vacancy (%)	4.91	5.94	5.71
Ad in English (%)	7.27	8.56	6.46
Ad in German (%)	82.52	83.34	86.18
Minimum FTE	83.84	83.29	80.56
Maximum FTE	93.78	93.79	92.7
Temporary position (%)	6.6	6.21	4.26
<i>Industry (%)</i> :			
Admin. and support services	8.4	3.16	6.89
Manufacturing	14.58	18	15.07
Construction	4.17	3.03	2.31
Trade	12.45	13.98	23.68
Hospitality	7.66	6.22	6.97
Public sector	25.88	27	16.46
<i>Occupation (%)</i> :			
Managers	3.77	4.11	1.11
Professionals	28.73	29.66	19.81
Technicians and associate professionals	19.52	21.01	20.61
Clerical support workers	14.11	16.43	29.42
Service and sales workers	16.65	14.94	18.57
Craft and related trades workers	10.74	8.28	7.74
Plant and machine operators, and assemblers	2.66	2.59	1.95
Elementary occupations	3.06	2.63	0.68
<i>Firm Size (%)</i> :			
1-10	17.41	4.22	1.03
11-100	34.91	25.99	12.59
101-1000	27.25	39.4	38.47
≥ 1001	20.43	30.39	47.91
Avg vacancy duration	88.96 (110.00)	84.83 (104.96)	79.32 (102.52)
Avg (in-sample) impressions	61.81 (145.18)	63.81 (137.57)	68.78 (142.79)
Avg (in-sample) views	2.5 (8.68)	2.55 (7.85)	2.9 (8.65)
Avg (in-sample) actions	0.26 (1.04)	0.26 (0.99)	0.32 (1.17)
<i>Number of job ads</i>	352,534	198,079	37,039

Notes: The table reports descriptive statistics at the job-ad level for the full and the experiment samples. The full sample in Column 1 is all ads with at least one impression that satisfy the sample restrictions described in Appendix C. Column 2 focuses on the ads matched with kununu benefit or wage information (the analysis sample), and Column 3 the ads with kununu wage information. “Has vacancy-posted wage” indicates that the ad contains a posted wage or wage range. The review wage is the mean monthly kununu wage (in CHF). “Reported WFH in vacancy” indicates that the vacancy text reports the possibility to work from home. Industry refers to the NACE section of the posting firm. Occupation refers to the one-digit ISCO group assigned to the ad. “In-sample” impressions, views, and actions count the analysis-sample impressions of the ad and the views and actions resulting from them. Standard deviations are in parentheses.

2.1.4 Users, Sessions, and Balance

Appendix Table A2 summarizes user and session characteristics in the full and analysis samples. The full sample consists of 275,983 users.⁷ Of these, 265,983 users saw at least one job ad with matched benefit information. These users and their 407,642 search sessions with at least one such ad form our analysis sample. Sessions last about 9 minutes on average and include 1.8 searches. Users see about 60 ad impressions per session, about half of which are matched job ads. About 77% of sessions occur within Switzerland. A bit more than 50% of sessions occur on mobile devices.

The occupational composition of sessions reflects the breadth of searches on the platform.⁸ Users most frequently search for professional occupations (26.6%), followed by clerical support (18.7%), technicians and associate professionals (17.7%), and service and sales occupations (14.0%). On average, each user conducts 1.5 search sessions during the experiment period. When we compare the occupational distribution of users in our sample with employment in Switzerland, we find that the experiment captures users across a broad range of occupations in proportions broadly related to their prevalence in the labor market.⁹

Appendix Tables A3, A4, and A5 report balance tests separately for the March, April, and May waves. Across observable user and pre-experiment activity measures, treatment arms are very similar within each wave. The joint tests generally do not reject equality across arms, consistent with successful random assignment.

2.1.5 Quality of the Information Provided

Because the displayed information comes from employee reviews, its accuracy is not guaranteed. Panel (b) of Figure 2 compares the median wages displayed in the experiment, aggregated to the two-digit ISCO occupation level, with official median wages in the same occupations, derived from the Swiss Earnings Structure Surveys 2018–2022. The relationship is close to one-for-one across low- and high-paying occupations, which suggests that kununu wages do not systematically over- or underestimate official wages. Appendix Figure A3 provides additional validation by comparing the distribution of kununu wages with official wage statistics by gender.

Validation of the benefit data is more limited because no official data source provides vacancy- or firm-level measures of benefit availability. For work from home, however, we can compare the benefit information displayed in the experiment with the vacancy text, since this is one of the few job benefits sometimes reported directly in job ads. Appendix Table A6 cross-tabulates whether the vacancy text mentions work from home with whether the experimental information identifies work from home as available at the posting firm. Since the absence of a statement in the vacancy text does not imply that work from home is unavailable, this comparison is most informative about false negatives. The comparison suggests that the 50% threshold used in a few treatment arms may be too restrictive. Under this threshold, among ads whose text mentions work from home and for which benefit information is available, the false-negative rate is 64.0%. The share of all

⁷The full sample reflects a set of sample restrictions. We exclude heavy users and potential bots using predefined activity criteria. We also omit sessions in which a user was exposed to two or more distinct treatments, a scenario occurring in fewer than 0.05% of cases. Appendix C details the data-cleaning steps and sample restrictions.

⁸We assign an occupation to a session when more than half of the impressions in that session belong to the same one-digit ISCO occupation. Using this rule, we can assign an occupation to about 74% of sessions.

⁹Appendix Figure A2 plots the number of assigned users in our sample against employment by two-digit ISCO occupation, as reported by the Federal Statistical Office. The relationship is strongly positive.

ads for which the experimental information contradicts a positive statement in the vacancy text is 2.6%. Under the 20% threshold used in most treatment arms, the corresponding false-negative rate is only 17.2%.

2.2 Experiment 2: Belief-Elicitation Survey Experiment

Experiment 1 identifies job seekers’ behavioral responses to displayed wage and benefit information, but it does not directly reveal the belief changes underlying these responses. Experiment 2 complements the field evidence by eliciting job seekers’ beliefs about wages and benefits, as well as how they update them in response to information.

2.2.1 Experimental Design

The follow-up survey invited job seekers recruited on Jobchannel’s platforms to participate in a short online study after providing informed consent. Respondents first reported the occupation and Swiss region in which they were typically searching for work, using broad occupational fields based on three-digit ISCO categories and NUTS-2 regions. We used these answers to show participants a personalized set of real job ads drawn from the field-experiment sample. This personalization is a distinguishing feature of the survey design. The key question when interpreting the behavior in our field experiment is not whether job seekers can distinguish higher- from lower-paying jobs across the labor market as a whole. Rather, it is whether they can infer a job’s pay from the firm name, location, and job title shown in an ad that falls within their typical search region and occupation. The survey mirrors this setting by showing respondents vacancies from their selected occupation–region cell, closely resembling the vacancies they encounter on the platform.

Participants were then randomly assigned to a control group or to one of several treatment groups. The control group saw result-list previews of ads as they normally appear on the platform. The treatment groups saw previews with wage information, benefit information, or both, displayed in the same format as in the field experiment (see Figure 1). The benefit-only arm mirrors the situation for the ads in the wage-and-benefit arms that lack kununu wage information and thus display only benefit information. One treatment group also received an explicit statement that the wage and benefit information originated from kununu.

Before answering any belief questions, all participants viewed a short preview of three real ads to familiarize themselves with the display format. They then evaluated eight personalized job ads drawn from the experimental field sample.¹⁰ For each of the eight jobs, respondents in all groups stated their beliefs about the monthly wage they would earn if hired into the position. They also reported beliefs about the availability of four specific benefits: flexible working hours, work from home, firm-sponsored childcare, and a company car. In addition, they reported one further outcome, randomly assigned at the respondent level: their perceived probability of being hired, their assessment of the workplace atmosphere, or the expected wage five years after entry. Treated respondents who saw wage or benefit information also indicated where

¹⁰We selected ads from the field-experiment sample by forming 15 occupational groups from the 20 three-digit ISCO occupations with the most impressions across ads with kununu wage and benefit information, merging closely related occupations. We then crossed these groups with the five predominantly German-speaking NUTS-2 regions. Within each occupation–region cell, we randomly drew up to eight ads with wage information from distinct firms. Impressions of ads in these occupation–region cells account for 40.7% of the impressions across all matched ads in the field-experiment analysis sample and 60.4% of the impressions across ads with matched wage information. Appendix Table A8 compares the occupational distribution of field-experiment users with that of the survey sample and reports the balance of occupations across the survey arms. Appendix Table A7 shows that the survey ads are broadly similar to the field-experiment ads in terms of average wages and benefit availability.

they believed this information came from (the employer, an employer-review platform such as kununu, or other sources).

After completing the belief-elicitation tasks, participants rated the relevance of the shown ads and whether they thought job platforms should display wage and benefit information. The survey concluded with an attention check, a demographic module, three questions on how respondents interpret the kununu benefit icons, and an employment and wage-history module. See Appendix G for the full survey script.

2.2.2 Sample Description and Balance

Our survey analysis sample consists of 1,484 respondents who completed the full survey and passed the attention check. Because the field experiment observes users through platform cookies and contains only limited demographic information, the survey provides useful evidence on the characteristics of Jobchannel users who participate in the belief-elicitation experiment.

Table A9 reports descriptive statistics for respondents and tests balance across treatment arms. Respondents are on average 43 years old, and 50.4% are women. The majority live in Switzerland. About 12.7% have at least one child below age six. The sample is relatively highly educated: 52.2% report tertiary education, and slightly less than half are employed at the time of the survey. The average respondent has 19.3 years of work experience in Switzerland. Among respondents reporting a current or previous wage, the average full-time monthly wage is approximately CHF 7,100, close to the Swiss median full-time wage of CHF 7,024 in 2024 according to the Swiss Earnings Structure Survey. The table also suggests that random assignment produced similar respondent characteristics across treatment arms. The joint balance tests generally do not reject equality across arms.

3 Conceptual Framework

In this section, we introduce a simple conceptual framework for understanding how wage and benefit information in vacancies affects job search. Consider a job seeker i who browses a set of job advertisements on a job platform and decides whether to apply to a given vacancy j . The platform does not display information about offered wages or benefits, so the individual makes this decision based on her priors about the vacancy’s characteristics. The propensity of individual i to apply to vacancy j is given by

$$y_{ij} = \varphi_j + \beta_w \log w_{ij}^{\text{prior}} + \sum_k \beta_b^k b_{ij}^{k,\text{prior}} + \epsilon_{ij}, \quad (1)$$

where $\log w_{ij}^{\text{prior}}$ denotes the expected log wage offered by vacancy j , and $b_{ij}^{k,\text{prior}}$ is the prior belief on the availability of benefit k . The term φ_j is a vacancy fixed effect that captures vacancy-specific determinants of application behavior unobserved by the econometrician, while ϵ_{ij} represents idiosyncratic factors affecting individual i ’s propensity to apply to vacancy j . The parameter β_w measures the semi-elasticity of applications with respect to the own wage, and is one of the key parameters of interest in our study. In classic monopsony models of differentiated jobs, this parameter is closely related to the firm labor supply elasticity—the key determinant of firms’ monopsony power [Azar and Marinescu, 2024]. When β_w is low, workers are relatively unresponsive to wage differences, allowing firms to exercise greater monopsony power and sustain larger wage markdowns. Conversely, a high β_w indicates that workers are highly responsive to wages, reducing

employers' wage-setting power. The other key parameters of interest are β_b^k , which capture the importance of each benefit k in a job seeker's application decision. The ratio β_b^k/β_w represents the willingness to pay (WTP) for benefit k . It measures the wage a worker would be willing to forgo to receive benefit k .

In the experiment, job seekers are shown information about the wage level, $\log w_j$, and the availability of benefits, b_j , associated with vacancy j . A key challenge is that job seekers may not fully update their beliefs in response to the information provided. Individuals may rationally discount the sometimes noisy benchmark wages provided, may not fully trust the information displayed, anchor their expectations around their current pay [Altmann et al., 2026, Andresen et al., 2026, Jäger et al., 2024], or believe that the ultimately negotiated wage could differ from the posted wage. To capture these possibilities, we assume that job seekers update their prior beliefs about wages and benefits according to

$$\log w_{ij}^{\text{post}} = \alpha_w \log w_{ij} + (1 - \alpha_w) \log w_{ij}^{\text{prior}}, \quad (2)$$

$$b_{ij}^{k,\text{post}} = \alpha_b b_{ij}^k + (1 - \alpha_b) b_{ij}^{k,\text{prior}}, \quad (3)$$

where α_w and α_b govern how much wage and benefit information is incorporated into posterior beliefs. When they are equal to one, job seekers fully update their beliefs to the wage and benefit information. This framework assumes that job seekers update their beliefs only about the job attributes for which information is explicitly provided, and do not revise their beliefs about other attributes. In our setting, such belief spillovers are less likely because information about different benefits is provided simultaneously alongside wage information. In addition, job seekers observe the firm's name and the whole vacancy text, allowing them to form beliefs based on more than the posted wage alone. Nevertheless, in Section 6 we study belief spillovers in detail, but, consistent with our assumption here, do not find evidence for economically meaningful spillovers.

Our experimental design gives rise to two distinct sources of identifying variation, which we can use to estimate job seekers' preferences for wages and benefits.

Approach based on within-ad variation The first approach exploits that the experiment randomly varied the wage and benefit information shown to job seekers for the same vacancy. The application semi-elasticity with respect to posted wages, η_w^{within} , can be estimated by regressing an indicator for applications y_{ij} on the log wage displayed to job seekers in the treatment groups:

$$y_{ij} = \varphi_j + \eta_w^{\text{within}} \log w_{ij} + \sum_k \eta_k^{\text{within}} b_{ij}^k + \varepsilon_{ij}, \quad (4)$$

where $\log w_{ij}$ denotes the wage information shown to individual i for vacancy j , and b_{ij}^k is an indicator equal to one if benefit k is shown as being offered by vacancy j . The term φ_j is a vacancy fixed effect that absorbs all vacancy-specific determinants of application behavior.

Because workers respond to their posterior beliefs rather than directly to the wage information shown, the posted-wage semi-elasticity η_w^{within} does not necessarily identify the structural application semi-elasticity with respect to beliefs about own wages, β_w . As we derive in Appendix D, the semi-elasticity of applications with respect to posted wages from Equation (4)—identified conditional on vacancy fixed effects φ_j and the

vector of benefit information shown $\mathbf{b}_{ij}$ —can be expressed as:

$$\begin{aligned}\eta_w^{\text{within}} &= \frac{\text{Cov}(y_{ij}, \log w_{ij} \mid \mathbf{b}_{ij}, \varphi_j)}{\text{Var}(\log w_{ij} \mid \mathbf{b}_{ij}, \varphi_j)} \\ &= \alpha_w \beta_w.\end{aligned}\tag{5}$$

The within-ad posted-wage semi-elasticity combines two channels: (i) how much higher wages encourage job seekers to apply (β_w), and (ii) how much job seekers update their beliefs in response to the wage information shown (α_w). If job seekers do not update their prior beliefs, the experiment would yield a zero semi-elasticity of applications even though $\beta_w > 0$.

Approach based on across-ad variation The second approach exploits that the experiment revealed information about wages and benefits across different vacancies to treated job seekers, while withholding this information from the control group. For simplicity, the following exposition focuses only on wages. In this case, the econometric specification is:

$$y_{ij} = \varphi_j + \gamma_w T_i^w + \eta_w^{\text{across}} (\log w_j - \overline{\log w}) \times T_i^w + \varepsilon_{ij},\tag{6}$$

where φ_j is a vacancy fixed effect, T_i^w an indicator equal to one if individual i was shown the wage information and zero otherwise, and $(\log w_j - \overline{\log w})$ is the log wage of vacancy j , centered around the sample mean wage. With this recentring, γ_w measures the effect of revealing the average wage relative to the control group. This coefficient provides insight into job seekers' prior beliefs about the overall wage level. A positive estimate of γ_w implies that revealing the average wage increases applications, suggesting that workers systematically underestimate the prevailing wages. The converse holds if γ_w is negative. The coefficient η_w^{across} measures the extent to which providing wage information increases the responsiveness of applications to wage differences across vacancies. It is identified from comparing the treatment groups' differences in applications to higher-versus lower-paid jobs with the differences in application rates across vacancies of the control group. In Appendix D we formally decompose η_w^{across} into the structural parameters:

$$\begin{aligned}\eta_w^{\text{across}} &= \frac{\text{Cov}(\log w_j, \mathbb{E}[y_{ij} \mid j, T_i^w = 1] - \mathbb{E}[y_{ij} \mid j, T_i^w = 0])}{\text{Var}(\log w_j)} \\ &= \alpha_w \beta_w \left(1 - \frac{\text{Cov}(\log w_j, \log w_{ij}^{\text{prior}})}{\text{Var}(\log w_j)} \right) = \alpha_w \beta_w (1 - \theta_w).\end{aligned}\tag{7}$$

Thus, when using the across-ad variation, the posted-wage semi-elasticity depends on three parameters: (i) the preferences for wages β_w , (ii) updating α_w , and (iii) the extent of prior knowledge θ_w . The prior knowledge parameter θ_w is defined as $\text{Cov}(\log w_j, \log w_{ij}^{\text{prior}}) / \text{Var}(\log w_j)$ and measures how much the prior beliefs are aligned with actual wages. If $\theta_w = 1$, job seekers' priors covary one-for-one with wages. A value of $\theta_w = 0$ indicates that priors are uninformative. Intuitively, if job seekers are perfectly informed about wage differences across vacancies, the across-ad information treatment does not have any bite.

One testable prediction of this framework is that, if $\theta_w > 0$, the application semi-elasticity using the within-ad variation η_w^{within} should be larger than the semi-elasticity estimated from the across-ad variation

η_w^{across} . In fact, one can back out θ_w from relating the two semi-elasticities: $\theta_w = 1 - \frac{\eta_w^{\text{across}}}{\eta_w^{\text{within}}}$.

4 Wages

Next, we turn to how job seekers react to posted wages in vacancies. We first leverage the field experiment to estimate the revealed-preference effect of wage posting on job seekers' search behavior. We then use the follow-up survey experiment to decompose the estimates of η into their components: (i) belief updating (α), (ii) wage misperceptions ($1 - \theta$), and (iii) the own-wage elasticity of job applications (β).

4.1 Econometric Approach

The conceptual framework highlights two sources of wage variation that the field experiment can exploit to causally identify the posted-wage effect. In both approaches, we restrict the sample to ad impressions with kununu benefit information. Ads without matched wage information remain in the sample. While these ads do not help identify the wage coefficients, retaining them helps estimate other regression coefficients, including the benefit coefficients and the session fixed effects.

Within-ad variation The first source of variation is the within-ad variation in wages displayed on the same job ad. We focus on job seekers in the mean- and median-wage treatment arms (i.e., we drop the control group) and estimate the following empirical counterpart of Equation (4):

$$y_{s(i)j} = \eta_w^{\text{within}} \widetilde{\log(w_{s(i)j})} \times T_{s(i)j}^w + \sum_{k=1}^K \eta_k^{\text{within}} b_{s(i)j}^k \times T_{s(i)}^k + \varphi_j + \chi_s + \pi_r + \varepsilon_{s(i)j}, \quad (8)$$

where $y_{s(i)j}$ denotes one of our two binary outcome measures of job-search behavior by job seeker i in session s on vacancy j : viewing the ad or taking an ad action, both conditional on an ad impression. The variable $\widetilde{\log(w_{s(i)j})}$ is the log of the wage displayed for vacancy j . This is either the median wage or the mean wage, depending on the treatment arm, both recentered around the sample mean log wage (the log of CHF 6,185). The posted-wage indicator, $T_{s(i)j}^w$, equals one if job seeker i is shown wage information for vacancy j in session s , and zero otherwise. We account for the benefit information provided to job seekers by controlling for an ad-level dummy $b_{s(i)j}^k$, which equals one if a benefit k is displayed as available at the firm posting vacancy j . We interact this variable with a session-level indicator, $T_{s(i)}^k$, which equals one if session s belongs to the treatment arm that displays information about benefit k . The associated benefit effects are discussed in Section 5.1.¹¹ The terms φ_j , χ_s , and π_r denote vacancy, session, and rank (result-list position) fixed effects, respectively, included for the reasons discussed below. $\varepsilon_{s(i)j}$ is an error term.

The coefficient of interest is η_w^{within} . It captures the semi-elasticity of ad views and ad actions with respect to the posted wage. Because we focus on the wage treatment arms and control for vacancy fixed effects φ_j , we identify the coefficient only from ads with experimental variation in wages. Specifically, we compare the decisions of job seekers who saw the mean wage with those who saw the median wage for the same job ad.

¹¹Since we restrict the sample to vacancies with benefit information, whether benefit information is displayed depends only on the session-specific treatment arm. Thus, the indicator $T_{s(i)}^k$ carries no vacancy j index. In contrast, $b_{s(i)j}^k$ varies between ads in a session because firms differ in the benefits they offer. The posted-wage indicator, $T_{s(i)j}^w$, is vacancy j -specific, too, because many ads with benefit information had no matched wage information.

Across-ad variation Our second empirical approach contrasts job seekers’ behavior in the information treatment arms with that in the control group. The approach mirrors a continuous difference-in-differences (DiD) design. We compare the difference in click rates for job seekers viewing higher- versus lower-wage jobs in the treatment group and then subtract the difference for job seekers viewing the same ads in the control group. The sample is the mean-wage treatment arms and the control group and thus abstracts from the within-ad variation by dropping the median-wage treatment arms. The empirical counterpart of Equation (6) is

$$y_{s(i)j} = \gamma_w T_{s(i)j}^w + \eta_w^{\text{across}} \widetilde{\log(w_{s(i)j})} \times T_{s(i)j}^w + \sum_{k=1}^K \eta_k^{\text{across}} b_{s(i)j}^k \times T_{s(i)j}^k + \varphi_j + \chi_s + \pi_r + \varepsilon_{s(i)j} \quad (9)$$

The outcome $y_{s(i)j}$, the recentered displayed wage $\widetilde{\log(w_{s(i)j})}$, and the wage-treatment indicator $T_{s(i)j}^w$ are defined as in Equation (8). One difference is that the regression contains the uninteracted posted-wage indicator $T_{s(i)j}^w$. It enters because $T_{s(i)j}^w$ now equals zero not only for ads without kununu wage information but also for all impressions in the control group. Hence, the variable is no longer absorbed by the vacancy fixed effects. The associated coefficient γ_w captures the effect of displaying a wage equal to the sample mean relative to providing no wage information. Interpreted through the lens of our model (see Section 3), its sign indicates whether job seekers, on average, under- or overestimate prevailing wages.

The key element in Equation (9) is the vacancy fixed effects φ_j . They absorb all observed and unobserved vacancy characteristics that do not vary across treatment arms—the vacancy text, job title, firm identity, or job location. They ensure a causal interpretation of η_w^{across} by accounting for the possibility that ads that are more attractive in observed and unobserved dimensions might pay higher wages. Indeed, the possibility to control for vacancy fixed effects distinguishes our across-ad estimates from studies that estimated regressions similar to Equation (9) using observational data.

Our across- and within-ad specifications also control for session fixed effects χ_s and rank fixed effects π_r . The rank fixed effects account for the influence of an ad’s position in the result list on job seekers’ interest and mainly serve to increase precision. The session fixed effects imply that the DiD design focuses on wage differences across ads that appeared in the same session and thus in the same treatment arm. They also hold the search context constant, including the job seeker’s characteristics and preferences, as well as the characteristics of other job ads displayed simultaneously. We include session fixed effects for two reasons. First, they account for potential treatment effects on job seekers’ search behavior within a session, such as the total number of ads viewed or screened in detail. Second, they restrict attention to ads that job seekers most likely compared directly. This reduces the risk of comparing clicks that were based on different decision criteria. These criteria may change across sessions, for instance, due to changes in preferences or search scopes during the job search episode.

In principle, however, the session fixed effects are not required for identification. Because treatment arms are randomly assigned prior to exposure, the job seekers’ initial search queries and the platform’s resulting ad rankings are balanced in expectation. The displayed information could in principle affect scrolling or later queries. Robustness checks such as controlling for previously viewed wages and focusing on the 20 top-ranked ads per search indicate, however, that possible treatment effects on search behavior do not materially bias the main estimates. Consistent with this, omitting the session fixed effects yields quantitatively similar but

less precise estimates of γ_w and η_w^{across} .

4.2 Posted Wage Elasticities

Table 2 uses the analysis sample described in Section 2.1.4 to estimate the causal effect of posted wages on job seekers’ interest in job ads. Columns 1–4 use ad views as the outcome and Columns 5–6 use the composite indicator of ad actions, i.e., printing, sharing, saving the ad, applying through the platform, and navigating to the employer’s website. The unit of observation is an ad impression. Columns 1, 3, and 5 present estimates from the across-ad specification (Equation (9)). Columns 2, 4, and 6 report results from the within-ad specification (Equation (8)). Columns 3–4 further restrict the sample to ads for which kununu wage information is available and drop the session fixed effects. Throughout, we cluster standard errors at the user-by-month level, the level at which treatment assignment varies.

The table provides several important insights. First, ad views increase with posted wages. In our baseline across-ad specification (Column 1), the semi-elasticity of ad views with respect to log posted wages is 0.015. To facilitate comparison with previous research, we express our estimates henceforth as elasticities by dividing each coefficient by the mean of the outcome.¹² Given that job seekers view 3.86% of the ads shown in their result lists, the coefficient implies an elasticity of ad views with respect to posted wages of $\hat{\eta}_w^{\text{across}} = 0.393$.¹³ The elasticity directly measures the extent to which wage information shifts views toward higher-paying vacancies. Extrapolating our posted-wage elasticity implies that displaying wage information on all vacancies would increase ad views of vacancies in the top quintile of the wage distribution by approximately 30% relative to those in the bottom quintile (see Appendix Figure A8).

Second, the posted wage elasticity when exploiting within-ad variation is approximately twice as large as the corresponding across-ad elasticities, with implied elasticities of $\hat{\eta}_w^{\text{within}} = 0.83$ (Column 2) and 0.6 (Column 4). As shown in Section 3, comparing the two elasticities reveals how well informed job seekers are about pay differences across jobs when no wage is posted. In particular, the two baseline estimates suggest that $\hat{\theta}_w = 1 - (\hat{\eta}_w^{\text{across}}/\hat{\eta}_w^{\text{within}}) = 1 - 0.015/0.033 \approx 0.54$, which implies that a 10% wage gap between two jobs in the kununu data translates into a 5.4% perceived gap among job seekers without wage information.

Third, the estimated across- and within-ad elasticities are marginally smaller in Columns 3 and 4, which drop the session fixed effects. The coefficients are less precisely estimated, consistent with the role of session fixed effects discussed in Section 4.1: by comparing only ads plausibly evaluated against one another, they absorb differences in click propensities across job seekers and search contexts.

Fourth, the coefficients on the wage-shown indicator ($T_{s(i)j}^w$) are positive and sizable: displaying a wage equal to the sample mean strongly increases job seekers’ propensity to engage with an ad, relative to providing no wage information. Through the lens of our framework, the positive estimates indicate that job seekers without posted wage information underestimate the pay at the firms for which we could match kununu wage information on average.

¹²We use the control-group mean in the across-ad specifications and the estimation-sample mean in the within-ad specifications, whose sample excludes the control group.

¹³In a slight abuse of notation, we use η_w for this elasticity in what follows, whereas in Section 3 it denotes the regression coefficient. Since the rescaling applies to all outcome-side parameters, including β_w and β_b^k , while the belief parameters θ and α are unit-free, the identities linking η , α , θ , and β carry over, as do ratios of parameters such as the WTP. The one qualification is that the across-ad and within-ad elasticities are scaled by slightly different means, which is why we compute $\hat{\theta}_w$ below from the ratio of the coefficients rather than of the elasticities.

Table 2: Effect of Posted Wages on Job Seekers’ Ad Views and Ad Actions

	Ad views		Ad views, wage-info ads		Ad actions	
	(1) Across	(2) Within	(3) Across	(4) Within	(5) Across	(6) Within
$\log(w_j^k) \times T_{s(i)j}^w$	.015*** (.0016)	.033*** (.009)	.013*** (.0025)	.026*** (.0093)	.0014*** (.00045)	.0037 (.003)
$T_{s(i)j}^w$	.0089*** (.00049)		.0071*** (.00078)		.0012*** (.00016)	
Control mean	.038599	.038896	.03511	.038896	.003697	.003745
Posted wage elasticity	.393	.836	.373	.608	.383	.902
Model	LPM	LPM	LPM	LPM	LPM	LPM
Ad fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Rank fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Session fixed effects	Yes	Yes	No	No	Yes	Yes
Observations	10,565,784	10,921,672	2,145,045	2,192,893	10,565,784	10,921,672

Notes: This table reports estimates of the effect of posted wages (in logs, recentered around the sample mean) on job seekers’ interest in a job advertisement. The outcome is an indicator for viewing the ad (Columns 1–4) or for taking any ad action, i.e., printing, sharing, saving, applying through the platform, or navigating to the employer’s website (Columns 5–6). The unit of observation is an ad impression. The sample is the analysis sample defined in Section 2.1.4. Columns 1, 3, and 5 report the across-ad specification (Equation (9)), which uses the control group and all mean-wage treatment arms. Columns 2, 4, and 6 report the within-ad specification (Equation (8)), which uses the mean- and median-wage treatment arms. Columns 3–4 restrict the sample to ads with kununu wage information and drop the session fixed effects. The posted-wage elasticity at the bottom of the table divides the log-wage coefficient by the mean outcome. For the across-ad specification, the mean is computed in the control group. For the within-ad specification, it is computed in the estimation sample, which excludes the control group. Standard errors are clustered at the user-by-month level.

Fifth, the estimated wage elasticities are very similar when using ad actions as the outcome (Columns 5 and 6). The effect of higher posted wages on the probability of viewing an ad, therefore, translates one-to-one into a higher probability of taking an ad action. This pattern speaks against concerns that prominently displaying wages at the result-list stage prompts job seekers to click on high-wage ads, a decision they subsequently reverse after seeing the full job information. Indeed, we directly test for such a reversal by estimating the effect on ad actions conditional on viewing the full vacancy details. The wage effects are insignificant in this specification, which speaks against such a reversal (see Appendix Table A10).

The estimates of the posted-wage elasticities η in Table 2 fall within the typical range of the growing experimental and non-experimental literature on how job seekers respond to wages on online job boards. Observational studies of job-board data [e.g., Banfi and Villena-Roldan, 2019, Escudero et al., 2025, Marinescu and Wolthoff, 2020] and field experiments that randomize wages on real or fictitious vacancies [Abebe et al., 2021, Balgova et al., 2025, Belot et al., 2022, Dal Bó et al., 2013, He et al., 2023] typically find posted-wage elasticities of applicant interest or applications in the 0.5–1.0 range.

Robustness. The across-ad estimates (η_w^{across}) are robust to a range of further sensitivity and specification checks, presented in Appendix Tables A11 and A12. Specifically, Column 2 of Table A11 shows that the wage estimates are similar if we drop the ad fixed effects and instead control for the ad information visible at the search-results stage. The specification yields more precise estimates, though, because it compares

decisions across observably similar ads rather than focusing solely on decisions taken on the same ad.¹⁴ Column 3 shows similar wage estimates using a multinomial logit discrete-choice model, interpreting the job ads displayed in the same session as the job seeker’s choice set. We implement it via fixed-effects Poisson Pseudo-Maximum Likelihood (PPML), which yields coefficients identical to those of the conditional logit [Baker, 1994, Guimaraes, 2004] but remains computationally tractable even with our large number of fixed effects. The wage estimates are also virtually unchanged (but less precise) when we replace the session fixed effects with user-by-month or experiment-arm fixed effects (Columns 4 and 5). Finally, controlling for the wages job seekers already observed on prior impressions in the same session does not affect the wage estimates (Column 6).¹⁵

A related concern is that the displayed information affects how far job seekers scroll or whether they start further searches, so that later impressions are no longer balanced across treatment arms. Column 1 of Appendix Table A12 addresses this concern by restricting the sample to the 20 top-ranked ads per search, which are least affected by such responses. The wage coefficient is, if anything, somewhat larger than in the baseline, while the implied elasticity is somewhat smaller.

The other columns of Table A12 show that the estimated wage effects are similar when we exclude short sessions (in which users show little engagement with the displayed ads) or restrict the sample to full-time jobs (for which the displayed full-time wage corresponds to the position’s monthly gross pay). They are also similar when we restrict the sample to job seekers who were always assigned to the same treatment arm (for whom exposure to multiple arms cannot attenuate the experimental estimates). The wage elasticity is also unchanged in the April treatment arm, in which we displayed the wage to the right of the benefit icons rather than to the left.

Appendix Table A12 also probes the estimates’ sensitivity to the wage range considered. Column 5 excludes ads with posted wages more than a third above or below the median of their two-digit ISCO occupation. The across-ad posted-wage elasticity is essentially unchanged (0.38). Column 6 instead excludes ads in the bottom and top quintiles of the wage distribution. Within this central range, the elasticity is somewhat larger (around 0.53). This estimate provides an alternative basis to compute θ_w , since the within-ad elasticity is also identified from comparable wage ranges. Relating the two coefficients, the implied awareness parameter is somewhat smaller: $\hat{\theta}_w = 1 - 0.021/0.033 \approx 0.36$, compared with 0.54 in the full sample.

Directed Search. Our framework assumes that job seekers’ interest in a job increases linearly with its wage. This assumption may be violated if workers interpret higher wages as a signal of greater job competitiveness, as predicted by directed search models. In such models, job seekers apply for higher-paying jobs only up to the point where increased competition offsets the wage gain. Recent experimental evidence supports this mechanism [Belot et al., 2022, He et al., 2023].

As a first test of this mechanism, Appendix Figure A8 reports the effect of wages on ad views non-parametrically, using indicators for each quintile of the displayed wage distribution. The wage effect rises approximately linearly across the lower quintiles but flattens at the top quintile. This is the pattern the

¹⁴The close agreement between Columns 1 and 2 also suggests that, in this setting, a selection-on-observables strategy based on these controls would deliver a wage elasticity close to the experimental benchmark.

¹⁵These prior wages shape clicking behavior. However, treatment assignment does not affect which ads appear in a search or their ordering. The benchmark wages underlying the ads of job seekers’ earlier impressions are therefore approximately balanced across treatment arms and largely uncorrelated with our wage variables of interest.

competition channel would generate if higher wages no longer translate into higher interest at the top. The same pattern, however, is also consistent with other mechanisms,¹⁶ including random search if job seekers' reservation wages cluster around the upper end of the wage distribution. In the latter case, the plateau arises because a high posted wage exceeds nearly all job seekers' reservation wages.

To assess the competition channel directly, we draw on the follow-up belief-elicitation survey, which elicited respondents' perceived probability of being hired when applying to a vacancy (Appendix Table A15). We then regress this probability on the kununu wage, an indicator for whether the respondent was in a wage treatment arm, and their interaction. Consistent with Belot et al. [2022], we find a negative coefficient on log wage, suggesting that job seekers in the control group perceive higher-paying jobs as harder to get. However, the interaction with the wage-information treatment is not statistically significant and positive. Perceived hiring chances still decline with the wage among job seekers who receive wage information, but the decline is, if anything, smaller than in the control group. Higher posted wages therefore do not appear to strengthen the perceived competition for high-wage jobs.¹⁷ These results speak against the concern that our field-experimental posted-wage elasticities are attenuated because job seekers infer stronger competition from higher posted wages.

Heterogeneity. Figure 3 illustrates heterogeneity in the posted wage elasticity of ad views across occupations. Panel (a) reports estimates by broad occupational group (one-digit ISCO occupations). The figure suggests that the posted wage elasticity declines with the occupational skill level: it is highest among elementary occupations, plant and machine operators, and assemblers, and lowest among professionals, technicians, and associate professionals.

The upper part of Panel (b) suggests that the posted wage elasticity declines with the occupational wage level: it is highest in the lowest-paying tercile of occupations and lowest in the highest-paying tercile. The lower part of Panel (b) shows that the posted wage elasticity is statistically significantly higher in occupations with higher female employment shares. The gender-related heterogeneity aligns with Jalal [2026], who finds that pay disclosure raises women's applications by more than men's. It is also consistent with our survey-experimental evidence that women update wage beliefs more strongly to wage information than men.¹⁸ The pattern differs from studies that find little gender heterogeneity in the wage elasticity of applications or recruitment among individual job seekers [Balgova et al., 2025, Dal Bó et al., 2013].

4.3 Beliefs and Prior Updating


Our theoretical framework implies that the posted-wage elasticities η combine job seekers' own-wage elasticity of job applications β_w , their misperception of pay differences across jobs in the absence of wage information $1 - \theta_w$, and their belief updating when given wage information α_w . Based on the comparison of the across-ad and within-ad specifications, our preferred estimate of θ_w is 0.54. In this section, we use the belief-elicitation survey to estimate α_w and to provide an independent estimate of θ_w .

¹⁶The plateau could arise from diminishing marginal value of further wage increases at the top or from skepticism about implausibly high posted wages.

¹⁷These results are consistent with Christensen et al. [2026], who likewise find limited evidence that wage information affects beliefs about job-level competition.

¹⁸Appendix Table A19 reports how the components of η_w^{across} vary across respondents. Awareness of pay differences is similar by gender and across terciles of respondents' own wage ($\hat{\theta}_w$ between 0.41 and 0.48). Belief updating differs more: women update about twice as strongly as men ($\hat{\alpha}_w = 0.40$ versus 0.21).

Figure 3: Heterogeneity in the Posted Wage Elasticity of Ad Views by Occupation



Notes: Panel (a) shows the posted wage elasticity of ad views by broad occupational group (one-digit ISCO occupation). Panel (b) shows estimates by tertile of the occupational wage level (upper part) and by occupations with low, medium, and high shares of female workers (lower part). Across all specifications, we exploit only across-ad variation by excluding treatment arms with median wages. We calculate wage elasticities by dividing the wage coefficient by the average ad view rate in the control group. Whiskers denote 95% confidence intervals, and standard errors are clustered at the user-by-month level. The occupational wage level is the median wage in the three-digit ISCO occupation, and the female employment share is the share of women in employment in the three-digit ISCO occupation, both computed from the pooled Swiss Earnings Structure Survey (Lohnstrukturhebung) waves 2018, 2020, and 2022.


We start with a descriptive analysis of wage beliefs. Panels (a) and (b) of Figure 4 plot the distribution of the gap between a respondent's expected own pay in a job and the job's matched kununu wage, winsorized at $\pm 5,000$ CHF. We use the survey sample described in Section 2.2. In the control group (Panel (a)), priors are widely dispersed around the benchmark and, on average, fall below the mean full-time kununu wage by 377 CHF (approximately 5%). Panel (b) shows the same distribution for respondents who were informed about the kununu wage. Compared to the control group, providing wage information shifts beliefs toward the kununu benchmark and sharply reduces their dispersion. It also essentially closes the mean gap. Individual beliefs, however, still frequently depart from the displayed wage, indicating that job seekers update wage beliefs only partially. Panel (c) offers a complementary perspective by plotting job seekers' own-wage beliefs against the kununu wage. The binscatter for the control group indicates that job seekers without wage information can, on average, distinguish between high- and low-paying jobs. However, their beliefs lie below the benchmark wage for high-wage vacancies and above it for low-wage ones. Providing wage information (red line) steepens the relationship between the two wages and shifts beliefs toward the benchmark. However, the slope remains below the slope of one that full updating would imply.

Table 3 formalizes this visual relationship by reporting estimates from the following regression:

$$\widetilde{\log(w_{ij}^{\text{post}})} = \iota + \theta_w \widetilde{\log(w_j)} + \gamma_w T_i^w + \sigma_w \widetilde{\log(w_j)} \times T_i^w + \rho X_{ij} + \varepsilon_{ij}, \quad (10)$$

where $\widetilde{\log(w_{ij}^{\text{post}})}$ denotes respondent i 's posterior expected own log wage in job j , $\widetilde{\log(w_j)}$ is the corresponding kununu wage, and T_i^w indicates assignment to a wage-information treatment arm. Both log wages are recentered around the field-experiment mean log wage (the log of CHF 6,185). X_{ij} denotes user- and vacancy-level controls. θ_w measures the slope between wage beliefs and actual wages in the control group, and thus maps directly into the structural parameter θ_w from the conceptual framework. The coefficient

Figure 4: Belief-Elicitation Survey: Descriptive Evidence on Wage Misperceptions and Belief Updating



Notes: The figure describes respondents' wage beliefs and updating in response to information using data from the belief-elicitation survey. The sample is the analysis sample described in Section 2.2. The unit of observation is a respondent–ad pair. Wage beliefs are respondents' expected own monthly full-time wage in the job shown. The kununu wage is the matched job-level benchmark. Panels (a) and (b) show the distribution of the difference between the two, winsorized at $\pm 5,000$ CHF, for the control group and the wage treatment groups, respectively. The vertical line marks the mean difference. Panel (c) shows a binned scatter plot of own-wage beliefs against the kununu wage, both in logs, for the control group (blue) and the wage treatment groups (red), with linear fits. The gray line is the 45-degree line.

γ_w in (10) captures whether providing wage information shifts job seekers' wage expectations for a job with average pay. The interaction coefficient σ_w measures the extent to which posted kununu wages change the relationship between wage beliefs and the kununu wage. Using the belief-updating rule from Equation (2), we derive in Appendix D that

$$\sigma_w = \alpha_w(1 - \theta_w). \quad (11)$$

Intuitively, σ_w is a mix of the updating parameter α_w and the degree of prior misperception $1 - \theta_w$: when job seekers are perfectly informed about wage differences, the treatment cannot steepen the belief–wage relationship even under full updating. It follows that the belief-updating parameter can be recovered as $\hat{\alpha}_w = \hat{\sigma}_w / (1 - \hat{\theta}_w)$. It measures the share of the initial misperception that is corrected when wage information is provided.

Table 3: Estimates of Wage Misperceptions and Belief Updating

	(1)	(2)	(3)	(4)	(5)
$\widetilde{\log(w_j)}$	0.645*** (0.0252)	0.456*** (0.0209)	0.436*** (0.0211)		0.406*** (0.0224)
T_i^w	0.0447*** (0.0114)	0.0450*** (0.0107)	0.0441*** (0.0102)	0.0434*** (0.0105)	0.0206 (0.0143)
$\widetilde{\log(w_j)} \times T_i^w$	0.164*** (0.0380)	0.156*** (0.0363)	0.206*** (0.0369)	0.205*** (0.0376)	0.252*** (0.0442)
$\widetilde{\log(lastwage_i)}$			0.162*** (0.0226)	0.161*** (0.0231)	
$\widetilde{\log(lastwage_i)} \times T_i^w$			-0.122*** (0.0294)	-0.121*** (0.0300)	
Occupation*Region FE		Yes	Yes		Yes
Job FE				Yes	
Control + kununu info sample					Yes
$\hat{\alpha}_w$	0.463 (0.089)	0.287 (0.059)	0.365 (0.057)		0.424 (0.067)
Observations	11872	11872	11872	11848	6880
R-squared	0.283	0.327	0.351	0.476	0.315

Notes: The table shows regression estimates of Equation (10) based on data from the belief-elicitation survey at the respondent-ad level. The dependent variable is the respondent's expected own log wage in a job. $\widetilde{\log(w_j)}$ is the average kununu log wage for the job title and firm posting vacancy j . Both are recentered around the mean log wage in the field experiment. T_i^w equals one if job seeker i was assigned to a treatment arm with wage information. $\widetilde{\log(lastwage_i)}$ is the current (employed) or most recent (unemployed) monthly full-time wage. Columns 3 and 4 additionally include the recentered past pay and its interaction with T_i^w . Column 5 restricts the sample to the control group and the treatment arm that was explicitly informed that the wage information comes from kununu. The belief-updating parameter is $\hat{\alpha}_w = \hat{\sigma}_w / (1 - \hat{\theta}_w)$. Standard errors (in parentheses) are clustered by respondent. The constant is omitted.

Table 3 reports the results from estimating Equation (10). Because the survey respondents need not be representative of the job seekers in the field experiment, Appendix Table A17 also presents the same specifications after reweighting the survey sample to match the occupational distribution of field-experiment impressions. We reweight both to all impressions and to impressions that displayed wage information. The reweighted estimates are comparable to the unweighted ones, which is why we focus on the unweighted estimates in the following.

We begin with the estimates of θ_w . Column 1 suggests that a one-log-point increase in the kununu wage is associated with an increase of $\hat{\theta}_w = 0.65$ in the control group's wage priors. However, Column 1 largely captures whether job seekers can infer pay differences across occupations and regions. As discussed in Section 2.2, what matters for interpreting our field experiment is instead whether job seekers can identify a job's expected pay within their search scope. Our preferred specifications, presented in the remaining columns, therefore include occupation-by-region fixed effects and compare jobs only within occupation-and-region cells.

Column 2 indicates that control-group job seekers are relatively good at reading relative pay differences from job-ad previews in their search scope: a one-log-point increase in the kununu wage is associated with an increase of $\hat{\theta}_w = 0.46$ in expected pay beliefs. Our direct estimates of θ_w are thus remarkably similar to

the indirect values of θ_w implied by the comparison of posted-wage elasticities in the field experiment. The results also align with recent survey-experimental estimates that workers are broadly aware of pay differences across jobs.¹⁹

Next, we turn to the baseline effect of providing wage information on beliefs about the own wage. The coefficient on the wage-information indicator T_i^w implies that job seekers update wage beliefs upward when shown an average posted wage. Specifically, the estimate suggests that they underestimate full-time monthly wages in the survey jobs by 4.5% on average (Column 2). These results are consistent with our field-experimental finding that posting an average wage increases job seekers’ interest in job ads, which also indicates that job seekers underestimate average full-time wages (cf. Table 2).

Finally, we turn to the interaction term that measures whether job seekers update their wage beliefs in response to the posted wage. Indeed, in the wage-treatment arms, the relationship between wage beliefs and the kununu wage rises by a statistically significant $\hat{\sigma}_w = 0.156$, from $\hat{\theta}_w = 0.456$ to $\hat{\theta}_w + \hat{\sigma}_w = 0.61$. This implies that the belief-updating parameter is $\hat{\alpha}_w = \hat{\sigma}_w / (1 - \hat{\theta}_w) \approx 0.29$, which suggests that wage information closes about 29% of the gap between job seekers’ priors and the kununu benchmark. These findings confirm that job seekers update beliefs about their own pay meaningfully but incompletely to posted wages.

As discussed in Section 3, there are many reasons for incomplete updating. One prominent potential mechanism is that job seekers anchor their beliefs about their outside options to their current pay [Altmann et al., 2026, Andresen et al., 2026, Jäger et al., 2024]. Consistent with such anchoring, Appendix Table A14 shows that a respondent’s most recent wage is the primary source of heterogeneity in pay beliefs in the control group. Columns 3 and 4 of Table 3 examine how this relationship changes once job seekers observe benchmark wages, by interacting past pay with the wage-information treatment. In the control group, a 10% higher past wage is associated with 1.6% higher expected pay, conditional on the kununu wage. Among job seekers who receive wage information, this association falls by about three quarters. The wage benchmark thus appears to largely replace past pay as a guide to what a job pays. In addition, the result suggests that anchoring to past pay can account for only a limited part of the incomplete updating that we document.

Our evidence that workers update wage beliefs only partially to wage information aligns with recent survey and field-experimental literature documenting incomplete belief updating in response to different types of wage information, including high-quality wage benchmarks from administrative data [Altmann et al., 2026, Andresen et al., 2026, Bhatiya et al., 2026, Christensen et al., 2026]. Indeed, our estimate lies toward the upper end of those belief-updating estimates, which range from 0.15 to 0.35, suggesting that job seekers do not discount the fact that the wage information comes from an employer-review platform. Consistent with this, Column 5 of Table 3 shows that restricting the treatment sample to the arm explicitly informed about the source of the wage data leads, if anything, to higher estimates for α_w . The belief-updating parameter is also similar among survey respondents who rated the ads shown to them as relevant and among respondents who believed the wage information comes directly from the employer rather than from an employer-review platform (see Appendix Table A17).

¹⁹In a survey of full-time German workers, Caldwell et al. [2025] find a correlation between expected and actual firm pay premia of 0.42. In a survey of Danish early-career job seekers, Christensen et al. [2026] similarly estimate a regression slope of 0.54 for perceived versus actual wage gaps, essentially identical to our two estimates of θ . In contrast, Jäger et al. [2024] estimate a slope of only 0.107 between employed workers’ perceived and likely wage changes upon switching to a potential employer. Their results suggest that workers are quite badly informed about wage differences across potential outside options.

4.4 Wage Elasticities

Together, the results in the previous section show that job seekers have informative but imperfect priors about wages and update them in response to wage information, consistent with partial Bayesian learning. We can now combine our estimates of α_w and θ_w with the posted-wage elasticities, $\hat{\eta}_w^{\text{across}}$ and $\hat{\eta}_w^{\text{within}}$, to quantify job seekers' own-wage elasticity of job applications, β . Table 4 reports the results using the survey estimates $\hat{\theta}_w = 0.456$ and $\hat{\alpha}_w = 0.287$ from Column 2 of Table 3.

The first estimate divides the across-ad posted-wage elasticity by the product of the updating parameter and the misperception term, $\hat{\beta}_w^{\text{across}} = \hat{\eta}_w^{\text{across}} / [\hat{\alpha}_w(1 - \hat{\theta}_w)]$, yielding $\hat{\beta}_w = 2.51$ (95% CI: 1.53 to 4.13). The second divides the within-ad posted-wage elasticity by the updating parameter, $\hat{\beta}_w^{\text{within}} = \hat{\eta}_w^{\text{within}} / \hat{\alpha}_w$, yielding $\hat{\beta}_w = 2.91$ (95% CI: 1.49 to 5.67). The two routes deliver mutually consistent estimates that are several times larger than the posted-wage elasticities. Even the lower bounds of the confidence intervals are roughly two to four times as large as the posted-wage elasticities.

A potential concern is that the belief parameters are estimated from survey respondents, who may differ from the job seekers whose search behavior identifies the posted-wage elasticities in the field experiment. To assess this concern, Table 4 also reports $\hat{\beta}_w$ when we estimate the belief parameters after reweighting the survey sample to the occupational distribution of field-experiment impressions with wage information. The reweighting yields a somewhat larger updating parameter and correspondingly smaller implied wage elasticities of 1.8 and 2.1. Nevertheless, even those estimates remain substantially above the posted-wage elasticities.

Overall, our preferred estimates suggest own-wage elasticities of job applications on the order of 2.5–3. This range aligns closely with complementary evidence from Christensen et al. [2026], who estimate an own-wage application elasticity of approximately 3. These estimates suggest that job-seeker preferences are an important source of firms' wage-setting power. In addition, the gap between the own-wage and the posted-wage elasticities also suggests that, even if workers value wages highly, firms may retain wage-setting power when workers do not fully perceive wage differences across jobs. These results support a small but growing literature that studies the role of imperfect information about wages in shaping labor market power [Cheremukhin and Restrepo-Echavarria, 2025, Christensen et al., 2026].

5 Job Benefits

In this section, we turn to job seekers' valuation of job benefits. As with wages, the analysis proceeds in three steps. Section 5.1 uses the field experiment data to quantify how the availability of job benefits shapes job seekers' interest in jobs. Section 5.2 leverages the belief-elicitation survey to assess misperceptions about benefit availability at firms and to estimate updating in response to benefit information. Section 5.3 combines these ingredients to estimate job seekers' WTP for the twelve benefits studied in the paper.

5.1 Effects of Benefits on Job Search

Econometric Approach. We measure the effect of benefits on job seekers' viewing and clicking behavior using the across-ad specification of Equation (9). The coefficients of interest, η_k^{across} , are on the interactions between the benefit-treatment indicators $T_{s(i)}^k$ and the firm-benefit indicators $b_{s(i)j}^k$. $T_{s(i)}^k$ indicates whether

Table 4: Implied Own-Wage Elasticity of Job Applications $\hat{\beta}_w$

	(1) Across-ad	(2) Within-ad
Posted-wage elasticity $\hat{\eta}_w$	0.393 (0.040)	0.836 (0.226)
Belief-update factor $\hat{\sigma}_w$	0.156 (0.036)	
Belief-update factor $\hat{\alpha}_w$		0.287 (0.059)
Own-wage elasticity $\hat{\beta}_w$	2.51 (0.64)	2.91 (0.99)
95% CI	[1.53, 4.13]	[1.49, 5.67]
$\hat{\beta}_w$, survey reweighted to wage-information impressions	1.76	2.10

Notes: The table reports the own-wage elasticity of job seekers' ad views, $\hat{\beta}_w$, obtained by correcting the posted-wage elasticity for prior misperception and belief updating as estimated in the belief-elicitation survey. Column 1 divides the across-ad posted-wage elasticity $\hat{\eta}_w^{\text{across}}$ (Column 1 of Table 2) by $\hat{\sigma}_w = \hat{\alpha}_w(1 - \hat{\theta}_w)$. Column 2 divides the within-ad posted-wage elasticity $\hat{\eta}_w^{\text{within}}$ (Column 2 of Table 2) by $\hat{\alpha}_w = \hat{\sigma}_w / (1 - \hat{\theta}_w)$. Both belief parameters come from Column 2 of Table 3. Standard errors (in parentheses) are computed by the delta method, treating the two experiments as independent. The 95% confidence intervals are constructed on the log scale and back-transformed. The last row reports $\hat{\beta}_w$ when the belief parameters come instead from the specification that reweights the survey sample to the occupational distribution of field-experiment impressions with wage information ($\hat{\sigma}_w = 0.223$, $\hat{\alpha}_w = 0.399$, Appendix Table A17).

job seeker i was informed about benefit k in session s , and $b_{s(i)j}^k$ equals one if benefit k is available at the firm posting vacancy j . They capture the effect of informing job seekers that benefit k is available at the posting firm, relative to informing them that it is *not*.


As in the wage analysis, the treatment main effect $T_{s(i)}^k$ is absorbed by the session fixed effects. The coefficient η_k^{across} is identified from across-ad variation in benefit availability: the difference in viewing rates between ads with and without benefit k among treated job seekers, net of the corresponding difference among controls. We restrict the sample to benefit-treatment arms that classified a benefit as available at a firm if at least 20% of kununu reviewers confirmed it. For the three benefits where some arms instead used a 50% threshold—flexible working hours, work from home, and firm-sponsored childcare—we additionally report a within-ad estimate, η_k^{within} , based on Equation (8). This specification drops observations from the control group and exploits only the variation on the same ad generated by the two thresholds.

Main Results. Figure 5 reports the estimated effects of the twelve job benefits on the likelihood that job seekers view an ad. The corresponding across-ad estimates for ad actions, which are less precise but quantitatively similar, are reported in Appendix Figure A9.²⁰

The estimates suggest that some job benefits shape job seekers' interest in jobs. The largest effect on ad views is observed for the provision of a company car, which raises job seekers' probability of viewing an ad by 7.3%. The opportunity to work from home (+6.0%), firm-sponsored childcare (+4.4%), and a parking spot (+3.1%) also generate economically and statistically significant increases in ad views. These effects remain statistically significant after applying Holm's correction for multiple hypothesis testing across the twelve

²⁰Appendix Table A13 also shows that the estimated benefit effects are highly stable across a range of alternative regression specifications.

Figure 5: Effect of Job Benefits on Job Seekers' Ad Views



Notes: The figure reports the estimated effects of the twelve job benefits on the likelihood that a job seeker views an ad. The outcome is an indicator for viewing the ad. The unit of observation is an ad impression. Across-ad estimates (red) are obtained from Equation (9) and identified from variation in benefit availability between ads, restricting the sample to benefit-treatment arms that classified a benefit as available at a firm if at least 20% of its kununu reviewers mentioned it. Within-ad estimates (blue) are obtained from Equation (8) for the three benefits with experimentally varied availability thresholds (20% vs. 50%). All specifications include ad, rank, and session fixed effects. Coefficients are scaled by the control-group mean of the ad-view rate (0.039) and thus measure percentage changes relative to that baseline. Whiskers denote 95% confidence intervals. Standard errors are clustered at the user-by-month level. Holm-adjusted two-sided p -values, which account for the twelve across-ad hypotheses in the figure, are below 0.001 for the company car and work from home, 0.002 for childcare, 0.024 for the company doctor, and 0.053 for the parking spot. The underlying coefficients are reported in Column 1 of Appendix Table A13.


benefits, although the parking-spot effect is significant only at the 10% level (see the adjusted p -values in the notes to Figure 5). The point estimates for coaching opportunities, food allowances, flexibility, public transport connections, and health measures are positive but not precise enough to rule out zero effects at conventional levels. In contrast, a canteen, employee events, and a company doctor are associated with lower viewing rates. For the company doctor, this reduction is statistically significant. A likely explanation is that a company doctor is perceived as a disamenity, signaling hazardous or otherwise less desirable working conditions.

For the three benefits with within-ad variation, Figure 5 also shows that the within-ad and across-ad estimates are broadly similar. Through the lens of the conceptual framework, the close agreement is consistent with a small θ_b . This suggests that job seekers have relatively little prior knowledge about which firms in their search scope offer which benefits, a tentative result that the belief-elicitation survey confirms.

Heterogeneity. Figure 6 explores heterogeneity in the benefit effects across three occupational dimensions. Panel (a) shows that the benefit effects on ad views are, with few exceptions, similar across terciles of the occupational wage distribution. These results run counter to the notion that lower-paid workers react less to benefit information than higher-paid workers.

Panel (b) reports estimates split by the female employment share in the occupation. The clearest evidence

Figure 6: Heterogeneity of the Effect of Job Benefits on Ad Views



Notes: The figure reports heterogeneity in the estimated effects of job benefits on the likelihood that a job seeker views an ad. The unit of observation is an ad impression. All estimates are obtained from Equation (9), restricting the sample to benefit-treatment arms that classified a benefit as available at a firm if at least 20% of its kununu reviewers mentioned it. Specifications include ad, rank, and session fixed effects. Coefficients are scaled by the control-group mean of the ad-view rate, computed within the respective subsample in Panels (a) and (b) and over the full estimation sample in Panel (c). Panel (a) splits the estimates by tercile of the occupational wage level (median wage in the three-digit ISCO occupation). Panel (b) splits them by the female employment share in the three-digit ISCO occupation. See the notes to Figure 3 for the underlying data sources. Panel (c) splits the estimates by how common each benefit is in an occupation, measured as the share of kununu reviewers confirming the benefit among firms posting vacancies in an occupation, averaged across vacancies in a two-digit ISCO occupation. Occupations are classified as “common” (blue) or “not common” (red) relative to the unweighted mean across occupations. Whiskers denote 95% confidence intervals, and standard errors are clustered at the user-by-month level.

for heterogeneity arises for the opportunity to work from home, which has a significantly larger effect on ad views in female-dominated occupations. The effect rises from about 4% in occupations with a low share of female workers to about 13% in female-dominated occupations. This pattern is consistent with prior work suggesting that women may place a higher value on workplace flexibility and other non-wage job attributes [Burbano et al., 2024, Escudero et al., 2025, Le Barbanchon et al., 2021, Mas and Pallais, 2017, Wiswall and Zafar, 2018]. However, most benefits show little systematic evidence for gender-related heterogeneity, including other aspects of workplace flexibility such as flexible working hours and childcare.

Panel (c) splits the sample by how commonly each benefit is offered within an occupation. The figure also serves as a robustness check for benefits that may be available only to a subset of workers within a firm. While a canteen or a parking spot is typically an amenity offered to all workers within a firm, others, such as work from home, may be available only to some workers. The firm-level benefit effect may thus mask heterogeneity between workers who profit from it and workers who do not.²¹ Indeed, we find that the work-from-home effect is larger in occupations where work from home is more common. We observe a similar pattern for food allowances, which increase job seekers’ interest in occupations where they are common. For the remaining benefits, the differences between the two subsamples are small, or the effect is larger when the benefit is uncommon. Overall, except for work from home and possibly food allowances, the results broadly confirm our approach of treating benefits as firm-level amenities.

5.2 Benefit Misperceptions and Belief Updating

The conceptual framework implies that the across-ad benefit estimates from the previous section reflect a combination of job seekers’ valuation of job benefits (β_b^k), prior misperceptions about benefit availability at firms ($1 - \theta_b^k$), and belief updating to benefit information (α_b^k). As with wages, we now leverage the belief-elicitation survey to recover θ_b^k and α_b^k based on the following regression:

$$b_{ij}^{\text{post},k} = \theta_b^k b_j^k + \gamma_b^k T_i^b + \sigma_b^k b_j^k \times T_i^b + \varepsilon_{ij}. \quad (12)$$

Here, $b_{ij}^{\text{post},k}$ is respondent i ’s belief about whether firm j offers benefit k , where k is one of the three benefits covered by the survey module—flexible working hours (flx_j), work from home (wfh_j), and firm-sponsored childcare ($chld_j$). b_j^k is the corresponding indicator of whether the firm offers benefit k according to kununu. T_i^b equals one if survey respondent i was assigned to a benefit-information treatment arm. It follows that θ_b^k identifies the control group’s awareness of which firms offer benefit k . γ_b^k captures the change in beliefs when respondents learn that a firm does not offer benefit k , and σ_b^k measures the change when they learn that the firm *does* offer it relative to learning that it does not offer it. The belief-updating parameter is $\hat{\alpha}_b^k = \hat{\sigma}_b^k / (1 - \hat{\theta}_b^k)$. We estimate Equation (12) with and without additional controls for the other two benefits shown, the wage-information treatment, and the kununu wage.

Table 5 reports our baseline estimates of Equation (12). We focus on θ_b first. Column 1 shows that respondents in the control group rate the probability that a firm offers flexible working hours about 7.7 percentage points higher ($\hat{\theta}_b = 0.077$) when the firm actually offers it according to kununu than when it does not. The corresponding gap for work from home (Column 3) is somewhat larger but still modest, with

²¹Ideally, for such a robustness check, we would split occupations into those where a benefit (e.g., work from home) is essentially absent and those where a non-negligible share of workers receive it. However, our employer review data do not record reviewers’ occupations, so we cannot directly observe the within-firm distribution of benefit availability across occupations.

Table 5: Misperception about Benefits and Belief Updating

	(1) flx_{ij}^{post}	(2) flx_{ij}^{post}	(3) wfh_{ij}^{post}	(4) wfh_{ij}^{post}	(5) $chld_{ij}^{\text{post}}$	(6) $chld_{ij}^{\text{post}}$	(7) b_{ij}^{post}	(8) b_{ij}^{post}
T_i^b	-0.0323 (0.0245)	-0.0140 (0.0260)	-0.0851*** (0.0157)	-0.00884 (0.0226)	-0.148*** (0.0184)	-0.140*** (0.0287)	-0.110*** (0.0142)	-0.110*** (0.0152)
flx_j	0.0767*** (0.0156)	0.0381** (0.0161)		0.0530*** (0.0154)		0.0389*** (0.0141)		
$flx_j \times T_i^b$	0.215*** (0.0266)	0.233*** (0.0279)		-0.104*** (0.0225)		-0.0125 (0.0238)		
wfh_j		0.0854*** (0.0117)	0.148*** (0.0106)	0.108*** (0.0112)		-0.0157 (0.0108)		
$wfh_j \times T_i^b$		-0.0480*** (0.0181)	0.240*** (0.0199)	0.270*** (0.0213)		-0.0132 (0.0171)		
$chld_j$		-0.00472 (0.0126)		0.0294** (0.0116)	0.138*** (0.0125)	0.128*** (0.0128)		
$chld_j \times T_i^b$		-0.00248 (0.0200)		-0.00916 (0.0173)	0.286*** (0.0244)	0.293*** (0.0246)		
b_j							0.0892*** (0.00864)	0.0796*** (0.00868)
$b_j \times T_i^b$							0.275*** (0.0182)	0.276*** (0.0183)
$\widetilde{\log(w_j)}$		0.172*** (0.0318)		0.321*** (0.0286)		0.187*** (0.0310)		0.248*** (0.0208)
T_i^w		-0.0158 (0.0169)		-0.00693 (0.0125)		0.0206 (0.0192)		-0.000624 (0.0117)
$\widetilde{\log(w_j)} \times T_i^w$		-0.00923 (0.0468)		-0.0867** (0.0366)		-0.0213 (0.0540)		-0.0488 (0.0323)
Occ*Region FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
$\hat{\alpha}_b$	0.233 (0.027)	0.242 (0.027)	0.282 (0.022)	0.303 (0.022)	0.332 (0.026)	0.336 (0.026)	0.302 (0.019)	0.300 (0.019)
Observations	11872	11872	11872	11872	11872	11872	35616	35616
R-squared	0.119	0.127	0.278	0.289	0.0976	0.102	0.121	0.127

Notes: The table shows regression estimates of Equation (12) based on data from the belief-elicitation survey. The dependent variables are respondents' beliefs about the availability of flexible working hours (flx_{ij}^{post} , Columns 1–2), work from home (wfh_{ij}^{post} , Columns 3–4), and firm-sponsored childcare ($chld_{ij}^{\text{post}}$, Columns 5–6). Columns 7–8 pool the observations for the three benefits in a single regression (b_{ij}^{post}). T_i^w and T_i^b equal one if job seeker i was assigned to a treatment arm with wage or benefit information, respectively. flx_j , wfh_j , and $chld_j$ indicate the availability of the corresponding benefit at the firm posting vacancy j according to kununu. $\widetilde{\log(w_j)}$ is the average kununu log wage for the job title and firm posting vacancy j , recentered around the mean log wage in the field experiment. Columns 2, 4, and 6 additionally control for the availability of the other two benefits and their interactions with T_i^b , and for T_i^w , the kununu wage, and their interaction. Column 8 controls for the wage variables only. The belief-updating parameter is $\hat{\alpha}_b = \hat{\sigma}_b / (1 - \hat{\theta}_b)$. Standard errors (in parentheses) are clustered by respondent.

$\hat{\theta}_b = 0.148$. The estimate for childcare (Column 5) is $\hat{\theta}_b = 0.138$. These values are several times smaller than the wage-awareness estimate $\hat{\theta}_w = 0.46$ in Table 3, indicating that job seekers are considerably less informed about benefit availability than about pay. The result that θ_b is low is also consistent with the indirect evidence in Figure 5 that the across-ad and within-ad benefit estimates are similar. It is further consistent with the fact that survey respondents rate the usefulness of benefit information on job advertisements as highly as that of wage information (see Appendix Figure A12).

Next, we turn to belief updating. The interaction coefficient $\hat{\sigma}_b$ is positive and statistically significant for all benefits, and the implied updating parameter $\hat{\alpha}_b$ is quite homogeneous across them, falling between 0.23 and 0.33. These values are similar to the wage-updating parameter $\hat{\alpha}_w = 0.29$, suggesting that respondents internalize benefit information at a rate comparable to that for posted wage information.

The close agreement of $\hat{\theta}_b$ and $\hat{\alpha}_b$ across the three benefits motivates estimating a single θ_b and α_b for the three benefits by pooling responses on all of them. The pooled estimates, reported in Column 7, are $\hat{\theta}_b = 0.089$ and $\hat{\alpha}_b = 0.302$.

As for wages, the updating estimates are robust. They are similar when we add the controls in the even columns of Table 5 and when we restrict attention to respondents who found the ads that we showed to them relevant or who believed the information came from the employer. They are also similar when we reweight the survey sample to the occupational distribution of field-experiment impressions (see Appendix Table A18).²²

5.3 Willingness-to-Pay Estimates

We now have all the ingredients to quantify how much of a wage increase job seekers would forgo to obtain a specific benefit—their WTP for benefit k . We maintain the assumption throughout—supported by the comparable estimates of $\hat{\theta}_b$ and $\hat{\alpha}_b$ across the three benefits tested in Table 5—that the belief-updating and prior-knowledge parameters are common across the twelve benefits considered. Solving the expressions for η_w^{across} and η_k^{across} to recover the underlying preferences and taking the ratio yields

$$\text{WTP}^k = \frac{\beta_b^k}{\beta_w} = \frac{\eta_k^{\text{across}}}{\eta_w^{\text{across}}} \cdot \frac{\alpha_w (1 - \theta_w)}{\alpha_b (1 - \theta_b)}. \quad (13)$$

We recover $\hat{\alpha}_w$ and $\hat{\theta}_w$ from Column 2 of Table 3 and $\hat{\alpha}_b$ and $\hat{\theta}_b$ from Column 7 of Table 5. Standard errors for $\widehat{\text{WTP}}^k$ are computed by the delta method. Appendix E gives the explicit variance expression and describes how we recover the joint covariances.


Figure 7 presents the resulting WTP estimates and associated 95% confidence intervals for the 12 benefits considered in our experiment.²³ The highest-valued benefit is the provision of a company car, with a WTP of 10.4% of the wage. At the field-experiment mean wage of CHF 6,185, this amounts to roughly CHF 640 per month—an estimate that is broadly consistent with the cost of privately leasing a mid-size car in Switzerland.²⁴

²²In contrast, the estimates of θ_b fall when we control for the availability of the other benefits and the wage information. This is because control-group beliefs about one benefit correlate with the availability of the other benefits and the kununu wage. Job seekers' priors thus partly reflect a firm's general benefit generosity rather than knowledge about the specific benefit. Thus, the baseline estimates of θ_b may overstate benefit-specific knowledge.

²³This interprets the log-wage equivalent as a direct percentage.

²⁴All-inclusive leasing rates, including insurance, tax, and maintenance, amount to approximately CHF 500–700 per month

Figure 7: Job Seekers' Willingness to Pay for Job Benefits



Notes: The figure reports point estimates and 95% confidence intervals of the willingness to pay (WTP) for each of the 12 benefits, expressed as a percentage of the wage. The WTP is computed using the across-ad design as $\widehat{WTP}^k = (\hat{\eta}_k^{\text{across}} / \hat{\eta}_w^{\text{across}}) \cdot (\hat{\sigma}_w / \hat{\sigma}_b)$ per Equation (13). The first ratio is identified from the field experiment (Equation (9)), and the second ratio is identified from the belief-elicitation survey (Table 3 for $\hat{\sigma}_w$ and Column 7 of Table 5 for $\hat{\sigma}_b$). Confidence intervals are computed by the delta method following Equation (25), as detailed in Appendix E.

The second-highest-valued benefit is the opportunity to work from home, with a WTP of 8.3% of the wage (about CHF 510 per month). Based on the responses to the belief-elicitation survey, we should interpret this estimate as the value of the option to work from home for part of the week.²⁵ The estimate lies well within the 4–25% range reported in a growing literature that has measured the WTP for work from home using a variety of approaches [Andresen et al., 2026, Audoly et al., 2026, Colonnelli et al., 2023, Cordes and Müller, 2026, Cullen et al., 2025, Drake et al., 2025, Maestas et al., 2023, Mas and Pallais, 2017].

The third-most-valued benefit is firm-sponsored childcare, with a WTP of 6.1% of the wage. The belief-elicitation survey responses suggest that this estimate should be interpreted as an approximately 50% childcare subsidy rather than fully employer-financed or on-site care.²⁶ A back-of-the-envelope calculation suggests that this WTP is approximately consistent with such a subsidy.²⁷

A parking spot at the workplace is the fourth benefit with a positive and statistically significant WTP of 4.4%. Most other benefits have positive but statistically insignificant WTPs, including flexible working

for compact-to-midsize models such as a VW Golf or Skoda Octavia.

²⁵In the belief-elicitation survey, we showed respondents the kununu home-office icon and asked how many days they would expect to work from home in a full-time job. Of the 1,621 respondents, 2.3% expected no home office, 6.9% one day, 28.6% two days, 21.2% three days, 6.3% four days, and 21.2% a full five days, with the remaining 13.5% giving another answer.

²⁶We presented survey respondents the kununu childcare icon and asked what share of the costs of external childcare they would expect the employer to cover. Among the 1,361 respondents who gave a numerical answer, the mean expected subsidy was 50%, with a modal response of 50%.

²⁷The OECD estimates that net childcare costs for a two-earner couple with two children amount to 35% of average earnings in Switzerland [OECD, 2022]. A perceived 50% subsidy therefore implies a benefit worth 17.5% of earnings. Since 32.9% of respondents are aged 25–40, our proxy for those potentially benefiting, the population-weighted value is 5.8% of wages, close to our WTP estimate of 6.1%.

hours. This null result is consistent with Mas and Pallais [2017], who find that a majority of workers do not value schedule flexibility, but it contrasts with the sizable WTP in Wiswall and Zafar [2018] and Maestas et al. [2023]. A potential explanation for the null result is that a substantial minority of survey respondents read the icon as indicating employer discretion over working time rather than as freedom to set one’s own schedule. Employer discretion is a work arrangement that workers typically dislike [e.g., Mas and Pallais, 2017].²⁸

6 Are Wages a Signal for Other Job Attributes?

In many settings, wage information is more readily available than benefit information.²⁹ Do job seekers use available wage information to update their perceptions of other job aspects in such cases? This possibility has long been recognized as a potential source of bias in the stated-preference literature and in settings such as ours [Adrián et al., 2026]. If information on one job attribute spills over to beliefs about other attributes, WTP estimates capture not only the value of the attribute itself but also the value of changes in beliefs about these other attributes.

To tackle this question, we adapt our conceptual framework from Section 3 to allow job seekers to update their posterior beliefs about benefits in case only wage information is shown. Formally, the belief updating is now summarized with the following two equations:

$$\log w_{ij}^{\text{post}} = \alpha_w \log w_{ij} + (1 - \alpha_w) \log w_{ij}^{\text{prior}}, \quad (14)$$

$$b_{ij}^{\text{post}} = b_{ij}^{\text{prior}} + \alpha_{wb}(\log w_{ij} - \log w_{ij}^{\text{prior}}). \quad (15)$$

As before, the posterior wage belief ($\log w_{ij}^{\text{post}}$) is a convex combination of the job seeker’s prior belief ($\log w_{ij}^{\text{prior}}$) and the displayed wage ($\log w_{ij}$). But benefit beliefs are now updated in response to wage information rather than benefit information. The parameter α_{wb} captures these belief spillovers: a displayed wage above the job seeker’s prior raises the perceived likelihood that the benefit is available if $\alpha_{wb} > 0$ and lowers it if $\alpha_{wb} < 0$.

Our experimental setup allows us to quantify the extent of spillovers, α_{wb} , by leveraging the wage-only treatment arm in the survey experiment. Although the group did not receive information on benefits, we elicited beliefs on the availability of flexible working hours, work from home, firm-sponsored childcare, and a company car. In addition, we asked a random subset of respondents about their beliefs regarding work atmosphere and the probability of getting the job. We estimate whether information about wages affects benefit beliefs using the following regression:

$$b_{ij}^{\text{post},k} = \iota + \kappa^k \widehat{\log(w_j)} + \gamma_w^k T_i^w + \sigma_w^k \widehat{\log(w_j)} \times T_i^w + \gamma^k b_j^k + \varepsilon_{ij}. \quad (16)$$

Here, κ^k captures the relationship between beliefs about benefit k and wages in the absence of wage information, while σ_w^k captures how this relationship changes when individuals receive wage information. Appendix F

²⁸When asked to interpret the kununu icon in the belief-elicitation survey, 71.0% of the 1,621 respondents read it as “Employees are free to choose their own working hours,” while 20.7% read it as “The employer expects a high degree of flexibility from employees.”

²⁹For example, many US states and the 2026 EU Pay Transparency Directive mandate greater transparency regarding salaries, but not for benefits. In addition, many job board platforms have started displaying wage estimates on vacancies.

Table 6: Testing for Spillovers of Wage Information on Benefit Beliefs

	(1) flx_{ij}^{post}	(2) flx_{ij}^{post}	(3) wfh_{ij}^{post}	(4) wfh_{ij}^{post}	(5) $chld_{ij}^{post}$	(6) $chld_{ij}^{post}$	(7) car_{ij}^{post}	(8) car_{ij}^{post}	(9) b_{ij}^{post}	(10) atm_{ij}^{post}
$\log(\widetilde{w_j})$	0.236*** (0.0368)		0.289*** (0.0345)		0.230*** (0.0353)		0.253*** (0.0314)		0.242*** (0.0210)	
T_i^w	-0.0180 (0.0225)	-0.0193 (0.0232)	-0.00916 (0.0167)	-0.0105 (0.0172)	0.0385 (0.0282)	0.0381 (0.0292)	0.0317* (0.0170)	0.0321* (0.0176)	0.0107 (0.0141)	-0.0900 (0.122)
$\log(\widetilde{w_j}) \times T_i^w$	-0.0218 (0.0673)	-0.00418 (0.0689)	-0.0186 (0.0539)	0.00309 (0.0531)	-0.0331 (0.0867)	-0.0305 (0.0913)	0.0644 (0.0587)	0.0555 (0.0598)	-0.00183 (0.0429)	0.513 (0.370)
flx_j	0.0648*** (0.0150)									
wfh_j			0.111*** (0.0109)							
$chld_j$					0.128*** (0.0119)					
car_j							0.155*** (0.0128)			
b_j									0.152*** (0.00755)	
Occ*Region FE	Yes		Yes		Yes		Yes		Yes	
Job FE		Yes		Yes		Yes		Yes		Yes
Observations	6960	6917	6960	6917	6960	6917	6960	6917	27840	2294
R-squared	0.0979	0.215	0.288	0.469	0.0733	0.128	0.189	0.306	0.0901	0.213

Notes: The table shows results of estimating Equation (16) on data from the control and wage-only treatment arms of the belief elicitation survey. The dependent variables are respondents' beliefs about the probability that the firm posting the vacancy offers flexible working hours (flx_{ij}^{post} , Columns 1–2), work from home (wfh_{ij}^{post} , Columns 3–4), firm-sponsored childcare ($chld_{ij}^{post}$, Columns 5–6), and a company car (car_{ij}^{post} , Columns 7–8). Column 9 pools the observations for the four benefits in a single regression (b_{ij}^{post}). The dependent variable in Column 10 (atm_{ij}^{post}) is the respondents' belief about the workplace atmosphere in the firm posting the vacancy, measured on a scale from 0 to 7. car_j indicates the kununu-based availability of a company car at the firm posting vacancy j . All other variables are defined as in Tables 3 and 5. Standard errors (in parentheses) are clustered by respondent.

shows that, under the conceptual framework with the new updating rule, σ_w identifies:

$$\sigma_w = \alpha_{wb}(1 - \theta_w). \quad (17)$$

Thus, σ_w captures the strength of belief updating α_{wb} attenuated by a term that reflects individuals' prior knowledge about wage differences across jobs. Intuitively, wage information has little effect on benefit beliefs when job seekers already have accurate beliefs about wage differences across jobs. Since θ_w is substantially below one (around 0.5), we can test belief spillovers directly with the null hypothesis $\sigma_w = 0$.

Table 6 presents results of estimating Equation (16) for different benefits. The coefficient on the log wage, presented in the first row, shows that control-group job seekers expect benefits to be more likely available at firms with higher benchmark wages. Despite this, we find no evidence that wage information meaningfully changed job seekers' beliefs about benefit availability. All interaction terms between the displayed wage and the wage treatment (σ_w) are close to zero and statistically insignificant, even when we pool the benefits to maximize statistical precision (Column 9). In addition, we find no evidence that the posted wage level affected respondents' expected probability of receiving a job offer or their longer-run wage expectations (see Appendix Table A15). Appendix Table A16 further examines whether providing information about specific benefits generates belief spillovers to other benefits or to wages. Consistent with the previous results, we

find little evidence of cross-belief updating.³⁰

Although we cannot exclude the possibility that wages shape perceptions of other, unobserved job attributes, our tests generally fail to detect spillovers from wages or benefits to other job attributes. These findings are consistent with the limited evidence for belief spillovers in Christensen et al. [2026]. They contrast, however, with the belief spillovers from advertised amenities and posted wages documented by Adrjan et al. [2026]. One difference between the settings is that our job seekers take consequential decisions on real vacancies from named firms within their search scope, whereas Adrjan et al. [2026] study choices between hypothetical ads without a named employer. In such a setting, varying a posted wage or amenity may shift beliefs about the identity of the firm behind the ad. Because benefits are firm-level attributes, such a shift may move benefit beliefs more strongly than in our setting, in which the firm remains the same.

7 Implications for Inequality

A central question is whether non-wage benefits compress or amplify inequality across jobs. If higher-paying firms also offer more highly valued benefits, then conventional measures of wage inequality understate the true dispersion in total compensation and may mismeasure inequality between groups of workers. Indeed, models of imperfect competition generate dispersion in overall job value, so that better-paying jobs tend to be better along non-wage dimensions as well [see Mas, 2026]. Conversely, if high wages represent compensating differentials for offering fewer benefits, then focusing on wages alone overstates inequality in overall job value. In either case, adding the value of benefits to the wage recovers the underlying dispersion of job value. In this section, we thus combine our willingness-to-pay estimates with data on wages and benefit availability across jobs to quantify the extent to which non-wage benefits affect inequality among the sample of ads with matched kununu wage information.

To move from wage inequality to job-value inequality, we construct a money-metric measure of total compensation: the sum of a job’s pay and the monetized value of all benefits offered by the firm (both measured using kununu reviews). Specifically, for each job vacancy j with matched wage and benefit information, we define total job value as:

$$JV_j = w_j \left(1 + \sum_{b \in \mathcal{B}_j} \widehat{WTP}_b \right), \quad (18)$$

where w_j denotes the wage associated with vacancy j , $\mathcal{B}_j$ is the set of benefits available at j , and $\widehat{WTP}_b$ is the estimated willingness to pay for benefit b as a share of the wage, as estimated in Section 5.3. The summation includes all twelve benefits studied in the paper, including the few benefits with negative estimated valuations.

Table 7 compares the distribution of wages and total job value across the 36,504 job ads with wage information. First, the average job vacancy offers benefits with a monetized value of approximately 11 percent of wages, or CHF 741. Benefits thus constitute an economically important component of total compensation. Second, the dispersion of total job value exceeds the dispersion of wages across every measure reported. For

³⁰Information on flexible working hours and working from home does not significantly affect beliefs about company cars, wages, or the work atmosphere. Similarly, information on childcare does not significantly affect beliefs about wages or the work atmosphere. The only exception is a statistically significant effect of childcare on beliefs about the availability of a company car.

Table 7: Inequality in Wages and Total Job Value across Job Advertisements

	Mean	SD	P50	P90/P10	P50/P10	Gini	Obs.
Wage (CHF)	6616	2109	6008	2.14	1.34	0.172	36,504
Job Value (CHF)	7357	2456	6725	2.21	1.38	0.182	36,504

Notes: The table reports the distribution of the kununu wage and of total job value across job advertisements in the matched sample with wage information. Total job value equals the wage times one plus the summed willingness to pay, as a share of the wage, for the benefits the job offers, including the few benefits with negative estimated valuations. A benefit counts as available when at least 20% of a firm’s reviewers confirm it. P90/P10 and P50/P10 are ratios of the corresponding percentiles, and Gini is the Gini coefficient.

instance, the standard deviation increases from CHF 2,109 to CHF 2,456—a rise of approximately 16 percent. The Gini coefficient, which is scale-invariant and measures relative inequality, rises from 0.172 to 0.182 (+6 percent). The reason why job-value inequality exceeds wage inequality is that higher-paying jobs offer more benefits. As shown in Appendix Figure A11, the lowest-paying decile of firms offers on average around six benefits, whereas the highest-paying decile offers 8 benefits.

These results speak to a growing literature studying how amenities relate to wages across jobs and firms [e.g., Caldwell et al., 2026, Costas-Fernández et al., 2026, Lamadon et al., 2022, Ouimet and Tate, 2023, Sockin, 2024, Sorkin, 2018]. Our results align with the majority of papers in this literature, which find that accounting for non-wage compensation amplifies rather than compresses measured inequality. Indeed, the estimated relationship between job value and wages is consistent with Maestas et al. [2023], who also value a set of observed amenities. They find that the 90–10 differential in total compensation is about 11% larger than the corresponding wage differential. Our estimates are smaller, though, than those from more structural approaches that infer amenities from job transitions or offer data. Such approaches attribute sometimes more than half of the dispersion in job value to non-wage factors (see the estimates surveyed in Mas, 2026, Section 8.5).

8 Conclusion

This paper combines a large-scale field experiment on one of Switzerland’s largest job-board networks with a belief-elicitation experiment. Together, the two experiments provide new evidence on how job seekers value wages and non-wage benefits, how well informed they are about compensation differences across employers, and how they update their beliefs when provided with new information. The experiments provide several key insights.

First, revealing wage information has a highly significant and robust effect on job search behavior: a 1% increase in the posted wage raises the likelihood that a job seeker engages with an ad by 0.4%.

Second, in both the field and the survey experiment, we find that job seekers are already partially informed about wage differences across jobs in their search scope. Despite widespread pay opacity in Switzerland, control-group job seekers who did not receive wage information correctly identify roughly half of the differences in benchmark wages across ads.

Third, we find that job seekers update only partially toward posted wages, closing about 29% of the gap between their prior beliefs and the benchmark wage. Combining these belief-updating estimates with the field-experimental elasticities, we recover application elasticities with respect to job seekers’ beliefs about

their own pay of 2.5 and 2.9. This implies that job seekers are several times more responsive to actual wage differences than the posted-wage elasticity would suggest.

Fourth, we estimate job seekers’ willingness to pay for job benefits by combining the two experiments. The field experiment delivers the ratio of benefit-to-wage responses in actual search behavior, and the survey corrects for the fact that job seekers are less informed about benefit availability than about wages. Among the twelve benefits studied, a company car and work from home are the highest-valued benefits. Job seekers are willing to forgo 10.4% of their wage for a company car and 8.3% of their wages to work from home. The estimated WTP is 6.1% of wages for firm-sponsored childcare and 4.4% for a parking spot.

Finally, we find that job benefits constitute an important component of total compensation and amplify between-job inequality. In our sample of ads with wage information, total job value exceeds the wage by approximately 11% for the average job. Because higher-paying jobs tend to offer more valued benefits, inequality in job value is significantly larger than inequality in wages alone.

References

- Girum Abebe, A. Stefano Caria, and Esteban Ortiz-Ospina. The selection of talent: Experimental and structural evidence from Ethiopia. *American Economic Review*, 111(6):1757–1806, 2021.
- Pawel Adrjan, Mária Balgová, Simon Jäger, Jonas Jessen, and Jason Sockin. Job ads as signals: Evidence from a priced amenity and worker beliefs. Discussion Paper 184/26, RFBerlin, 2026.
- Steffen Altmann, Robert Mahlstedt, Malte Jacob Rattenborg, Alexander Sebald, Sonja Settele, and Johannes Wohlfart. Wage expectations and job search. CESifo Working Paper 12420, CESifo, Munich, 2026.
- Martin E. Andresen, Manudeep Bhuller, and Alfred Løvgren. Pay beliefs and the amenity-pay tradeoff. arXiv:2606.02503, 2026.
- Richard Audoly, Manudeep Bhuller, and Tore Adam Reiremo. The pay and non-pay content of job ads. Staff Report 1124, Federal Reserve Bank of New York, 2026.
- José Azar and Ioana Marinescu. Monopsony power in the labor market. In *Handbook of Labor Economics*, volume 5, pages 761–827. Elsevier, 2024.
- Michael Baker, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar. Pay transparency and the gender gap. *American Economic Journal: Applied Economics*, 15(2):157–183, 2023.
- Stuart G Baker. The multinomial-poisson transformation. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 43(4):495–504, 1994.
- Maria Balgova, Tsegay Tekleselassie, Lukas Hensel, and Marc J. Witte. Wage information and applicant selection. IZA Discussion Paper 18220, Institute of Labor Economics (IZA), Bonn, Oct 2025.
- Stefano Banfi and Benjamin Villena-Roldan. Do high-wage jobs attract more applicants? directed search evidence from the online labor market. *Journal of Labor Economics*, 37(3):715–746, 2019.
- Michele Belot, Philipp Kircher, and Paul Muller. How wage announcements affect job search—a field experiment. *American Economic Journal: Macroeconomics*, 14(4):1–67, 2022.

- Morten Bennedsen, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon. Do firms respond to gender pay gap transparency? *The Journal of Finance*, 77(4):2051–2091, 2022.
- Morten Bennedsen, Birthe Larsen, and Jiayi Wei. Gender wage transparency and the gender pay gap: A survey. *Journal of Economic Surveys*, 37(5):1743–1777, 2023.
- Apurav Yash Bhatiya, Bhaskar Chakravorty, Clément Imbert, and Roland Rathelot. How do beliefs and preferences over jobs affect enrollment in vocational training: Experimental evidence from India. *European Economic Review*, 184:105238, 2026. doi: 10.1016/j.euroecorev.2025.105238.
- Vanessa Burbano, Nicolas Padilla, and Stephan Meier. Gender differences in preferences for meaning at work. *American Economic Journal: Economic Policy*, 16(3):61–94, 2024. doi: 10.1257/pol.20220121.
- Sydnee Caldwell, Ingrid Haegle, and Jörg Heining. Why workers stay: Pay, beliefs, and attachment. NBER Working Paper 33445, National Bureau of Economic Research, 2025.
- Sydnee Caldwell, Ingrid Haegle, and Jörg Heining. Firm pay, amenities, and inequality. NBER Working Paper 35149, 2026.
- Anton Cheremukhin and Paulina Restrepo-Echavarria. An information-based theory of monopsony power. 2025.
- Asker Nygaard Christensen, Nikolaj Harmon, Sonja Settele, and Daphné Skandalis. Perceived wage differences and early career job search choices. Working paper, University of Copenhagen and University of Cologne, 2026.
- Emanuele Colonnelli, Timothy McQuade, Gabriel Ramos, Thomas Rauter, and Olivia Xiong. Polarizing corporations: Does talent flow to “good” firms? Technical report, National Bureau of Economic Research, 2023.
- Simon Cordes and Max Müller. Do firms know what workers want? CRC TR 224 Discussion Paper 739/2026, 2026.
- Julián Costas-Fernández, Sebastian Findeisen, Anna Raute, and Uta Schönberg. Family-friendly workplace policies. CEPR Discussion Paper 21499, 2026.
- Zoë Cullen. Is pay transparency good? *Journal of Economic Perspectives*, 38(1):153–180, 2024.
- Zoë Cullen, Bobak Pakzad-Hurson, and Ricardo Perez-Truglia. Home sweet home: How much do employees value remote work? In *AEA Papers and Proceedings*, volume 115, pages 276–281. American Economic Association, 2025.
- Ernesto Dal Bó, Frederico Finan, and Martín A Rossi. Strengthening state capabilities: The role of financial incentives in the call to public service. *The Quarterly Journal of Economics*, 128(3):1169–1218, 2013.
- Marshall Drake, Neil Thakral, and Linh T. Tô. Wage differentials and the price of workplace flexibility. Working paper, Boston University, 2025.

- Verónica Escudero, Hannah Liepmann, and Damián Vergara. Directed search, wages, and non-wage amenities: Evidence from an online job board. ILO Working Paper 136, 2025.
- Paulo Guimaraes. Understanding the multinomial-poisson transformation. *The Stata Journal*, 4(3):265–273, 2004.
- Andreas Gulyas, Sebastian Seitz, and Sourav Sinha. Does pay transparency affect the gender wage gap? evidence from austria. *American Economic Journal: Economic Policy*, 15(2):236–255, 2023.
- Haoran He, David Neumark, and Qian Weng. Do workers value flexible jobs? a field experiment. *Journal of Labor Economics*, 39(3):709–738, 2021.
- Haoran He, David Neumark, and Qian Weng. I still haven’t found what i’m looking for: Evidence of directed search from a field experiment. *The Economic Journal*, 133(649):258–280, 2023.
- Simon Jäger, Christopher Roth, Nina Roussille, and Benjamin Schoefer. Worker beliefs about outside options. *The Quarterly Journal of Economics*, 139(3):1505–1556, 2024.
- Amen Jalal. Screening women out? pay transparency in job postings. London School of Economics, Job Market Paper, 2026.
- Thibaut Lamadon, Magne Mogstad, and Bradley Setzler. Imperfect competition, compensating differentials, and rent sharing in the us labor market. *American Economic Review*, 112(1):169–212, 2022.
- Thomas Le Barbanchon, Roland Rathelot, and Alexandra Roulet. Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics*, 136(1):381–426, 2021.
- Nicole Maestas, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger. The value of working conditions in the United States and implications for the structure of wages. *American Economic Review*, 113(7):2007–2047, 2023.
- Alan Manning. Monopsony in labor markets: A review. *ILR Review*, 74(1):3–26, 2021.
- Ioana Marinescu and Ronald Wolthoff. Opening the black box of the matching function: The power of words. *Journal of Labor Economics*, 38(2):535–568, 2020.
- Alexandre Mas. Non-wage amenities. In Christian Dustmann and Thomas Lemieux, editors, *Handbook of Labor Economics*, volume 6, chapter 5, pages 373–446. Elsevier, 2026.
- Alexandre Mas and Amanda Pallais. Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–3759, 2017.
- OECD. PF3.4: Childcare support. OECD Family Database, April 2022. URL https://webfs.oecd.org/els-com/Family_Database/PF3-4-Childcare-support.pdf. Data refer to 2021 for most countries and 2020 for Switzerland. Switzerland figures are based on tariffs in Zürich. Source: OECD Tax and Benefit Models 2021.
- Paige Ouimet and Geoffrey Tate. Firms with benefits? nonwage compensation and implications for firms and labor markets. NBER Working Paper 31463, National Bureau of Economic Research, 2023.

Jason Sockin. Show me the amenity: Are higher-paying firms better all around? SSRN Working Paper No. 3957002, 2024.

Isaac Sorkin. Ranking firms using revealed preference. *The quarterly journal of economics*, 133(3):1331–1393, 2018.

Matthew Wiswall and Basit Zafar. Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics*, 133(1):457–507, 2018.

A Further Figures

Figure A1: Snapshots of kununu Employer Review Forms

(a) Wage Reviews



How fairly are you paid?

Your input counts. Help make salaries more transparent and fair. Your information is of course completely anonymous.

employment

Full time Part time

Salary Frequency

Yearly Monthly

Annual salary - full time

Gross (including other remunerati... Annual full-time gross salary including other remuneration

Your salary will be treated confidentially and displayed separately from your employer rating.

Job title

Search for positions

This is mine **Experience**

current job Ex-job Please choose

(b) Benefits Review

What benefits are offered to employees in the company?

Barrier-free Employer-funded pension Company doctor

Coaching Work cell phone Food allowance

Company car Flexible working hours Health measures

Good transport connection Home office Dog allowed

childcare Internet use canteen

Employee participation Employee events Discounts


parking spot

Would you recommend this employer to a friend?

Yes No


Notes: The figure shows snapshots of the kununu wage and benefit review pages. Workers can choose either a monthly or annual pay frequency to report their wages. Since one or two additional monthly payments are common in Switzerland, reviewers who select the monthly payment are asked about the number of monthly payments in a year. For this study, we divide the annual payment for all wage reviews by 12 to obtain a consistent monthly measure of wages. On the benefit review page, reviewers click the benefits offered to them. The order of the benefits is randomized for each reviewer.

Figure A2: Users in Experiment vs. Employment in Switzerland by Two-Digit ISCO Occupation




Notes: The figure plots the number of users in the analysis sample against employment in Switzerland by two-digit ISCO occupation, as reported by the Federal Statistical Office, together with a fitted regression line. Users are assigned to an occupation if more than half of the job ads shown to them are associated with that occupation.

Figure A3: Distribution of kununu Wages: Comparison to Official Statistics



Notes: The figure compares the distribution of monthly wages reported by kununu reviewers with the official wage distribution in Switzerland, separately for women (left panel) and men (right panel). Blue bars show the distribution of gross monthly wages of full-time employees in 2020 according to the Swiss Earnings Structure Survey (SESS, “Lohnstrukturerhebung”), private and public sectors combined. Red bars show the distribution of monthly wages (in CHF) among kununu reviews of employers in Switzerland submitted between April 2019 and March 2021, computed as the reported annual pay divided by twelve. The sample excludes students and interns as well as reviewers who reported leaving the firm before 2019. Wages are grouped in bins of CHF 2,000 and top-coded above CHF 20,000. Bar labels report the share of each bin in percent, and N below each panel denotes the number of kununu wage observations.

Figure A4: Difference between Log Mean and Log Median kununu Wages of Job Ads



Notes: The figure shows the distribution of the difference between the log mean and the log median kununu wage of job ads in the analysis sample with kununu median and mean wage information (N=25,115). The log difference is winsorized at ± 0.2 . The unit of observation is a job ad. The red vertical line marks the (unwinsorized) average log difference across ads.

Figure A5: Screenshot of Result List for Treatment Group: Desktop Users




Figure A6: Screenshot of Single Ad View for Treatment Group: Desktop Users





Figure A7: Screenshots of the Experiment for Treatment Group: Mobile Users

(a) Result List



(b) Single Ad View






Figure A8: Effect of the Posted Wage on Ad Views by Quintiles of Wages




Notes: The figure shows the effect of wages on the likelihood of viewing an ad by quintiles of the (recentered) log wage. The regression model, sample, and dependent variable are the same as in Column 1 of Table 2 except for the wage variables. Estimates are measured relative to ads in the bottom wage quintile, which is normalized to zero. Whiskers denote 95% confidence intervals, and standard errors are clustered at the user-by-month level.

Figure A9: Effect of Job Benefits on Job Seekers' Interest in Job Ads: Ad Views vs. Ad Actions



Notes: The figure shows reduced-form estimates of the effect of each benefit on two outcomes: viewing a job ad (the outcome used in Figure 5) and taking any ad action—printing, sharing, saving, applying through the platform, or navigating to the employer’s website. Both series are estimated from the across-ad specification on the same sample and with the same set of controls as Figure 5. The only difference is the outcome variable. All coefficients are normalized by the mean of the respective outcome in the control group, so the two series are reported on the same relative-effect scale.


Figure A10: Belief about Information Source by Treatment Arm




Notes: The figure uses data from the belief-elicitation survey. The sample is the survey analysis sample. Respondents were asked where they believed the wage and benefit information displayed on the job-ad previews came from. The figure displays the percentage of respondents who attributed the information to each source category, by survey treatment arm.

Figure A11: Job Benefits by Wage Decile, across Job Advertisements

(a) Value of Benefits




(b) Number of Benefits



Notes: The unit of observation is a job advertisement in the matched experimental sample (vacancies with both kununu wage and benefit information). Job ads are sorted into deciles of the kununu wage, measured at the firm–job-title level. Panel (a) plots the average value of the benefits a job offers, expressed as a percentage of the wage, using our estimated willingness to pay for each of the twelve benefits. Panel (b) plots the average number of benefits offered. We count a benefit as available when at least 20% of a firm’s reviewers confirm it.

Figure A12: Stated Usefulness of Wage and Benefit Information in Job Advertisements



Notes: At the end of the belief-elicitation survey, respondents were asked whether they would find it useful if job platforms displayed information about wages and about additional benefits in job advertisements. Respondents answered on a seven-point scale, where 1 indicates they would not find the information useful at all and 7 that they would find it very useful. The figure shows the distribution of answers, for wage information in Panel (a) and for benefit information in Panel (b). Bars report the percentage of respondents choosing each rating in the control group and in the pooled treatment arms, with the percentages displayed above the bars. The sample is the survey analysis sample of 1,484 respondents. For respondents in the arm that was explicitly informed that the information comes from kununu, the question referred to information from kununu.com, and their answers are pooled with the remaining treatment arms.

B Further Tables

Table A1: Experimental Design by Wave

Panel A: March				
Treatment group	Wage info	Benefit 1	Benefit 2	Benefit 3
Control	None	–	–	–
Wage only (avg.)	Average	–	–	–
Wage only (med.)	Median	–	–	–
Family	Average	Flexible hours	Work from home	Childcare
Commute	Average	Parking	Public transport	Company car
Nutrition	Average	Canteen	Food allowance	Coaching
Health	Average	Childcare	Health measures	Company doctor
Work environment	Average	Flexible hours	Coaching	Employee events
Panel B: April				
Treatment group	Wage info	Benefit 1	Benefit 2	Benefit 3
Control	None	–	–	–
Wage only (avg.)	Average	–	–	–
Wage only (med.)	Median	–	–	–
Family (20%)	Average	Flexible hours	Work from home	Childcare
Family (50%)	Average	Flexible hours	Work from home	Childcare
Family (med., 20%)	Median	Flexible hours	Work from home	Childcare
Family (med., 50%)	Median	Flexible hours	Work from home	Childcare
Family (reordered)	Average	Childcare	Work from home	Flexible hours
Panel C: May				
Treatment group	Wage info	Benefit 1	Benefit 2	Benefit 3
Control	None	–	–	–
May 1	Average	Public transport	Canteen	Health measures
May 2	Average	Company car	Company doctor	Employee events
May 3	Average	Childcare	Food allowance	Employee events
May 4	Average	Flexible hours	Work from home	Parking
May 5	Average	Work from home	Parking	Company car
May 6	Average	Public transport	Food allowance	Health measures
May 7	Average	Company doctor	Canteen	Coaching

Table A2: Summary of Session and User Characteristics

	Full Sample	Analysis Sample (Matched Sample)
A. Session Statistics		
Avg session length (min)	8.7 (14.54)	8.91 (14.65)
Avg number of searches	1.76 (1.37)	1.78 (1.39)
Avg total impressions per session	57.43 (96.74)	59.82 (98.35)
Avg in-sample impressions per session	50.87 (87.54)	31.01 (56.28)
Avg total views per session	2.51 (4.38)	2.6 (4.47)
Avg in-sample views per session	2.05 (3.48)	1.24 (2.27)
Avg total actions per session	0.26 (1.27)	0.27 (1.31)
Avg in-sample actions per session	0.21 (1.12)	0.13 (0.69)
<i>Region (%)</i>		
Zurich	32.4	31.73
Switzerland	78.82	77.09
Germany/France/Italy/Austria	7.19	6.98
<i>Device properties (%)</i>		
Mobile	54.2	54.12
Desktop	44.07	44.11
Language: DE	83.32	83.71
Language: EN	6.91	6.87
Has assigned occupation (%)	74.15	74.13
<i>Assigned occupation (%):</i>		
Managers	5.19	5.19
Professionals	26.57	26.58
Clerical support workers	18.75	18.72
Technicians and associate professionals	17.72	17.73
Service and sales workers	13.96	13.96
Craft and related trades workers	5.95	5.94
Plant and machine operators, and assemblers	4.76	4.76
Elementary occupations	6.54	6.56
<i>Number of sessions</i>	428,304	407,642
B. User Statistics		
Avg sessions per user	1.55	1.53
Median of sessions per user	1	1
Has a platform login ID (%)	3.89	4.34
<i>Number of users</i>	275,983	265,983

Notes: The table reports descriptive statistics on job search sessions (Panel A) and users (Panel B) in the field experiment. The unit of observation is a session in Panel A and a user (cookie identifier) in Panel B. The first column includes all sessions and users satisfying the data-cleaning restrictions described in Appendix C. The second column additionally restricts the sample to sessions and users with at least one impression of a job ad matched with kununu review information, our final analysis sample. “Total” impressions, views, and actions count all activity in a session, whereas “in-sample” counts include only impressions satisfying the sample restrictions of the respective column (impressions appearing in the search-result list in the first column and, among those, impressions on ads matched with kununu review information in the second), together with the views and actions resulting from them. The number of searches is the number of distinct search-result pages opened in a session. Region, device, and language shares are based on Google Analytics geolocation and device information. A session is assigned an (one-digit ISCO) occupation if more than half of the job ads shown in the session are associated with that occupation. Occupation shares are computed among sessions with an assigned occupation. “Has a platform login ID” is the share of users registered on one of the platform’s job boards. Standard deviations are in parentheses.

Table A3: Balance Tests for Users' Characteristics: March

	0	1	2	3	4	5	6	7	p-value
Share in Zurich	0.301 (0.004)	0.299 (0.004)	0.299 (0.004)	0.294 (0.004)	0.293 (0.004)	0.299 (0.004)	0.298 (0.004)	0.296 (0.004)	0.894
Share in Switzerland	0.727 (0.004)	0.721 (0.004)	0.725 (0.004)	0.721 (0.004)	0.719 (0.004)	0.725 (0.004)	0.723 (0.004)	0.727 (0.004)	0.763
Share in Germany/France/Italy/Austria	0.071 (0.002)	0.076 (0.002)	0.072 (0.002)	0.069 (0.002)	0.073 (0.002)	0.071 (0.002)	0.071 (0.002)	0.070 (0.002)	0.560
Device: mobile	0.567 (0.004)	0.569 (0.004)	0.579 (0.004)	0.573 (0.004)	0.570 (0.004)	0.566 (0.004)	0.571 (0.004)	0.567 (0.004)	0.500
Device: desktop	0.412 (0.004)	0.410 (0.004)	0.403 (0.004)	0.409 (0.004)	0.410 (0.004)	0.415 (0.004)	0.412 (0.004)	0.414 (0.004)	0.570
Device language: DE	0.838 (0.003)	0.837 (0.003)	0.842 (0.003)	0.838 (0.003)	0.841 (0.003)	0.842 (0.003)	0.843 (0.003)	0.843 (0.003)	0.771
Hour of the first session: 8-14	0.373 (0.004)	0.364 (0.004)	0.374 (0.004)	0.376 (0.004)	0.376 (0.004)	0.375 (0.004)	0.366 (0.004)	0.369 (0.004)	0.234
Hour of the first session: 14-20	0.344 (0.004)	0.348 (0.004)	0.350 (0.004)	0.341 (0.004)	0.341 (0.004)	0.338 (0.004)	0.344 (0.004)	0.344 (0.004)	0.539
Share of users active in February	0.078 (0.003)	0.081 (0.003)	0.076 (0.002)	0.076 (0.002)	0.079 (0.003)	0.079 (0.003)	0.079 (0.003)	0.081 (0.003)	0.707
<i>For users active in Feb:</i>									
Prob. of opening an ad after impression in Feb	0.065 (0.002)	0.065 (0.002)	0.069 (0.002)	0.063 (0.002)	0.061 (0.002)	0.064 (0.002)	0.067 (0.002)	0.065 (0.002)	0.687
Prob. of action on an ad after view in Feb	0.121 (0.005)	0.131 (0.005)	0.123 (0.005)	0.129 (0.005)	0.130 (0.005)	0.130 (0.005)	0.116 (0.005)	0.126 (0.005)	0.433
Number of Users	13,014	13,418	13,565	13,324	13,346	13,280	13,381	13,263	

Notes: The table reports balance tests for user characteristics across the experiment arms of the March wave. The unit of observation is a user (cookie identifier) in the final analysis sample who was active in March. Randomization occurred at the user-cookie level. Columns 0–7 denote the treatment arms, with 0 the control group (see Table A1 for the treatment-arm schedule). Entries are arm-specific means, with standard errors in parentheses. Location, device, and language shares are based on Google Analytics geolocation and device information. The hour-of-first-session variables refer to the time of day at which the user's first session in the month started. The bottom panel is computed for the subsample of users who were already active on the platform in February 2023, before the start of the experiment. The two probabilities measure pre-experiment click-through behavior from ad impressions to ad views and from ad views to ad actions. The final column reports the p -value of an F -test of equal means across all arms.

Table A4: Balance Tests for Users' Characteristics: April

	0	1	2	3	4	5	6	7	p-value
Share in Zurich	0.286 (0.004)	0.279 (0.004)	0.279 (0.004)	0.281 (0.004)	0.282 (0.004)	0.286 (0.004)	0.271 (0.004)	0.283 (0.004)	0.212
Share in Switzerland	0.711 (0.004)	0.711 (0.004)	0.707 (0.004)	0.707 (0.004)	0.712 (0.004)	0.715 (0.004)	0.699 (0.004)	0.708 (0.004)	0.305
Share in Germany/France/Italy/Austria	0.078 (0.003)	0.080 (0.003)	0.076 (0.003)	0.077 (0.003)	0.078 (0.003)	0.079 (0.003)	0.081 (0.003)	0.077 (0.003)	0.887
Device: mobile	0.555 (0.005)	0.556 (0.005)	0.561 (0.005)	0.560 (0.005)	0.557 (0.005)	0.552 (0.005)	0.560 (0.005)	0.548 (0.005)	0.515
Device: desktop	0.424 (0.005)	0.424 (0.005)	0.421 (0.005)	0.421 (0.005)	0.424 (0.005)	0.429 (0.005)	0.421 (0.005)	0.432 (0.005)	0.690
Device Language: DE	0.840 (0.003)	0.841 (0.003)	0.842 (0.003)	0.848 (0.003)	0.849 (0.003)	0.847 (0.003)	0.839 (0.003)	0.840 (0.003)	0.240
Hour of the first session: 8-14	0.372 (0.005)	0.378 (0.005)	0.372 (0.005)	0.370 (0.005)	0.381 (0.005)	0.379 (0.005)	0.379 (0.005)	0.372 (0.005)	0.497
Hour of the first session: 14-20	0.345 (0.004)	0.343 (0.005)	0.346 (0.005)	0.345 (0.004)	0.340 (0.004)	0.339 (0.004)	0.335 (0.004)	0.346 (0.005)	0.588
Share of users active in February	0.046 (0.002)	0.041 (0.002)	0.046 (0.002)	0.045 (0.002)	0.042 (0.002)	0.046 (0.002)	0.045 (0.002)	0.044 (0.002)	0.392
<i>For users active in Feb:</i>									
Prob. of opening an ad after impression in Feb	0.083 (0.006)	0.066 (0.005)	0.085 (0.005)	0.084 (0.005)	0.083 (0.006)	0.067 (0.003)	0.075 (0.005)	0.075 (0.005)	0.029
Prob. of action on an ad after view in Feb	0.116 (0.010)	0.114 (0.011)	0.135 (0.011)	0.124 (0.010)	0.121 (0.011)	0.124 (0.011)	0.116 (0.010)	0.115 (0.010)	0.896
Number of Users	11,159	11,044	11,098	11,172	11,188	11,330	11,109	11,089	

Notes: The table reports balance tests for user characteristics across the experiment arms of the April wave. Samples, variables, and tests are defined as in the notes to Table A3.

Table A5: Balance Tests for Users' Characteristics: May

	0	1	2	3	4	5	6	7	p-value
Share in Zurich	0.267 (0.004)	0.265 (0.004)	0.267 (0.004)	0.263 (0.004)	0.267 (0.004)	0.270 (0.004)	0.268 (0.004)	0.270 (0.004)	0.940
Share in Switzerland	0.717 (0.004)	0.704 (0.004)	0.712 (0.004)	0.708 (0.004)	0.708 (0.004)	0.711 (0.004)	0.707 (0.004)	0.710 (0.004)	0.448
Share in Germany/France/Italy/Austria	0.076 (0.002)	0.077 (0.002)	0.073 (0.002)	0.080 (0.002)	0.074 (0.002)	0.078 (0.002)	0.082 (0.002)	0.075 (0.002)	0.143
Device: mobile	0.559 (0.004)	0.566 (0.005)	0.562 (0.004)	0.562 (0.004)	0.569 (0.004)	0.555 (0.004)	0.560 (0.004)	0.557 (0.005)	0.328
Device: desktop	0.424 (0.004)	0.416 (0.005)	0.418 (0.004)	0.419 (0.004)	0.412 (0.004)	0.425 (0.004)	0.417 (0.004)	0.426 (0.004)	0.312
Device Language: DE	0.841 (0.003)	0.836 (0.003)	0.841 (0.003)	0.836 (0.003)	0.836 (0.003)	0.843 (0.003)	0.839 (0.003)	0.847 (0.003)	0.224
Hour of the first session: 8-14	0.369 (0.004)	0.365 (0.004)	0.374 (0.004)	0.372 (0.004)	0.370 (0.004)	0.366 (0.004)	0.368 (0.004)	0.379 (0.004)	0.398
Hour of the first session: 14-20	0.339 (0.004)	0.342 (0.004)	0.332 (0.004)	0.336 (0.004)	0.340 (0.004)	0.334 (0.004)	0.345 (0.004)	0.335 (0.004)	0.433
Share of users active in February	0.032 (0.002)	0.032 (0.002)	0.030 (0.002)	0.030 (0.002)	0.033 (0.002)	0.034 (0.002)	0.030 (0.002)	0.029 (0.002)	0.291
<i>For users active in Feb:</i>									
Prob. of opening an ad after impression in Feb	0.073 (0.005)	0.071 (0.005)	0.065 (0.004)	0.069 (0.005)	0.079 (0.005)	0.077 (0.005)	0.070 (0.006)	0.073 (0.005)	0.643
Prob. of action on an ad after view in Feb	0.127 (0.012)	0.120 (0.012)	0.120 (0.012)	0.113 (0.012)	0.142 (0.012)	0.114 (0.011)	0.121 (0.012)	0.092 (0.010)	0.291
Number of Users	12,297	11,936	12,453	12,365	12,251	12,485	12,242	12,175	

Notes: The table reports balance tests for user characteristics across the experiment arms of the May wave. Samples, variables, and tests are defined as in the notes to Table A3.

Table A6: Correspondence between Work-from-Home Information from Vacancy Texts and kununu

Vacancy text	Kununu	Number of job ads	FNR
<i>Review threshold: 20%</i>			
NO INFO	FALSE	65,832	
NO INFO	TRUE	89,920	
TRUE	FALSE	1,839	FNR: 0.172
TRUE	TRUE	8,864	FN share: 0.006
<i>Review threshold: 50%</i>			
NO INFO	FALSE	67,689	
NO INFO	TRUE	14,008	
TRUE	FALSE	3,883	FNR: 0.640
TRUE	TRUE	2,187	FN share: 0.026
<i>Likely telecommutable, Thr: 20%</i>			
NO INFO	FALSE	30,655	
NO INFO	TRUE	59,332	
TRUE	FALSE	1,464	FNR: 0.155
TRUE	TRUE	7,984	FN share: 0.010

Notes: The table cross-tabulates, at the job-ad level, the work-from-home information contained in the vacancy text against the kununu benefit information displayed in the experiment. “Vacancy text” equals TRUE if the platform’s job-ad database records the possibility to work from home for the ad, and NO INFO if the ad contains no work-from-home information. “Kununu” equals TRUE if the share of the firm’s reviewers confirming the benefit on kununu reaches the review threshold. The first panel uses the analysis-sample impressions of the treatment arms displaying work from home under the 20% threshold. The second panel uses those of the arms with the 50% threshold. The third panel repeats the 20% threshold exercise, restricting the sample to likely telecommutable occupations, defined as two-digit ISCO occupations in which at least 5% of job ads mention work-from-home availability explicitly in the vacancy text. Ads without kununu work-from-home information are excluded from the displayed rows. FNR is the false-negative rate: among ads whose vacancy text records work from home, the share that kununu classifies as not offering the benefit. The FN share relates the number of these false negatives to all job ads in the respective panel, including ads without kununu information.

Table A7: Descriptive Statistics on Wages and Benefits in Experimental and Survey Samples

	Main sample	With wage information	Survey sample
<u>Wages</u>			
Wage	6,614 (11)	6,614 (11)	6,525 (75)
<u>Benefits</u>			
Flexible hours	.83 (.00091)	.91 (.0016)	.82 (.016)
Work from home	.58 (.0012)	.66 (.0027)	.56 (.021)
Childcare	.14 (.00083)	.15 (.002)	.21 (.017)
Parking spot	.72 (.0012)	.77 (.0025)	.8 (.017)
Public transport	.8 (.0011)	.89 (.0019)	.83 (.016)
Company car	.19 (.001)	.21 (.0025)	.21 (.018)
Canteen	.56 (.0013)	.64 (.0029)	.56 (.021)
Food allowance	.37 (.0013)	.43 (.003)	.4 (.021)
Coaching	.45 (.0013)	.58 (.003)	.49 (.022)
Health measures	.5 (.0013)	.57 (.003)	.5 (.022)
Company doctor	.15 (.00093)	.11 (.0019)	.14 (.015)
Employee events	.86 (.00091)	.95 (.0014)	.92 (.012)
Observations	195848	36997	583

Notes: The table reports the wage and benefit characteristics of job ads in the field-experiment analysis sample (Column 1), analysis-sample ads with kununu wage information (Column 2), and job ads shown in the belief-elicitation survey (Column 3). The unit of observation is a job ad. The wage is the monthly kununu mean wage (in CHF), and the benefit indicators equal one if the benefit was displayed as available for the ad. For each characteristic, we compute averages among job ads that appeared at least once in a treatment arm displaying that type of information; matched ads whose analysis-sample impressions all occurred in the control group or in arms not displaying the respective information do not contribute. For this reason, the number of job ads in Column 1 is smaller than the number of matched ads in Table 1, as just over 2,000 job ads have impressions only in the control arms. Standard errors of the mean are in parentheses.

Table A8: Survey Experiment Respondents' Occupational Distribution and Balance Table

	Field Experiment	Full Sample	Control	T1a: Wage + Benefits	T1b: Wage Only	T1c: Benefits Only	T2: Wage + Benefits (kunumu)	p-value
<i>Occupational field (%):</i>								
Sales and customer advisory	23.95	12.26 (32.81)	12.59 (1.29)	11.28 (2.27)	9.48 (2.02)	11.93 (2.20)	15.42 (2.55)	0.454
Administrative and specialised secretaries, clerks	11.42	17.99 (38.42)	18.36 (1.51)	15.90 (2.62)	18.48 (2.68)	17.43 (2.57)	18.91 (2.77)	0.927
Cleaning and facility services	8.19	4.11 (19.86)	2.58 (0.62)	7.69 (1.91)	4.27 (1.39)	5.05 (1.49)	4.48 (1.46)	0.066
Insurance and pension advisory, field sales	5.04	1.82 (13.37)	1.82 (0.52)	3.08 (1.24)	1.42 (0.82)	0.92 (0.65)	1.99 (0.99)	0.573
Truck and bus drivers	7.39	0.67 (8.18)	1.06 (0.40)	0.00 (0.00)	0.95 (0.67)	0.46 (0.46)	0.00 (0.00)	0.018
Logistics and transport workers	3.43	6.81 (25.19)	6.07 (0.93)	7.69 (1.91)	7.58 (1.83)	9.17 (1.96)	4.98 (1.54)	0.436
Customer information clerks	4.02	3.37 (18.05)	4.55 (0.81)	3.08 (1.24)	2.37 (1.05)	1.38 (0.79)	2.99 (1.20)	0.088
Marketing, public relations and real estate professionals	7.28	9.16 (28.86)	10.17 (1.18)	8.72 (2.02)	7.58 (1.83)	8.72 (1.91)	8.46 (1.97)	0.792
Health professionals	3.18	10.11 (30.15)	9.86 (1.16)	9.23 (2.08)	9.00 (1.97)	12.84 (2.27)	9.95 (2.11)	0.740
Cooks, kitchen staff and service staff	10.13	4.51 (20.77)	3.49 (0.72)	6.15 (1.72)	6.64 (1.72)	4.59 (1.42)	3.98 (1.38)	0.359
Finance professionals	1.33	6.54 (24.73)	7.59 (1.03)	7.69 (1.91)	5.69 (1.60)	2.75 (1.11)	6.97 (1.80)	0.018
Teachers	3.16	3.30 (17.87)	2.28 (0.58)	4.10 (1.42)	2.37 (1.05)	5.96 (1.61)	3.98 (1.38)	0.174
Technical and engineering professionals	5.41	12.26 (32.81)	12.59 (1.29)	8.72 (2.02)	16.11 (2.54)	10.09 (2.04)	12.94 (2.37)	0.166
Machinery mechanics and repairers	2.66	1.68 (12.87)	2.12 (0.56)	2.05 (1.02)	0.47 (0.47)	0.92 (0.65)	1.99 (0.99)	0.162
Software and applications developers and analysts	3.40	5.39 (22.59)	4.86 (0.84)	4.62 (1.51)	7.58 (1.83)	7.80 (1.82)	2.99 (1.20)	0.123
Number of respondents	130,565	1484	659	195	211	218	201	

Notes: The table reports the occupational distribution of respondents in the belief-elicitation survey and tests random assignment to the treatment arms. The unit of observation is a respondent. The first column, “Field Experiment,” reports, for comparison, the share of users in the main field experiment assigned to each occupation, based on the modal occupation among the (non-missing) job ads shown to the user and restricted to users whose assigned occupation falls in one of the occupational fields selected for the survey. This column reports only a point estimate, as it is an external benchmark rather than a quantity estimated within the survey sample. The second column reports the share of survey respondents in each occupation for the full survey sample, with the standard deviation in parentheses. The next five columns report the corresponding share for each treatment arm, with the robust standard error in parentheses, from a regression of the respective occupation indicator on treatment-arm indicators. The last column reports the p -value of a joint test that the arm shares are equal. Shares are in percent.

Table A9: Survey Experiment Respondents' Summary Statistics and Balance Table

	Full Sample	Control ^a	T1a: Wage + Benefits	T1b: Wage Only	T1c: Benefits Only	T2: Wage + Benefits (kununu)	p-value
<i>Demographics:</i>							
Age	43.163 (11.982)	43.272 (0.462)	43.206 (0.845)	42.716 (0.877)	43.129 (0.831)	43.274 (0.821)	0.987
Female	0.504 (0.500)	0.478 (0.019)	0.518 (0.036)	0.474 (0.034)	0.555 (0.034)	0.552 (0.035)	0.143
Lives in Switzerland	0.929 (0.257)	0.932 (0.010)	0.903 (0.021)	0.948 (0.015)	0.936 (0.017)	0.920 (0.019)	0.495
Has a child aged 0–6	0.127 (0.333)	0.138 (0.013)	0.092 (0.021)	0.104 (0.021)	0.147 (0.024)	0.124 (0.023)	0.268
Has a child aged 7–18	0.244 (0.430)	0.241 (0.017)	0.251 (0.031)	0.251 (0.030)	0.243 (0.029)	0.239 (0.030)	0.997
Secondary education	0.410 (0.492)	0.405 (0.019)	0.431 (0.036)	0.398 (0.034)	0.408 (0.033)	0.418 (0.035)	0.965
Tertiary education	0.522 (0.500)	0.543 (0.019)	0.482 (0.036)	0.502 (0.034)	0.500 (0.034)	0.532 (0.035)	0.514
<i>Job history:</i>							
Duration of search (weeks)	17.090 (23.321)	17.839 (0.946)	14.585 (1.469)	16.810 (1.509)	16.821 (1.663)	17.647 (1.602)	0.452
Work experience in Switzerland	19.316 (13.434)	19.751 (0.541)	18.505 (0.989)	19.261 (1.022)	18.766 (0.930)	19.336 (0.967)	0.796
Employed	0.441 (0.497)	0.448 (0.019)	0.462 (0.036)	0.393 (0.034)	0.459 (0.034)	0.433 (0.035)	0.596
Monthly full-time wage (CHF)	7,122 (3,499)	7,317 (141)	6,829 (220)	7,039 (243)	7,026 (246)	6,960 (235)	0.355
Full-time equivalents (%)	83.951 (23.305)	85.404 (0.865)	82.923 (1.732)	82.332 (1.678)	81.917 (1.756)	84.095 (1.526)	0.250
<i>Sample:</i>							
Included in final sample	0.915 (0.278)	0.899 (0.011)	0.929 (0.018)	0.942 (0.016)	0.916 (0.018)	0.931 (0.017)	0.191
Number of respondents	1484	659	195	211	218	201	

Notes: The table reports summary statistics on demographics and job history for respondents in the belief-elicitation survey, and tests random assignment to the treatment arms. The unit of observation is a respondent. The first column reports the mean for the full survey sample, with the standard deviation in parentheses. The next five columns report the mean for each treatment arm, with the robust standard error in parentheses, from a regression of the respective variable on treatment-arm indicators. The last column reports the p -value of a joint test that the arm means are equal. The monthly full-time wage and the full-time equivalent are the respondent's current values if currently employed, and the corresponding last values from a previous job otherwise. The bottom panel reports the share of randomized respondents who enter the final analysis sample, that is, who did not decline consent and who answered the attention check correctly; it is the only row computed outside that sample. All remaining rows, and the number of respondents reported at the foot of the table, refer to the final sample.

Table A10: Effect of the Wage on Ad Actions Conditional on Viewing the Ad

	(1) Across	(2) Within
$\widetilde{\log(w_j^k)} \times T_{s(i)j}^w$	.014 (.014)	.068 (.077)
$T_{s(i)j}^w$	.0087** (.0042)	
Control mean	.074148	.077944
Posted wage elasticity	.192	.87
Model	LPM	LPM
Ad fixed effects	Yes	Yes
Rank fixed effects	Yes	Yes
Session fixed effects	Yes	Yes
Observations	287,009	299,044

Notes: This table reports estimates of the effect of posted wages (in logs, recentered around the sample mean) on job seekers' probability of taking any ad action—printing, sharing, saving, applying through the platform, or navigating to the employer's website—conditional on the ad having been viewed. We exclude impressions in which the ad was saved to the watchlist directly from the result list, since this action does not require viewing the ad. Column 1 reports estimates from the across-ad specification, which exploits variation in posted wages across treatment arms. Column 2 reports estimates from the within-ad specification (Equation (8)), which exploits experimental variation in average versus median wages within the same ad. The posted-wage elasticity, η , reported at the bottom of the table, is calculated by dividing the log-wage coefficient by the mean outcome. For the across-ad specification, the mean is computed using the control group. For the within-ad specification, the mean is computed using the estimation sample, as the control group is excluded. Standard errors are clustered at the user-by-month level.

Table A11: Robustness of Wage Estimates to Different Model Specifications

	(1)	(2)	(3)	(4)	(5)	(6)
$\widetilde{\log(w_j^k)} \times T_{s(i)j}^w$	.015*** (.0016)	.015*** (.0016)	.03*** (.0034)	.016*** (.0016)	.013*** (.0025)	.015*** (.0016)
$T_{s(i)j}^w$	.0089*** (.00049)	.0092*** (.00049)	.016*** (.00092)	.0075*** (.00086)	.0062*** (.00088)	.0083*** (.00049)
$\widetilde{\log(w_j^k)}$		.0022 (.0015)				
Ad has kununu wage		-.00042 (.00047)				
Saw a prior wage						-.03*** (.00048)
Log of last prior wage						.0012*** (.00042)
Mean log of prior wages						.005*** (.0014)
Control mean	.0386	.0387	.0647	.0387	.0389	.0386
Posted wage elasticity	.393	.389	.458	.409	.33	.384
Model	LPM	LPM	Poisson	LPM	LPM	LPM
Ad fixed effects	Yes	No	Yes	Yes	Yes	Yes
Rank fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Session fixed effects	Yes	Yes	Yes	No	No	Yes
Ad characteristics and occupation	No	Yes	No	No	No	No
User * month FE	No	No	No	Yes	No	No
Experiment arm FE	No	No	No	No	Yes	No
Observations	10,565,784	10,443,993	6,257,436	10,575,688	10,587,447	10,565,784

Notes: This table reports specification checks for the effect of posted wages (in logs, recentered around the sample mean) on job seekers' interest in a job ad. The outcome is an indicator for viewing an ad, and the unit of observation is an ad impression. All columns are variants of the across-ad specification in Equation (9). Column 1 reports the baseline linear probability model (LPM) with ad, rank, and session fixed effects. Column 2 replaces the ad fixed effects with fixed effects for the ad characteristics visible in the search-results list: the posting firm, the job location, the firm's kununu rating category, the uninteracted recentered log kununu wage, an indicator for ads with kununu wage information, and the detailed occupation, which captures the content of the job title. Column 3 reports average marginal effects from a Poisson model with the baseline fixed effects. Column 4 replaces the session fixed effects with user-by-month fixed effects. Column 5 replaces the session fixed effects with experiment-arm-by-month fixed effects. Column 6 augments the baseline specification with three controls summarizing the wages the job seeker has already seen in the same session: an indicator for any prior wage-shown impression, the recentered log wage of the most recent such impression, and the running average recentered log wage across all of them, each constructed from impressions preceding the current one. The posted-wage elasticity, reported at the bottom of the table, is the log-wage coefficient divided by the average ad-view rate in the control group. Standard errors are clustered at the user-by-month level.

Table A12: Robustness of Wage Effect to Different Sample Restrictions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Top 20 ranks	No short sessions	Full-time	Same treatment	No outliers	Wage quintiles 2-4	April: reordered icons
$\log(w_j^k) \times T_{s(i)j}^w$	.019*** (.0043)	.016*** (.0016)	.017*** (.0024)	.016*** (.0017)	.015*** (.0018)	.021*** (.0034)	.016*** (.0039)
$T_{s(i)j}^w$	.012*** (.0013)	.0091*** (.0005)	.009*** (.00077)	.0091*** (.00055)	.0086*** (.0005)	.011*** (.00061)	.0065*** (.0013)
Control mean	.0667	.0392	.0388	.0379	.0386	.0386	.0388
Posted wage elasticity	.279	.397	.435	.435	.378	.535	.407
Model	LPM	LPM	LPM	LPM	LPM	LPM	LPM
Ad fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Rank fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Session fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,527,514	10,186,919	4,328,519	7,951,644	10,396,050	9,835,871	916,998

Notes: This table reports robustness checks for the effect of posted wages (in logs, recentered around the sample mean) on job seekers' interest in a job ad. The outcome is an indicator for viewing an ad, and the unit of observation is an ad impression. All columns estimate the across-ad specification in Equation (9) on restricted samples. The corresponding baseline estimates are reported in Column 1 of Table A11. Column 1 restricts the sample to impressions among the 20 top-ranked ads per search. Column 2 excludes sessions shorter than 60 seconds. Column 3 restricts the sample to full-time jobs. Column 4 restricts the sample to job seekers who are always assigned to the same treatment arm. Column 5 excludes ads with posted wages more than a third above or below the two-digit ISCO occupation median. Column 6 excludes matched ads in the bottom and top quintiles of the recentered log-wage distribution and keeps all other impressions. Column 7 restricts the sample to April impressions from the control group and a treatment arm in which the wage was displayed after rather than before the benefit icons, checking whether the placement of the wage information influences the estimated effect. The posted-wage elasticity, reported at the bottom of the table, is the log-wage coefficient divided by the average ad-view rate in the control group. Standard errors are clustered at the user-by-month level.

Table A13: Robustness of Benefit Effects on Ad Views

	(1)	(2)	(3)	(4)	(5)
	Baseline	No session FE	Poisson	No short sessions	Same treatment arm
$flx_j^k \times T_{s(i)}^b$	.00021 (.00043)	-.00033 (.00037)	.00038 (.00073)	.00028 (.00044)	-.0001 (.0005)
$hok_j^k \times T_{s(i)}^b$	.0023*** (.00035)	.0022*** (.00043)	.0038*** (.00058)	.0024*** (.00035)	.0028*** (.0004)
$chld_j^k \times T_{s(i)}^b$	.0017*** (.00044)	.00086* (.00046)	.0035*** (.00077)	.0017*** (.00045)	.0022*** (.0005)
$prk_j^k \times T_{s(i)}^b$	.0012*** (.00045)	.0013*** (.00045)	.0019** (.00074)	.0012*** (.00045)	.0012** (.0005)
$pub_j^k \times T_{s(i)}^b$	.00035 (.0005)	.00061 (.00047)	.0006 (.00088)	.00035 (.00051)	.00055 (.00057)
$car_j^k \times T_{s(i)}^b$	.0028*** (.00057)	.0029*** (.00067)	.0047*** (.00088)	.0029*** (.00058)	.0029*** (.00064)
$cant_j^k \times T_{s(i)}^b$	-.00039 (.00041)	.00035 (.00044)	-.000099 (.00069)	-.00045 (.00042)	-.00065 (.00047)
$food_j^k \times T_{s(i)}^b$	.00049 (.00041)	-.00037 (.00052)	.00091 (.00069)	.0005 (.00042)	.00064 (.00047)
$coach_j^k \times T_{s(i)}^b$	.00054 (.00039)	-.00011 (.00046)	.001 (.00065)	.00056 (.00039)	.00049 (.00044)
$hlth_j^k \times T_{s(i)}^b$	.00014 (.00041)	-.00031 (.00049)	.00035 (.00069)	.00017 (.00041)	-.000053 (.00047)
$doc_j^k \times T_{s(i)}^b$	-.0017*** (.00057)	-.0017*** (.00058)	-.0023** (.001)	-.0017*** (.00058)	-.0018*** (.00063)
$evnt_j^k \times T_{s(i)}^b$	-.0011* (.00057)	.00052 (.00048)	-.0015 (.00097)	-.0011* (.00058)	-.0011 (.00064)
Control mean	.0386	.0389	.0636	.0392	.0379
Model	LPM	LPM	Poisson	LPM	LPM
Ad fixed effects	Yes	Yes	Yes	Yes	Yes
Rank fixed effects	Yes	Yes	Yes	Yes	Yes
Session fixed effects	Yes	No	Yes	Yes	Yes
Observations	11,690,672	11,714,389	7,040,247	11,269,622	8,773,842

Notes: This table reports robustness checks for the effect of job benefits on job seekers' interest in a job advertisement. The outcome is an indicator of ad view, and the unit of observation is an ad impression. All specifications exploit across-ad variation in benefit availability and correspond to the across-ad estimates (η_k^{across}) shown in red in Figure 5. Column 1 reports the baseline specification (a linear probability model with ad, rank, and session fixed effects). Column 2 removes the session fixed effects. Column 3 reports estimates from a Poisson model. Column 4 excludes sessions shorter than 60 seconds. Column 5 restricts the sample to job seekers who are always assigned to the same treatment arm. Each row reports the coefficient on the interaction between the kununu availability of the benefit and the benefit-information treatment indicator; the log kununu wage and the wage-shown indicator are included as controls but not reported. Coefficients are on the ad-view probability scale; the mean ad-view rate is reported at the bottom of each column. Standard errors are clustered at the user-by-month level.

Table A14: Heterogeneity in Wage Beliefs among Job Seekers without Wage Information

	(1)	(2)	(3)
Female	-0.0311 (0.0231)	0.0120 (0.0183)	0.00885 (0.0192)
Lives in Switzerland	0.0219 (0.0604)	-0.00874 (0.0595)	-0.00118 (0.0610)
Primary education	-0.0137 (0.108)	0.0803 (0.108)	0.0857 (0.114)
Secondary education	0.0944 (0.0775)	0.0898 (0.0719)	0.0906 (0.0747)
Tertiary education	0.213*** (0.0804)	0.128* (0.0741)	0.129* (0.0770)
Employed	-0.0168 (0.0206)	-0.000553 (0.0172)	-0.00301 (0.0179)
log(lastWage)	0.185*** (0.0404)	0.114*** (0.0339)	0.113*** (0.0358)
Last job FTE	0.0000909 (0.000649)	-0.0000826 (0.000534)	-0.0000540 (0.000557)
Work experience in Switzerland	0.000167 (0.000977)	0.00102 (0.000956)	0.000947 (0.00100)
Has a child under 6	-0.00306 (0.0295)	0.0332 (0.0248)	0.0317 (0.0258)
Has a child 7-18 yo	0.0391 (0.0254)	0.0427* (0.0228)	0.0443* (0.0236)
Occupation*Region FE		Yes	
Job FE			Yes
Observations	5272	5272	5221
R-squared	0.120	0.283	0.483

Notes: The table uses data from the belief-elicitation survey to analyze heterogeneity in elicited beliefs about respondents' own pay in the control group. The dependent variable is respondents' expected own log wage in a job. $\log(\text{lastWage})$ measures the most recent monthly full-time wage (for unemployed job seekers) or the current wage (for employed job seekers). Standard errors (in parentheses) are clustered by respondent.

Table A15: Effects of Wage and Benefit Information on Perceived Offer Likelihood and Expected Future Wages

	(1) Pr(Offer _{ij})	(2) Pr(Offer _{ij})	(3) Pr(Offer _{ij})	(4) log(Wage5yr _{ij})	(5) log(Wage5yr _{ij})	(6) log(Wage5yr _{ij})
$\widetilde{\log(w_j)}$	-0.0968*** (0.0347)		-0.111*** (0.0346)	0.483*** (0.0373)		0.485*** (0.0372)
T_i^w	-0.0216 (0.0216)	-0.0216 (0.0228)	-0.0144 (0.0237)	0.0440** (0.0203)	0.0446** (0.0213)	0.0429** (0.0218)
$\widetilde{\log(w_j)} \times T_i^w$	0.0426 (0.0629)	0.0590 (0.0689)	0.0467 (0.0626)	0.0101 (0.0707)	0.00264 (0.0775)	0.0101 (0.0692)
T_i^b			-0.00376 (0.0321)			-0.00547 (0.0269)
wfh_j			0.0163 (0.0100)			-0.00908 (0.0106)
$wfh_j \times T_i^b$			-0.0109 (0.0183)			-0.00643 (0.0175)
flx_j			0.0220 (0.0138)			-0.0238** (0.0117)
$flx_j \times T_i^b$			-0.0114 (0.0255)			0.0173 (0.0236)
$chld_j$			0.00870 (0.0106)			0.0324*** (0.0108)
$chld_j \times T_i^b$			0.00722 (0.0192)			-0.0102 (0.0197)
Occupation*Region FE	Yes		Yes	Yes		Yes
Job FE		Yes			Yes	
Observations	3992	3941	3992	3632	3549	3632
R-squared	0.130	0.208	0.132	0.460	0.534	0.462

Notes: The table shows the effects of wage and benefit information treatments in the belief-elicitation survey on respondents' beliefs about the competitiveness of a job vacancy and their future wages. The dependent variable in Columns 1–3 is the respondents' perceived likelihood of getting the job, and in Columns 4–6 the respondents' expected own log wage in the job after five years. The variables are defined as in Table 5. All log wages are recentered around the mean log wage in the field experiment. Standard errors (in parentheses) are clustered by respondent.

Table A16: Spillovers of Benefit Information on Wage, Company Car, and Work Atmosphere Beliefs

	(1) $\log(w_{ij}^{\text{post}})$	(2) $\log(w_{ij}^{\text{post}})$	(3) car_{ij}^{post}	(4) car_{ij}^{post}	(5) atm_{ij}^{post}
$\widetilde{\log(w_j)}$	0.366*** (0.0208)		0.188*** (0.0284)		
T_i^b	-0.0206 (0.0253)	-0.0217 (0.0255)	-0.00262 (0.0330)	0.00690 (0.0344)	0.131 (0.202)
flx_j	-0.0304*** (0.00927)		-0.0572*** (0.0150)		
$flx_j \times T_i^b$	-0.00803 (0.0219)	-0.0187 (0.0225)	-0.0226 (0.0305)	-0.0376 (0.0313)	-0.230 (0.187)
wfh_j	0.000255 (0.00873)		0.0713*** (0.0107)		
$wfh_j \times T_i^b$	-0.0178 (0.0188)	0.00132 (0.0171)	-0.0121 (0.0197)	-0.00192 (0.0203)	-0.0955 (0.111)
$chld_j$	0.0223** (0.00882)		-0.0683*** (0.0110)		
$chld_j \times T_i^b$	0.0149 (0.0199)	0.0125 (0.0192)	0.0687*** (0.0221)	0.0614*** (0.0224)	0.0889 (0.133)
Occ*Region FE	Yes		Yes		
Job FE		Yes		Yes	Yes
Observations	7016	6973	7016	6973	2400
R-squared	0.278	0.459	0.151	0.270	0.216

Notes: The table repeats the exercise in Table 6, testing for spillovers from the benefit-only treatment arm instead of the wage-only arm. The sample includes the control and benefit-only treatment arms. The dependent variable is the respondents' belief about the log wage in vacancy j in Columns 1–2, about the availability of a company car in Columns 3–4, and about the workplace atmosphere in Column 5. T_i^b is a dummy variable equal to one if job seeker i was assigned to the benefit-only treatment arm, which included flexible working hours, work from home, and firm-sponsored childcare, but not a company car or the wage. flx_j , wfh_j , and $chld_j$ indicate the kununu-based availability of these three benefits in the firm posting vacancy j . Their interactions with T_i^b capture the corresponding spillover onto beliefs about the wage, the company car, and the atmosphere. All other variables and the fixed-effect specifications are as defined in the notes to Table 6.

Table A17: Robustness Checks for Wage Belief Updating

	(1)	(2)	(3)	(4)	(5)
$\widetilde{\log(w_j)}$	0.456*** (0.0209)	0.413*** (0.0425)	0.497*** (0.0644)	0.462*** (0.0241)	0.439*** (0.0267)
T_i^w	0.0450*** (0.0107)	0.0465** (0.0212)	0.0688*** (0.0232)	0.0603*** (0.0139)	0.0591*** (0.0145)
$\widetilde{\log(w_j)} \times T_i^w$	0.156*** (0.0363)	0.157** (0.0794)	0.116 (0.0986)	0.161*** (0.0430)	0.223*** (0.0457)
Occ*Region FE	Yes	Yes	Yes	Yes	Yes
$\hat{\alpha}_w$	0.287 (0.059)	0.267 (0.119)	0.230 (0.172)	0.299 (0.071)	0.399 (0.069)
Observations	11872	4840	1952	11872	11872
R-squared	0.327	0.287	0.389	0.294	0.264

Notes: This table reports robustness checks for wage belief updating. The dependent variable is the respondents' expected own log wage in a job. The variables are defined as in Table 3. Column 1 reports the estimates in the main survey sample. Column 2 restricts the sample to survey respondents who found the job ads shown to them in the survey relevant (score of 4 or above on a 1–7 scale). Column 3 restricts the sample to survey respondents who assumed that the information on job ads was reported by the employers. Column 4 reweights observations to match the occupation distribution of job ad impressions in the field experiment. Column 5 reweights observations to match the occupation distribution of impressions for job ads with available wage information in the field experiment.

Table A18: Robustness Checks for Benefits Belief Updating

	(1)	(2)	(3)	(4)	(5)
T_i^b	-0.110*** (0.0142)	-0.0992*** (0.0233)	-0.187*** (0.0558)	-0.107*** (0.0187)	-0.114*** (0.0201)
ben_j^k	0.0892*** (0.00864)	0.0781*** (0.0145)	0.00806 (0.0398)	0.0589*** (0.0141)	0.0582*** (0.0164)
$ben_j^k \times T_i^b$	0.275*** (0.0182)	0.250*** (0.0287)	0.367*** (0.0482)	0.275*** (0.0238)	0.273*** (0.0261)
Occ*Region FE	Yes	Yes	Yes	Yes	Yes
$\hat{\alpha}_b$	0.302 (0.019)	0.271 (0.029)	0.370 (0.037)	0.293 (0.023)	0.290 (0.025)
Observations	35616	14520	5856	35616	35616
R-squared	0.121	0.115	0.193	0.127	0.117

Notes: This table reports robustness checks for benefits belief updating. We pool respondents' beliefs about the availability of the three benefits in one regression, as in Column 7 of Table 5, which Column 1 reproduces. The variables are defined as in Table 5. Columns 2–5 are defined as in the notes to Table A17.

Table A19: Heterogeneity in Wage Belief Updating

	(1) Female	(2) Male	(3) Wage Tercile 1	(4) Wage Tercile 2	(5) Wage Tercile 3
$\widetilde{\log(w_j)}$	0.423*** (0.0301)	0.478*** (0.0298)	0.412*** (0.0381)	0.420*** (0.0363)	0.466*** (0.0344)
T_i^w	0.0570*** (0.0149)	0.0309* (0.0162)	0.106*** (0.0204)	0.0443*** (0.0159)	-0.0204 (0.0197)
$\widetilde{\log(w_j)} \times T_i^w$	0.230*** (0.0496)	0.108** (0.0524)	0.146** (0.0672)	0.279*** (0.0597)	0.202*** (0.0599)
Occupation*Region FE	Yes	Yes	Yes	Yes	Yes
$\hat{\alpha}_w$	0.398 (0.072)	0.207 (0.093)	0.248 (0.103)	0.481 (0.084)	0.379 (0.097)
Observations	5984	5712	4216	3704	3952
R-squared	0.299	0.360	0.283	0.305	0.366

Notes: This table reports heterogeneity in wage belief updating. The dependent variable is the respondent's expected log wage in a job. All variables are defined as in Table 3. Columns 1 and 2 report estimates for female and male respondents, respectively. Columns 3–5 report estimates by tercile of reported last/current wage.

C Data Cleaning and Sample Construction

This appendix describes how we construct the analysis sample of ad impressions from the raw data. The raw data are Google Analytics event logs of Jobchannel’s job boards, which record for each user the ads shown in the search results (impressions), ad views, and ad actions, together with the wage and benefit information displayed for each ad. A session combines a user’s cookie identifier with the Google Analytics session identifier. We apply the following restrictions to construct the clean impression sample:

1. *Cookie consent.* We keep users who at least once accepted statistics cookies, both across all recorded events and among impression, view, and action events, as only those users would enter the field experiment.
2. *Bots and heavy sessions.* We exclude sessions with browsing patterns indicative of automated or non-standard use: sessions with at least 500 impressions but view or open rates below 0.2%; sessions with at least 50 impressions and view or open rates above 90%; sessions with more than six searches within less than two minutes; and all sessions of users with a positive Google spam score. We further exclude heavy sessions, defined as sessions with more than 1,000 impressions and view or open rates above 5%, or more than 100 impressions and view or open rates above 50%, as well as sessions shorter than five seconds. These thresholds are chosen based on the empirical distribution of session-level activity: the impression cutoffs of 100, 500, and 1,000 correspond approximately to the 90th, 99.5th, and 99.9th percentiles of impressions per session, and the 50% view- and open-rate cutoffs to the top two percent of session-level rates.
3. *Technical problems.* We exclude March 5–6 and April 21–24, when the platform experienced technical problems, and sessions spanning the change of the experimental setup between March and April or between April and May.
4. *Display errors.* We exclude sessions in which the displayed information was inconsistent with the assigned treatment arm, and impressions for which available wage or benefit information was not displayed due to technical errors (less than 0.05% of the sessions).
5. *Search-result impressions.* We keep impressions of job ads on search-results pages, where the experimental wage and benefit information was displayed. Impressions on promoted job ads and impressions on domains that aggregate ads from other job boards are excluded.
6. *Opt-out.* We exclude impressions of ads posted by firms that opted out of the experiment (less than 0.01% of the impressions of in-sample sessions).

The event-level regressions impose two further conditions. First, we exclude users assigned to more than one treatment arm within a calendar month (0.09% of users). Second, we exclude the few remaining impressions in which wage or benefit information was shown in the control group or deviated from the assigned arm. Together, these conditions drop approximately 0.05% of the clean-sample observations. The restriction to matched ads with benefit and wage reviews is described in the main text.

D Derivations of the Conceptual Framework

This section presents the derivations of the decompositions of the reduced-form semi-elasticities into the structural parameters.

Approach based on within-ad variation Let us consider the standard OLS expression for the semi-elasticity of applications with respect to wages η_w^{within} conditional on vacancy fixed effects φ_j and the vector of all benefit information shown ($\mathbf{b}_{ij}$) from Equation (4):

$$\begin{aligned}\eta_w^{\text{within}} &= \frac{\text{Cov}(y_{ij}, \log w_{ij} \mid \mathbf{b}_{ij}, \varphi_j)}{\text{Var}(\log w_{ij} \mid \mathbf{b}_{ij}, \varphi_j)} \\ &= \frac{\text{Cov}(\varphi_j + \beta_w (\alpha_w \log w_{ij} + (1 - \alpha_w) \log w_{ij}^{\text{prior}}) + \sum_k \beta_b^k (\alpha_b b_{ij}^k + (1 - \alpha_b) b_{ij}^{k, \text{prior}}) + \epsilon_{ij}, \log w_{ij} \mid \mathbf{b}_{ij}, \varphi_j)}{\text{Var}(\log w_{ij} \mid \mathbf{b}_{ij}, \varphi_j)} \\ &= \alpha_w \beta_w.\end{aligned}\tag{19}$$

The second line substitutes the belief-updating rules from Equations (2) and (3) into Equation (1) of the conceptual framework. The final equality follows from the random assignment of the information shown. In particular, conditional on the benefit information displayed and fixed vacancy effects, the wage information shown, $\log w_{ij}$, is independent of workers' prior beliefs, $\log w_{ij}^{\text{prior}}$ and $b_{ij}^{k, \text{prior}}$, as well as the idiosyncratic error term ϵ_{ij} . This implies that the within-ad estimate of the wage semi-elasticity combines two channels: (i) how much higher wages encourage job seekers to apply (β_w), and (ii) how much job seekers update their beliefs in response to the wage information shown (α_w).

Approach based on across-ad variation The second approach exploits that our experiment revealed information about wages and benefits across different vacancies to treated job seekers, while we withheld this information from the control group. For simplicity, we will focus only on wages for now. In this case, the econometric specification is:

$$y_{ij} = \varphi_j + \gamma_w T_i^w + \eta_w^{\text{across}} (\log w_j - \overline{\log w}) \times T_i^w + \varepsilon_{ij},\tag{20}$$

where φ_j is a vacancy fixed effect, T_i^w an indicator equal to one if individual i was shown the wage information and zero otherwise, and $(\log w_j - \overline{\log w})$ is the log wage of vacancy j , centered around the sample mean wage. With this recentering, γ_w measures the effect of revealing the average wage on applications relative to the control group. This coefficient provides insight into job seekers' prior beliefs about the overall wage level. A positive estimate of γ_w implies that revealing the average wage increases applications, suggesting that workers systematically underestimate prevailing wages. The converse holds if γ_w is negative. The coefficient η_w^{across} measures the extent to which providing wage information increases the responsiveness of applications to wage differences across vacancies. It is identified from comparing the treatment groups' differences in applications to higher- versus lower-paid jobs with the baseline differences in application rates across vacancies of the control group.

The equation implies that the treatment effect for vacancy j is given by

$$\mathbb{E}[y_{ij} \mid j, T_i^w = 1] - \mathbb{E}[y_{ij} \mid j, T_i^w = 0] = \gamma_w + \eta_w^{\text{across}} (\log w_j - \overline{\log w}) \quad (21)$$

Thus, we can derive η_w^{across} in the following way:

$$\begin{aligned} \eta_w^{\text{across}} &= \frac{\text{Cov}(\log w_j, \mathbb{E}[y_{ij} \mid j, T_i^w = 1] - \mathbb{E}[y_{ij} \mid j, T_i^w = 0])}{\text{Var}(\log w_j)} \\ &= \frac{\text{Cov}(\log w_j, \mathbb{E}[\varphi_j + \beta_w (\alpha_w \log w_j + (1 - \alpha_w) \log w_{ij}^{\text{prior}}) + u_{ij} \mid j, T_i^w = 1])}{\text{Var}(\log w_j)} \\ &\quad - \frac{\text{Cov}(\log w_j, \mathbb{E}[\varphi_j + \beta_w \log w_{ij}^{\text{prior}} + \epsilon_{ij} \mid j, T_i^w = 0])}{\text{Var}(\log w_j)} \\ &= \frac{\text{Cov}(\log w_j, \beta_w (\alpha_w \log w_j - \alpha_w \log w_{ij}^{\text{prior}}))}{\text{Var}(\log w_j)} \\ &= \alpha_w \beta_w \left(1 - \frac{\text{Cov}(\log w_j, \log w_{ij}^{\text{prior}})}{\text{Var}(\log w_j)} \right) = \alpha_w \beta_w (1 - \theta_w) \end{aligned} \quad (22)$$

The first line states the OLS formula for the regression coefficient η_w^{across} . The second line plugs in our conceptual framework. The third line follows from the random assignment of displayed wages, which implies that the treatment indicator is independent of vacancy fixed effects, prior beliefs, and the residual term. As a result, the vacancy fixed effects and residuals cancel in the treatment versus control expressions, leaving only the updating term. The fourth line simplifies the expression and introduces the definition of $\theta_w \equiv \frac{\text{Cov}(\log w_j, \log w_{ij}^{\text{prior}})}{\text{Var}(\log w_j)}$. θ_w measures how well job seekers were informed about wage differences in the first place. If prior beliefs exactly match wages, $\theta_w = 1$. In contrast, $\theta_w = 0$ indicates that prior beliefs are unrelated to actual wages. Thus, when using the across-ad variation to estimate the semi-elasticity of applications with respect to wages shown, the semi-elasticity depends on three parameters: (i) the preferences for wages β_w , (ii) updating α_w , and (iii) the extent of prior knowledge θ_w . Intuitively, if job seekers are perfectly informed about wage differences across vacancies, the across-ad information treatment does not have any bite.

Decomposing parameters using survey experiment To decompose the semi-elasticity of applications into preferences, belief updating, and wage misperception, we complement the field experiment with the survey experiment described in Section 2.2. As detailed above, we randomize survey participants into a control group that observes vacancies without wage and benefit information, while the treatment groups observe vacancies with wage and benefit information. We then elicit wage and benefit beliefs from both groups. For expositional purposes, we only focus on wages here.

Consider Equation (10) in simplified notation and without additional controls:

$$\log w_{ij}^{\text{post}} - \overline{\log w} = \iota + \gamma_w T_i + \theta_w (\log w_j - \overline{\log w}) + \sigma_w T_i \times (\log w_j - \overline{\log w}) + \epsilon_{ij} \quad (23)$$

where ι measures the misperception about the wage level at the average-paying job held by job seekers in

the absence of the information treatment. θ_w measures how well wage beliefs and benchmark wages align among untreated job seekers. γ_w and σ_w in turn estimate how much the treatment group updates their beliefs based on the information shown in the treatment condition.

How do these regression coefficients map into the parameters of our conceptual framework? θ_w directly follows from the definition. σ_w from the regression can be expressed as:

$$\begin{aligned}\sigma_w &= \frac{\text{Cov}(\log w_j, \mathbb{E}[\log w_{ij}^{\text{post}} \mid j, T_i = 1] - \mathbb{E}[\log w_{ij}^{\text{post}} \mid j, T_i = 0])}{\text{Var}(\log w_j)} \\ &= \frac{\text{Cov}(\log w_j, \mathbb{E}[\alpha_w \log w_j + (1 - \alpha_w) \log w_{ij}^{\text{prior}} \mid j, T_i = 1] - \mathbb{E}[\log w_{ij}^{\text{prior}} \mid j, T_i = 0])}{\text{Var}(\log w_j)} \\ &= \alpha_w (1 - \theta_w).\end{aligned}\tag{24}$$

The first line uses the OLS expression for the coefficient, the second line plugs in the updating framework from Equation (2), and the third line uses the definition of θ_w . Thus, σ_w is a mix between the updating parameter α_w and the level of misperception $1 - \theta_w$. Intuitively, when job seekers are perfectly informed about wages, our experiment would not show any elasticity of beliefs even under full updating. Equation (24) shows that the strength of belief updating α_w can be estimated by $\frac{\sigma_w}{1 - \theta_w}$ from Equation (23).

E Standard Errors for the Willingness-to-Pay Estimates

This appendix gives the explicit delta-method variance expression for $\widehat{\text{WTP}}^k$ used in Section 5.3, and explains how the joint covariances of the field-experiment and belief-survey estimators are recovered.

Writing the belief-regression interaction coefficients as $\sigma_w \equiv \alpha_w (1 - \theta_w)$ and $\sigma_b \equiv \alpha_b (1 - \theta_b)$, the WTP can be expressed equivalently as $\widehat{\text{WTP}}^k = (\hat{\eta}_k^{\text{across}} / \hat{\eta}_w^{\text{across}}) \cdot (\hat{\sigma}_w / \hat{\sigma}_b)$. A first-order expansion of this ratio gives

$$\begin{aligned}\text{Var}(\widehat{\text{WTP}}^k) &\approx \left(\widehat{\text{WTP}}^k\right)^2 \left[\underbrace{\frac{\text{Var}(\hat{\eta}_k^{\text{across}})}{(\hat{\eta}_k^{\text{across}})^2} + \frac{\text{Var}(\hat{\eta}_w^{\text{across}})}{(\hat{\eta}_w^{\text{across}})^2} - \frac{2 \text{Cov}(\hat{\eta}_k^{\text{across}}, \hat{\eta}_w^{\text{across}})}{\hat{\eta}_k^{\text{across}} \hat{\eta}_w^{\text{across}}}}_{\text{field experiment}} \right. \\ &\quad \left. + \underbrace{\frac{\text{Var}(\hat{\sigma}_w)}{\hat{\sigma}_w^2} + \frac{\text{Var}(\hat{\sigma}_b)}{\hat{\sigma}_b^2} - \frac{2 \text{Cov}(\hat{\sigma}_w, \hat{\sigma}_b)}{\hat{\sigma}_w \hat{\sigma}_b}}_{\text{belief survey}} \right].\end{aligned}\tag{25}$$

The two groups of terms correspond to the two independent experiments that contribute to $\widehat{\text{WTP}}^k$: the field experiment and the belief-elicitation survey. Within the field experiment, we jointly estimate $\hat{\eta}_k^{\text{across}}$ and $\hat{\eta}_w^{\text{across}}$ from the regression specification in Equation (9) and account for their full joint covariance. The wage-belief and benefit-belief regressions are estimated on the same survey respondents, so $\hat{\sigma}_w$ and $\hat{\sigma}_b$ are not independent. We recover their joint covariance by stacking the two equations and clustering at the respondent level (Stata's `suest`). The two experiments are independent by design because they use different samples, so no cross-block covariance appears. The variances of $\hat{\sigma}_w$ and $\hat{\sigma}_b$ are obtained from the respective belief regression by propagating the joint covariance of $(\hat{\alpha}, \hat{\theta})$ through $\text{Var}(\hat{\sigma}) = (1 - \hat{\theta})^2 \text{Var}(\hat{\alpha}) + \hat{\alpha}^2 \text{Var}(\hat{\theta}) -$

$$2(1 - \hat{\theta}) \hat{\alpha} \text{Cov}(\hat{\alpha}, \hat{\theta}).$$

F Belief Spillovers: Derivations

This appendix presents the derivations from Section 6. First, consider the belief spillover regression, abstracting from the control variables.

$$b_{ij}^{\text{post},k} = \iota + \kappa^k \widetilde{\log(w_j)} + \gamma_w^k T_i^w + \sigma_w^k \widetilde{\log(w_j)} \times T_i^w + \varepsilon_{ij} \quad (26)$$

where κ^k measures the relationship between beliefs on benefit k and wages, and σ_w^k measures how much this relationship changes with wage information provision. Using the definition of σ_w , we derive that:

$$\begin{aligned} \sigma_w &= \frac{\text{Cov}(b_{ij}^{\text{post}}, \log w_j | T_i^w = 1)}{\text{Var}(\log w_j | T_i^w = 1)} - \frac{\text{Cov}(b_{ij}^{\text{post}}, \log w_j | T_i^w = 0)}{\text{Var}(\log w_j | T_i^w = 0)} \\ &= \frac{\text{Cov}(\alpha_{wb}(\log w_{ij} - \log w_{ij}^{\text{prior}}) + b_{ij}^{\text{prior}}, \log w_j | T_i^w = 1)}{\text{Var}(\log w_j | T_i^w = 1)} - \frac{\text{Cov}(b_{ij}^{\text{prior}}, \log w_j | T_i^w = 0)}{\text{Var}(\log w_j | T_i^w = 0)} \\ &= \frac{\text{Cov}(\alpha_{wb}(\log w_{ij} - \log w_{ij}^{\text{prior}}), \log w_j)}{\text{Var}(\log w_j)} \\ &= \alpha_{wb} \left(1 - \frac{\text{Cov}(\log w_{ij}^{\text{prior}}, \log w_j)}{\text{Var}(\log w_j)} \right) \\ &= \alpha_{wb} (1 - \theta_w) \end{aligned} \quad (27)$$

The first line starts with the OLS definition of the coefficient σ_w , the second line plugs in the updating formula for b_{ij}^{post} , the third line takes advantage of the fact that treatment assignment is random, hence population variance and covariances across the groups are the same, the fourth and fifth lines simplify the expression.

G Experiment 2: Survey Instrument

This appendix reproduces the survey instrument used in the belief-elicitation experiment (Experiment 2). The original survey was programmed in Qualtrics and administered in German. Below we provide an English translation, together with a description of the randomization and survey flow.

Welcome to the ETH Job Survey!

With this short survey, we would like to better understand which aspects of jobs are particularly important for job seekers when choosing between vacancies. The survey takes only a few minutes.

Your answers will be stored anonymously and confidentially (click here for more information).

You can stop the survey at any time. As a thank you for participating, we will raffle Amazon vouchers with a total value of CHF 10,000 (1 voucher of CHF 1,000 and 90 vouchers of CHF 100).

You will receive more information about the lottery at the end of the survey.

- Have you read this information, and do you want to participate in the survey?

- Yes, I would like to participate.
- No, I do not want to participate.

[If No: - Are you sure that you want to make this choice? If you do not participate in the survey, you will lose the chance to win an Amazon voucher.

- I would like to participate in the survey.
- I do not want to participate in the survey.]

Only respondents who agreed to participate proceeded.

- Which occupational field best describes the occupation in which you have most recently been searching for a job?

- | | |
|--|-------------------------------------|
| • Insurance and pension advisory, field sales | • Teachers |
| • Administrative and specialized secretaries, clerks | • Health professionals |
| • Marketing, PR and real estate professionals | • Customer information clerks |
| • Cooks, kitchen staff and service staff | • Sales and customer advisory |
| • Truck and bus drivers | • Logistics and transport workers |
| • Finance professionals | • Cleaning and facility services |
| • Software and applications developers and analysts | • Machinery mechanics and repairers |
| • Technical and engineering professionals | |

- In which of these regions would you be most likely to search for a job?

- | | |
|--|--|
| • Espace Mittelland (BE, FR, JU, NE, SO) | • Eastern Switzerland (AI, AR, GL, GR, SG, SH, TG) |
| • Northwestern Switzerland (AG, BL, BS) | |
| • Zurich Region (ZH) | • Central Switzerland (LU, NW, OW, UR, SZ, ZG) |

Respondents were then randomly assigned to one of three main groups.

- **Control group:** Job ads shown without any additional information on wages or benefits beyond the standard platform layout.
- **Treatment Group 1:** Job ads shown with additional information from kununu. Within this group, respondents were further randomized into three subgroups:

- **TG1a:** Additional information on wages and benefits.
- **TG1b:** Additional information on wages only.
- **TG1c:** Additional information on benefits only.
- **Treatment Group 2 (TG2):** Job ads shown with wage and benefit information explicitly labeled as coming from kununu.

We will now show you various job advertisements in the occupational field you selected that were posted on Jobchannel’s job platforms.

[TG2 only: In addition, you will see information in a gray bar on wages and additional benefits that employees reported when rating their employer on kununu.com.]


For each job, we would like to know your personal assessment of what wage and which of the listed additional benefits you consider realistic for the job.

We are interested in your answers even if the job would not be suitable for you personally.

Below you can see a preview of how the jobs will be displayed:

Each respondent then saw a preview of three sample job ads.

Figure A13: Example of Job-Ad Previews



Notes: The figure shows an English-language illustration of the job-ad previews displayed to respondents assigned to treatment group TG2. In the first preview, annotations explicitly identify the displayed wage and benefit information as coming from kununu. Respondents assigned to TG1a saw wage and benefit information without these annotations. Respondents assigned to TG1b saw wage information only, respondents assigned to TG1c saw benefit information only, and respondents in the control group saw no additional wage or benefit information.

After viewing the example previews, respondents were shown eight real job advertisements drawn from the field-experiment sample and matched to the occupational field and region selected above. The advertisements were presented one at a time. For each advertisement, respondents saw the job-ad preview corresponding to their treatment assignment and answered the following wage and benefit questions.

- What do you personally consider realistic?

- 1. What do you think your monthly gross wage in this job would be on a full-time basis, including one-twelfth of any special payments and one-twelfth of any 13th-month salary?

CHF _____

- 2. What do you think about the following additional benefits? Would the employer...

	Yes	No
...allow you to set your own working hours?	☐	☐
...allow you to work from home?	☐	☐
...make a financial contribution toward the costs of external childcare?	☐	☐
...provide you with a company car?	☐	☐

Each respondent was also randomly assigned to one of the following three versions of Question 3. The assigned version was shown for each of the eight job advertisements.

- 3. What do you think your chances, in percent, would be of getting this job if you applied?

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

No chance

Absolutely certain

- 3. How would you rate the workplace atmosphere at this company?

Please try to answer even if you do not know the company.

Very poor (0) 1 2 3 4 5 6 Very good (7)

- 3. What do you think your monthly gross wage in this job would be **five years after starting the job**, on a full-time basis, including one-twelfth of any special payments and one-twelfth of any 13th-month salary?

CHF _____

After evaluating the eight job advertisements, respondents who had received wage and/or benefit information were asked where they thought this information came from.

- Where do you think the information [TG1a and TG2: on wages and additional benefits] [TG1b: on wages] [TG1c: on benefits] comes from?

- From the employer
- From an employer-review platform such as kununu, Indeed, Glassdoor, etc.
- Other: _____

All respondents were then asked a few follow-up questions about the jobs they had just seen and about the usefulness of displaying additional wage and benefit information.

- How relevant were the jobs you just saw compared with the jobs you typically search for?

Not at all relevant (1) 2 3 4 5 6 Very relevant (7)

- Would you find it useful if job platforms displayed information about additional benefits [TG2: from kununu.com] in job advertisements? [For example, flexible hours, childcare, company car]

Not at all useful (1) 2 3 4 5 6 Very useful (7)

- Would you find it useful if job platforms displayed information about wages [TG2: from kununu.com] in job advertisements?

Not at all useful (1) 2 3 4 5 6 Very useful (7)

The survey then continued with an attention check and questions about respondents' demographic characteristics and job search.

- Previous studies have shown that some people do not take the time to read what is displayed in surveys. To show that you have carefully read our questions, please select Red and Yellow below.

☐ Red ☐ Blue ☐ Green ☐ Yellow ☐ Black

- Could you tell us your year of birth?

- What is your gender?

- Female
- Male
- Other
- Prefer not to say

- In which country do you live?

[If Switzerland: - What is your postal code?]

- How many children aged 0–6 do you have?

0 1 2 3 4 5+

- How many children aged 7–18 do you have?

0 1 2 3 4 5+

- What is your highest completed level of education?

- None
- Attended compulsory school but did not complete it
- Compulsory school
- One-year training / bridge program
- Specialized secondary school / diploma school
- Apprenticeship in a company (EFZ, EBA) / pre-apprenticeship / vocational school / commercial school
- Academic baccalaureate / teacher-training college
- Vocational / specialized baccalaureate
- Professional examination with federal certificate / higher professional examination with federal diploma / master craftsman diploma
- Higher technical school (HF)
- Bachelor's degree: university, ETH, university of applied sciences, or university of teacher education (including FH/PH diploma)
- Master's degree: university, ETH, university of applied sciences, or university of teacher education (including licentiate / university or ETH diploma)
- Doctorate / habilitation
- Prefer not to say

- For how many weeks have you been looking for a new job? _____ weeks

Respondents were next shown three benefit symbols used by kununu and asked how they would interpret them.

- What do you understand by the following symbol if you see it for an open position?



Which statement fits better?

- Employees can freely choose their working hours and arrange them according to their preferences.
 - The employer expects employees to be highly flexible and determines when they work.
 - Other: _____
- What would you expect if you saw the following symbol for an open position?



How many days could you work from home in a full-time job?

- Not at all
 - 1 day
 - 2 days
 - 3 days
 - 4 days
 - 5 days
 - Other: _____
- What would you expect if you saw the following symbol for an open position?



The employer covers the following share of employees' external childcare costs:

- 0%
- 10%
- 20%
- 30%
- 40%
- 50%
- 60%
- 70%
- 80%
- 90%
- 100%
- Other: _____

The final substantive questions concerned respondents' employment status and wage history.

- For how many years have you been working in Switzerland? _____ years
- Are you currently employed (as an employee or self-employed)? • Yes • No

[If Yes: - How much do you currently earn per month in gross terms, scaled up to a full-time position (including one-twelfth of any special payments and one-twelfth of any 13th-month salary)? CHF _____]

[If Yes: - What is your current employment percentage? _____%]

[If No: - How much did you earn per month in gross terms in your last job, scaled up to a full-time position (including one-twelfth of any special payments and one-twelfth of any 13th-month salary)? CHF _____]

[If No: - What was your employment percentage in your last job? _____%]

The survey concluded with the lottery registration and debriefing.

Thank you very much for your participation!

As a thank you, we offer you the opportunity to take part in our lottery. In total, we are raffling Amazon vouchers with a total value of CHF 10,000 (1 voucher of CHF 1,000 and 90 vouchers of CHF 100). The winners will be notified by email.

Please enter your email address below to register for the lottery.

By providing your email address, you agree that we store your email address in accordance with the General Data Protection Regulation for the protection of personal data:

- Your email address will be downloaded from Qualtrics by the researchers and stored separately from your anonymous survey responses on the researchers' computers.
- Your email address will be deleted from Qualtrics and from the researchers' computers once the results of the Amazon voucher lottery have been communicated.
- Your email address will not be passed on to third parties. You can contact us at *email address* to request access to your email address, to correct it, or to have it deleted.

- Would you like to register for the lottery?

- Yes, with this email address: _____
- No

To conclude, we would like to inform you about the purpose of this study.

We are interested in understanding how well informed job seekers are about the wages and additional benefits that can be expected in different jobs.

For this reason, some respondents in this survey saw additional information about wages and additional benefits for open positions, while other respondents saw only the usual information. All respondents were asked to express their views about job search.

- Would you like to receive a summary of the results after this study is completed?

Your email address will be deleted from Qualtrics and from the researchers' computers once the summary has been sent.

- Yes, please send it to this email address: _____
- No

- May we contact you for possible future scientific surveys?

Your email address will be deleted from Qualtrics and from the researchers' computers within two years.

- Yes, at this email address: _____
- No

Thank you for participating in this survey. Your response has been recorded.