The Role of Wages and Benefits in Job Search*

Andreas Beerli†

Stefano Fiorin‡

Andreas Gulyas¶

Daniel Kopp†

Mahsa Khoshnama§

Michael Siegenthaler†

August 17, 2026

Abstract

This paper analyzes the importance of wages and job benefits for job seekers based on a large-scale field experiment. To quantify the elasticity of job seekers' applications with respect to posted wages and their willingness to pay for 12 job benefits, we randomly provide users of 112 Swiss job boards with supplementary information on wages and benefits associated with the positions they explore—information sourced from the market-leading employer review platform. We interpret the field-experimental estimates by measuring job seekers' prior and posterior beliefs about wages and benefits in jobs in a complementary survey experiment. The revealed-preference estimates suggest that a 1% increase in the posted wage increases the likelihood that job seekers view and apply to an advertisement by 0.4%. These estimates imply that job seekers' own-wage elasticity of job applications is approximately 3, because they update their expectations about their own pay in response to posted wages only partially and have prior information about ex ante wage differences across jobs. In addition, 4 of the 12 job benefits are highly valued by job seekers: a company car, work from home, firm-sponsored childcare, and a parking spot are valued at approximately 10%, 8%, 6%, and 4% of wages, respectively. The average position offers job benefits worth 11% of wages. Since higher-paying companies tend to offer more benefits, inequality in job value is significantly higher than inequality in wages.

Keywords: Job benefits, Job search, Wage elasticity of applications, Wage transparency
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."

† KOF Swiss Economic Institute, ETH Zürich
‡ Bocconi University
¶ University of Vienna
§ Princeton University

1 Introduction

A fundamental question in labor economics is how strongly job seekers respond to information about wages and benefits in jobs. This response depends, in turn, on three factors: 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 when presented with information on wages and benefits. These factors 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 to separately identify these three forces. 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 is sourced from the market-leading employer-review platform and allows us to study 12 non-wage benefits, including rarely studied ones such as company cars, firm-sponsored childcare, and health measures, alongside work from home, flexible hours, and other common amenities. A key contribution is that our novel experimental design recovers job seekers' wage elasticity of job applications and their WTP for each benefit entirely from revealed-preference variation on a real job search platform, where 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 field experiment. We model a job seeker's decision to apply for a vacancy as depending on their underlying preferences for wages and benefits. When job seekers receive wage or benefit information, they form posterior beliefs as a weighted average of the information shown and their prior beliefs. 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 of the wage or of different benefits. Recovering these structural preference parameters requires addressing two identification challenges: job seekers may not fully update their beliefs in response to the information shown, and they may already be partially informed about compensation differences across employers, attenuating the response to the information treatment. We tackle both challenges by complementing the field experiment with a survey that elicits job seekers' prior and posterior beliefs about wages and benefits for job ads tailored to their usual search scope.

The field experiment is implemented on Jobchannel, one of the largest private job-board networks in Switzerland, comprising 112 job platforms that together list several hundred thousand vacancies per month. For three months in 2023, users of these platforms were randomly assigned to either a control group, which experienced the platform and search-results page as before, or to different treatment arms where job seekers saw additional information on wages and benefits for job advertisements both on the search-results page and when viewing the full ad page. This 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. We define a benefit as available if the share of reviewers confirming that it is available 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. One concern is that our treatments may generate curiosity clicks that do not lead to any further actions taken by job seekers. Therefore, our secondary outcome is an indicator for taking any ad action on an opened ad, 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 experimental design generates two distinct sources of identifying variation. The first is *within-vacancy variation*: for the same job advertisement, some job seekers are randomly shown the average wage reported on kununu for the ad's job title in the advertising firm, while others are shown the median wage. An analogous design applies to benefits, where we experimentally varied the thresholds for whether a benefit was confirmed to be available. The second source is *across-vacancy variation*, implemented through a difference-in-differences design that compares 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 the information is randomized across job seekers rather than across vacancies. This allows us to control for vacancy fixed effects, which absorb all unobserved vacancy characteristics that do not vary across treatment arms, such as the vacancy text and the firm's attractiveness.

We complement the field experiment with a belief-elicitation survey experiment conducted among approximately 1,600 job seekers whom we recruit from the jobchannel platforms. The survey experiment mirrors the field-experimental design: respondents are shown eight real job ads personalized to their typical search region and occupation. The control group sees result-list previews of job ads as the control-group job seekers in the field experiment. The treatment groups see ad previews with information on wages and benefits from our field experiment. After each ad, respondents report their beliefs about the expected monthly wage they would earn if hired and whether they would expect to get certain benefits. The control group's answers reveal how well-informed job seekers are about wage and benefit differences across jobs in their typical search scope. The treatment groups' answers inform us about how much job seekers update their prior beliefs when provided with our wage and benefit information.

We first analyze the effects of posting a wage on job seekers' real job search decisions. We document several important insights. First, our field-experimental estimates reveal that posting wage information has a significant and robust effect on job seekers' search behavior. Exploiting the across-vacancy variation—which compares the differential response of treated and control job seekers to high- versus low-wage ads—we estimate a posted-wage elasticity of ad views of approximately 0.4, robust to a wide range of specification checks. Extrapolating this elasticity implies that full wage transparency would increase applications to 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 ad actions as the outcome, indicating that higher posted wages translate one-to-one into consequential search actions. Furthermore, using both the survey and field experiment, we find no evidence that wage information shifts job seekers' beliefs about the availability of not-shown benefits, workplace atmosphere, or the likelihood of getting a job offer. 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 the within-vacancy variation—which compares job seekers in the mean versus median

wage treatment groups—yields an elasticity that is more than twice as large as the across-ad elasticity, approximately 0.83. Our conceptual framework rationalizes this difference. Intuitively, since job seekers in both groups receive wage information, the within-vacancy elasticity does not depend on how much job seekers know about wage differences across vacancies. Instead, it depends only on how much they update their beliefs in response to the wage information we provide. Thus, the gap between the within-vacancy and across-vacancy elasticities identifies the degree to which job seekers misperceive wage differences across jobs. The estimated gap is approximately 0.46, which implies that in the absence of posted wage information, job seekers underestimate wage differences across ads roughly by half. Strikingly, the belief-elicitation survey provides an independent estimate with a similar magnitude. Together, these findings suggest that, despite widespread pay opacity in Switzerland—where only about one percent of job ads disclose wages—job seekers possess substantial information about wage differences within their relevant search market.

Third, we use the belief-elicitation survey to estimate belief updating and to back out the structural wage preference parameters 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 groups' priors and the benchmark wage, consistent with partial but meaningful Bayesian updating. The survey also shows that on average, job seekers in the control group underestimate full-time monthly wages by approximately 5 percent, which aligns with the field-experimental finding that posting an average wage increases ad views. Combining the field-experimental elasticities with survey-based estimates of belief updating and prior misperception, we recover the structural wage elasticity, which we term the own wage elasticity of job applications. We find an own wage elasticity of 2.5 when using the across-vacancy variation and a very similar elasticity of 2.9 when using the within-vacancy estimate. Hence, these results imply that job seekers are several times more responsive to actual wage differences than the posted-wage elasticity alone 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, and the survey delivers a correction for the fact that wages and benefits may be subject to different degrees of prior misperception and belief updating. The correction is quantitatively important: the survey reveals that job seekers are considerably less informed about benefit availability than about wages, implying that a naive 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 of the wage. We also find statistically significant WTP for firm-sponsored childcare (6.1 percent) and parking spots (4 percent). These WTP estimates for the main benefits align closely with objective cost benchmarks. We find small and statistically insignificant valuations for the other benefits. A notable outlier is 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—defined as the sum of wages and the monetized value of benefits in our

sample. We find that the average job offers benefits valued at approximately 11 percent of wages, making non-wage benefits an economically important component of total compensation. Since higher-paying firms also offer more benefits in the employer-review data, inequality in total job value is significantly larger than inequality in wages alone. For example, the standard deviation of total job value exceeds that of wages by approximately 16 percent. Hence, accounting for non-wage benefits amplifies, rather than compresses, compensation inequality across jobs, complementing evidence that amenities and wages tend to co-move 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., 2024, 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 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. Our estimates of the own-wage elasticity of 2.5–3 are similar to the elasticity of 3 obtained by Christensen et al. [2026], who adopt a belief-elicitation approach similar to ours.

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 [Burbano et al., 2024, Cullen et al., 2025, Drake et al., 2025, Le Barbanchon et al., 2021, Maestas et al., 2023, Mas and Pallais, 2017, see Mas, 2026 for a recent overview]. Our paper is distinguished by its revealed-preference approach on a real job search platform, which allows us to estimate willingness to pay for a comprehensive set of twelve benefits simultaneously, in an environment where real job seekers make consequential search decisions over real vacancies. By combining field-experimental estimates with survey-based measures of belief updating, we account for the fact that job seekers may not fully incorporate benefit information into their decisions—a correction that is typically absent from stated-preference or hypothetical-choice approaches. Because job seekers are largely unaware of differences in benefit availability across jobs within their typical search scope, a naive comparison of posted-benefit and posted-wage responses would substantially overstate the valuation of benefits.

In addition, we contribute to a growing literature on imperfect information in labor markets. Survey evidence indicates that workers hold systematically biased beliefs about pay at other firms and anchor their outside options to their current wage [Andresen et al., 2026, Caldwell et al., 2025, Jäger et al., 2024]. 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—the slope of control-group wage beliefs on actual kununu wages is 0.46—and that this estimate is independently corroborated by the gap between the within-vacancy and across-vacancy elasticities in the field experiment. Second, when provided with wage information, job seekers close only about 29 percent of the gap between their priors and the benchmark wage, consistent with anchoring to current or most recent pay. Third, we provide the first evidence on worker beliefs about job benefits, showing that workers are considerably less informed about differences in benefit provision than about wage differences.

Our results also have important 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 preferences over non-wage amenities and idiosyncratic job characteristics [Azar and Marinescu, 2024, Manning, 2021]. Our decomposition suggests that both preferences 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, so that imperfect information constitutes a distinct 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.[1] This 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], but much less is known about how wage transparency shapes workers' job search behavior—a first-order mechanism through which these policies are intended to operate. Our field experiment provides direct evidence on this channel. The across-vacancy wage elasticity quantifies the extent to which posting wages redirects applications toward higher-paying vacancies, an explicit goal of the pay transparency policies recently introduced by the European Union and a growing number of U.S. states. The large-scale nature of our experiment also allows us to estimate posted-wage elasticities separately by occupation and labor-market segment. The posted-wage elasticity is roughly twice as large in the lowest-paying as in the highest-paying occupation tercile. The elasticity is also significantly higher in occupations with a high share of female workers. These patterns show that job seekers in lower-paid and female-dominated occupations are more responsive to posted wage information, suggesting greater scope for wage transparency to redirect job search in these occupations.

The paper proceeds as follows. The next section describes the experimental design. Section 3 presents the conceptual framework. Section 4 analyzes the effects of wage information on job search behavior and recovers the own-wage elasticity of applications. Section5 estimates the WTP for twelve job benefits. Section 6 tests for belief spillovers across job attributes. Section 7 examines implications for job-value inequality, and Section 8 concludes.

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 to measure how strongly job seekers react to differences in wages and benefits between otherwise identical vacancies. 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, exposes them to a personalized set of real job ads under analogous information treatments, and elicits their beliefs to document both their priors and how these beliefs update when information is provided.

[1] See Bennedsen et al. [2023] and Cullen [2024] for a literature review.

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% 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.chur, jobs-zuerich) or sector-specific focus (e.g., healthcare, IT, hospitality). Together, these platforms list several hundred thousand vacancies per month and generate millions of monthly ad impressions.

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 consists of adding information about wages and benefits to job ads for some users. [2] 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 contain information about wages, job characteristics, workplace culture, and the availability of benefits at their firms. Panel a of Appendix Figure A1 shows the wage-review form, while Panel b shows the benefit-review form.[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 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 Figure A4). The same information appeared on the corresponding vacancy page (see Figure A5).[4]

(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.

[2] Unlike designs that require posting fictitious vacancies to generate experimental wage or benefit variation, our intervention concerns only real job advertisements and consists of displaying 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. The intervention appends wage and benefit information that is generally publicly accessible but not salient at the search-results stage. Importantly, users in the control group retained access to the same information during the experiment, as users generally have in the status quo before and after the intervention: the job ads in the search results list already include a link to the employer's profile on the same review platform. Therefore, we do not restrict access to information for any users; if anything, treated users are (at least weakly) better off through more salient disclosure.

[3] Up 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.

[4] Appendix Figure A6 provides analogous visualizations for mobile users.

Depending on the treatment arm, the displayed value corresponded either to the *average* or to the *median* reported wage. Importantly, the two wage-information treatments are visually and linguistically identical: in both treatment arms, the platform 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 numerical wage value displayed. This feature of the design generates within-vacancy variation in the wage shown to job seekers: for the same job ad, some users were shown the average kununu wage, while others were shown the median wage.

(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 reviews submitted by employees on kununu and was shown only when one review was available for the firm posting the ad. As a result, benefit information was available for 56% of vacancies during the experiment.

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 included 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. Generally, this threshold was 20%, but increased to 50% in some treatment arms. When the threshold was met, the corresponding icon and text appeared in black. When it was not met, the icon and text were shown in gray. As with the average-versus-median wage variation, the two thresholds generate within-vacancy 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 information bar appeared only when at least one type of kununu information was available for the ad. Benefit 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.

For desktop users, the platform disclosed the source and construction of the displayed information through mouse-over tooltips, allowing job seekers to assess its quality and decide whether to use it or discount it. Hovering over the wage value or benefit icons showed that the information was based on kununu reviews and, for benefits, stated the relevant availability threshold. Survey evidence described below suggests that roughly three out of four respondents correctly understood the information as coming from an employer-review website rather than directly from the employer.

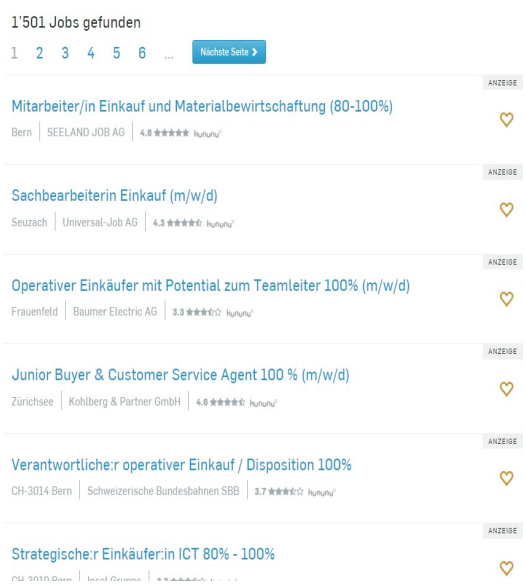
The field experiment ran on Jobchannel's job boards between March 6 and May 31, 2023. The experiment was implemented in three monthly waves, with eight approximately equally sized groups in each month (always including a control group). Table A1 summarizes the full treatment-arm schedule. The March wave 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. The May wave returned to the full set of twelve benefits, displayed average wages, and varied the benefit combinations (always using the 20% threshold) so that each benefit appeared in multiple combinations across the experiment, rather than always being coupled with the same two benefits.

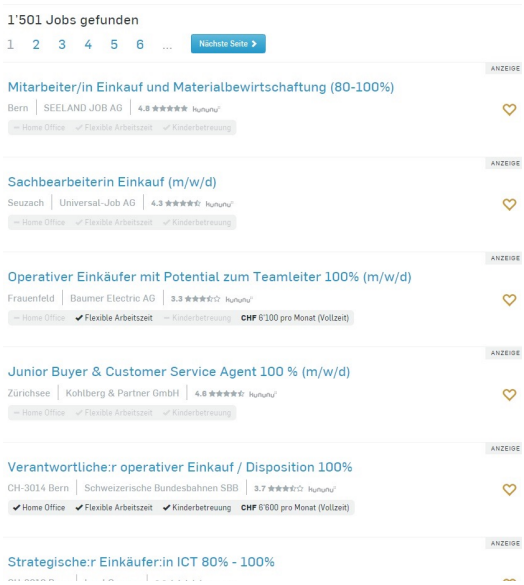
Assignment to experimental conditions occurred only for users who accepted the platform's statistics cookies. Jobchannel's A/B testing framework operates only for these consenting users and is used for routine website optimization, including experimentation. As a result, both randomization and outcome 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; randomization occurred at this user-cookie level, so that each user saw a consistent randomized information regime throughout their search. [5]

Figure 1: Screenshots of the Experiment

(a) Control Group



(b) Treatment Group



2.1.2 Outcome Measures

While 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 provides revealed-preference evidence from users' detailed interactions with real vacancies on the platform. Our primary outcome is

[5] This 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.

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*. It captures whether the information shown in the search-results list changes users' immediate interest in a vacancy. This outcome is closely aligned with the stage of the search process at which the treatment is shown: users first encounter the wage and benefit information directly in the result list and then decide whether to inspect the vacancy in more detail.

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, on which applications do not occur exclusively through the platform itself. Some users apply through the platform, but many applications take place only after users navigate to the employer's own website, and we cannot observe these applications directly.

Nevertheless, we also study a broader indicator for taking any *ad action*, which is our secondary outcome. 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 downstream search behaviors that indicate stronger interest than simply opening the vacancy page. The ad-action outcome also helps address possible concerns with interpreting ad views, such as the possibility that they merely reflect curiosity clicks. 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. 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 at least benefit information could be constructed from kununu reviews. To assess the coverage of job ads with matched benefit information relative to the broader Swiss vacancy market, Panel a in Figure 2 compares the number of matched job ads active on March 31, 2023 with the number of job openings reported by the Federal Statistical Office on the same date. The comparison is constructed at the two-digit NACE industry level. While the public sector is overrepresented in our matched sample relative to official job-opening statistics,[7] the number of matched job ads is otherwise strongly correlated with official vacancies 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. By construction, all matched ads have benefit information, and 18.9% also have wage 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

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

[7] This discrepancy may reflect measurement error in industry coding in our data.

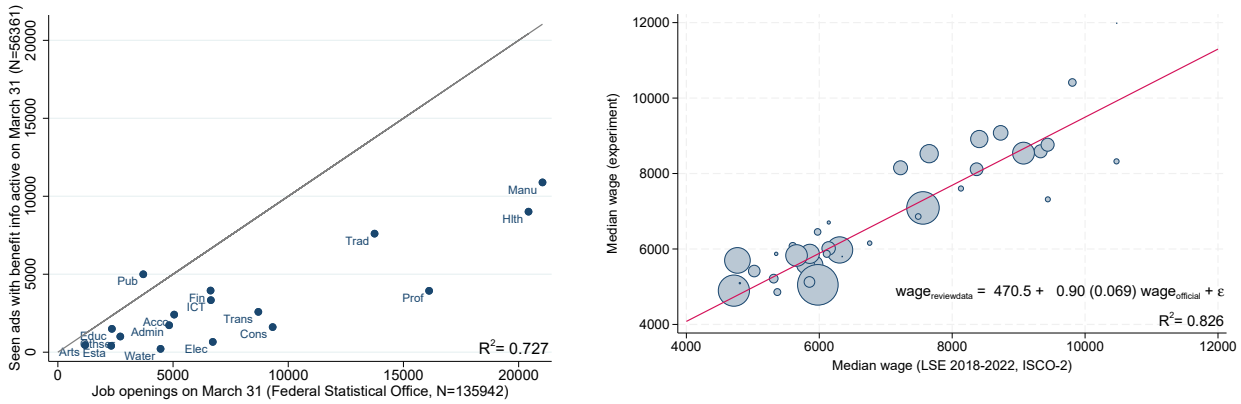
least three reviews in the underlying review data. On average, each matched ad receives about 64 in-sample impressions, 2.5 in-sample views, and 0.29 in-sample actions.

The wage analysis relies on the subset of matched ads for which kununu wage information is available (Column (3) of Table 1). These ads are posted by even larger firms, as expected given the requirement of at least three wage reviews at the firm–job-title level. They are also more concentrated in some industries, especially trade, and less concentrated in the public sector, and—in terms of occupations— overrepresent clerical support and service and sales occupations.

Figure 2: Representativeness of the Experiment and Accuracy of Displayed Wages

(a) Matched job ads viewed during experiment vs. job openings in Switzerland on March 31, by industry

(b) Median wages of job ads viewed in the experiment vs. official median wages, by occupation



Notes: Panel (a) plots the number of matched job ads in our sample active on March 31, 2023 against the number of job openings reported by the Federal Statistical Office on the same date, by two-digit NACE industry, together with a 45-degree reference line. Panel (b) compares the wages displayed in the experiment with official wage statistics at the ISCO 2-digit occupation level. Each marker is an occupation. The vertical axis shows the median kununu wage of matched job ads viewed during the experiment, averaged across ads within the occupation. The horizontal axis shows the official median monthly full-time-equivalent gross 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 in the occupation, and the inset reports the coefficient (standard error) and R^2 from the correspondingly weighted regression.

2.1.4 Users, Sessions, and Balance

Appendix Table A6 shows the summary statistics for the users and sessions in the field experiment. After imposing a set of sample restrictions,[8] our analysis sample consists of 265,983 users, who in total made 407,642 search sessions.

Appendix Table A6 summarizes user and session characteristics in the full and analysis samples. It shows that sessions in the analysis sample last about 9 minutes on average and include 1.8 searches. Users view approximately 60 ads per session, of which about half are matched job ads. About 77% of sessions are conducted from within Switzerland. A bit more than 50% of sessions occur on mobile devices, and most use the German-language version of the platform.

[8]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.

Table 1: Summary of Job Ad Characteristics

	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	6,614	6,614
	(2,107)	(2,107)	(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
Industry (%):			
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
Occupation (%):			
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
Firm Size (%):			
1-10	17.41	4.22	1.03
11-100	34.91	25.99	12.59
101-1000	27.25	39.4	38.47
$\geq$ 1001	20.43	30.39	47.91
Avg vacancy duration	88.96	84.83	79.32
	(110.00)	(104.96)	(102.52)
Avg (in-sample) impressions	61.81	63.81	68.78
	(145.18)	(137.57)	(142.79)
Avg (in-sample) views	2.5	2.55	2.9
	(8.68)	(7.85)	(8.65)
Avg (in-sample) actions	0.26	0.26	0.32
	(1.04)	(0.99)	(1.17)
Number of job ads	352,534	198,079	37,039

Notes: The table reports descriptive statistics at the job-ad level. The unit of observation is a job ad with at least one impression satisfying the sample restrictions described in Appendix C. Column 1 includes all such ads, Column 2 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 itself states a wage. The review wage is the monthly kununu mean review wage (in CHF). "Reported WFH in vacancy" indicates that the vacancy text reports the possibility to work from home. Ad language, work quotas (minimum and maximum posted full-time-equivalent percentages), temporary status, firm size (number of employees of the posting firm), and vacancy duration (days between the posting and the removal of the ad) come from the platform's job-ad database. Industry refers to the NACE section of the posting firm and occupation to the one-digit ISCO group assigned to the ad; only the most common categories are shown, and shares are computed among ads with non-missing information. "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.

The occupational composition of sessions reflects the breadth of searches on the platform.[9] 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 comparing the occupational distribution of users in our analysis sample with employment in Switzerland, we find that the experiment captures users across a broad range of occupations in proportions that are broadly related to their prevalence in the labor market.[10]

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 do not reject equality across arms, consistent with successful random assignment.

2.1.5 Quality of the information provided

The wages displayed in the experiment align closely with official wage statistics. Panel b of Figure 2 compares median displayed wages, aggregated to the two-digit ISCO occupation level, with official median wages in the same occupations, as derived from the Swiss Earnings Structure Surveys 2018–2022. The relationship is close to one-for-one, providing no indication of systematic bias in the wage information used in the experiment. 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 to the vacancy text, since this benefit is among the benefits most commonly reported directly in job ads. Appendix Table A2 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. The comparison supports the 20% threshold used in most treatment arms. Under this threshold, among ads whose text mentions work from home and for which benefit information is available, the false-negative rate is 17.2%. Because most vacancy texts do not mention work from home even when the arrangement may be available, the share of all ads for which the experimental information contradicts a positive statement in the vacancy text is only about 1%. The correspondence weakens under the 50% threshold, suggesting that this threshold may be too restrictive for some benefits.[11]

We note that, while some employers could provide wage or benefit details in the vacancy text, and users could access kununu information through the employer-review link on the platform, this information was not displayed directly in the search-results list. At a minimum, the intervention therefore made wage and benefit information more salient by moving selected kununu information forward in the search process and displaying it directly in the result list.

At the same time, the information may not have been fully incorporated into users' beliefs. The wage and benefit information came from employee reviews rather than directly from the employer, and users may therefore have discounted it. In addition, some users may have perceived the information as unreliable, for

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

[10] Appendix 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. For this comparison, users are assigned to a two-digit ISCO occupation if more than 50% of the job ads they viewed across all sessions are associated with that occupation.

[11] The third panel of Appendix Table A2 repeats the exercise for occupations in which work from home is more frequently mentioned in vacancy texts and shows the same pattern, with stronger correspondence under the 20% threshold.

example when it differed from details provided in the vacancy itself. For this reason, it is useful to measure directly how job seekers interpret the displayed information and how strongly they update their beliefs in response to it. This is the purpose of Experiment 2, which we present in the next section.

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 directly measuring job seekers' beliefs about wages and benefits, as well as their updating 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 searching for work, using ISCO-3 occupation categories and NUTS-2 regions, thereby replicating the core elements of a typical search on the platform. These answers allowed the survey to display a personalized set of job ads that closely resemble the vacancies respondents encounter on the platform.

Participants were then randomly assigned to a control group, which saw result-list previews of ads as they normally appear on the platform, or to treatment groups that displayed wage information, benefit information, or both, and—depending on the treatment—an explicit statement that this 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.[12]

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 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 a demographic module, an employment and wage-history module, and an attention check. See Appendix G for the full survey script.

[12] We 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, and crossed them 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. For the field experiment, we assigned each user the modal three-digit ISCO occupation across the ads viewed in their sessions. The table restricts attention to users whose assigned occupation falls in one of the occupational fields selected for the survey.

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 large majority live in Switzerland. About 12.7% have at least one child below age six, while 24.4% have at least one child aged seven to eighteen. The sample is relatively highly educated: 52.2% report tertiary education.

The job-history variables indicate that respondents are actively searching. The average reported search duration is 17.1 weeks. Slightly less than half of respondents are employed at the time of the survey, and respondents report on average 20.3 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.

Random assignment produced similar respondent characteristics across treatment arms. The joint balance tests generally do not reject equality across arms, consistent with successful random assignment.

2.2.3 Perceived source of information

Appendix Figure A10 shows that, among respondents shown wage information, roughly 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. Only 22% attributed the information to the employer, consistent with the fact that Swiss employers rarely report wages directly in vacancy texts. By contrast, in the benefit-only treatment arm, which was included in the survey but not in the field experiment, 48% of respondents attributed the information to the employer. This is consistent with the fact that some benefits, such as work from home, are more commonly disclosed directly in vacancy texts.

In the field experiment, the platform disclosed the source of the displayed information, but job seekers may not have paid close attention to this disclosure. The survey evidence nevertheless suggests that the vast majority of respondents interpreted wage information as coming from reviews rather than directly from the employer. This interpretation may lead job seekers to discount the information, providing an additional motivation for using the belief-elicitation survey to measure how strongly respondents update their beliefs when wage and benefit information is shown.

In the next section, we introduce a simple conceptual framework that links the experimental wage treatments in Experiments 1 and 2 to the parameters of interest, and then turn to the reduced-form evidence on wage beliefs and search behavior.

3 Conceptual Framework

To guide our empirical analysis, we present a simple framework for understanding the effects of wage and benefit information in vacancies on 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 propensity of individual i to apply

to vacancy j is given by

$$y_{ij} = \varphi_j + \beta_w \log w_j + \sum_k \beta_b^k b_j^k + \epsilon_{ij}, \tag{1}$$

where $\log w_j$ denotes the log wage offered by firm j, and b_j^k is an indicator equal to one if vacancy j offers benefit k. The term φ_j is a vacancy fixed effect that captures all vacancy-specific determinants of application behavior that are 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 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 elasticity is 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 value of β_w indicates that workers are highly responsive to wages, leading to lower wage-setting power of employers. 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 amount of wages a worker would be willing to forgo in exchange for receiving 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. There are several reasons why this may occur. For example, individuals may have anchored expectations around their current pay [Altmann et al., 2026, Andresen et al., 2026, Jäger et al., 2024], may not fully trust the information displayed, may anticipate discrimination in the hiring process, or may 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

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

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

where α_w and α_b govern the extent to which information about wages and benefits is incorporated into posterior beliefs. When they are equal to one, job seekers fully update their beliefs to the wage and benefit information, respectively, whereas lower values indicate limited updating. 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. 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. In Section 6 we study belief spillovers in detail, but 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-vacancy variation The first approach exploits within-vacancy variation by randomly varying the wage and benefit information shown to job seekers for the same vacancy. The application 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:

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

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.

As workers respond to their posterior beliefs rather than directly to the wage information shown, $\log w_{ij}$, the posted wage elasticity η_w^{within} does not necessarily identify the structural application elasticity with respect to beliefs about own wages, β_w.

As we derive in Appendix D, the elasticity of applications with respect to posted wages from regression equation (4)—identified conditional on vacancy fixed effects φ_j and the vector of benefit information shown $\mathbf{b}_{ij}$—can be expressed as:

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

This implies that the within-vacancy estimate of the wage 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). Intuitively, if job seekers do not update their prior beliefs in response to the wage information shown, the experiment would yield a zero elasticity of applications even though $\beta_w > 0$.

Approach based on across-vacancy 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 will focus only on wages for now. In this case, the econometric specification is:

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

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 $\left(\log w_j - \overline{\log w}\right)$ 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 the prevailing wages. The converse holds in case of a negative γ_w. 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.

Without loss of generality, let us consider the case of two vacancies, one paying $\log w_1$ and the other paying $\log w_2 > \log w_1$. Using the OLS formula, η_w^{across} can be written as follows:

$$\begin{aligned}\eta_w^{\text{across}} &= \frac{E[y_{ij}|T_i^w = 1, \log w_2] - E[y_{ij}|T_i^w = 1, \log w_1] - (E[y_{ij}|T_i^w = 0, \log w_2] - E[y_{ij}|T_i^w = 0, \log w_1])}{\log w_2 - \log w_1} \\ &= \alpha_w \beta_w (1 - \theta_w), \end{aligned} \tag{7}$$

where $\theta_w \equiv \left(E\left[\log w_2^{\text{prior}}\right] - E\left[\log w_1^{\text{prior}}\right]\right) / (\log w_2 - \log w_1)$ measures how well informed job seekers are about wage differences across vacancies in the absence of posted wage information. If $\theta_w = 1$, job seekers' priors fully reflect actual wage differences. A value of $\theta_w = 0$ indicates that priors are uninformative, and a negative value would imply that job seekers systematically mistake high-paying vacancies for low-paying ones. Thus, when using the across-vacancy variation to estimate the elasticity of applications with respect to wages shown, the 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 vacancy information treatment does not have any bite.

One testable prediction of this framework is that, if $\theta > 0$, the application elasticity using the within vacancy variation η_w^{within} should be larger than the elasticity estimated from the across vacancy variation η_w^{across}. In fact, one can back out θ_w from relating the two elasticities: $\theta_w = 1 - \frac{\eta_w^{\text{across}}}{\eta_w^{\text{within}}}$.

Next, we turn to how job seekers react to posted wages in vacancies.

4 Wages

In this section, we leverage the field experiment to estimate the revealed-preference effect of wage posting on job seekers' behavior, and then use the follow-up survey experiment to decompose the estimate of η into its 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 that the field experiment generates two sources of wage variation that can be exploited for causal identification.

Within-ad variation The first source of variation is the *within-ad variation* in wages displayed on the same job ad. Throughout, we restrict the sample to impressions of ads with kununu benefit information. Ads without matched wage information remain in the sample. The empirical counterpart of Equation (4) is

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

where $y_{s(i)j}$ is a binary outcome capturing job-search behavior by job seeker i in session s on vacancy j, such as viewing the ad or taking an action on the vacancy page. The variable $\log(\widetilde{w_{s(i)j}})$ is the log of the

wage displayed for vacancy j. This is either the median wage (median-wage condition) or the mean wage (mean-wage condition), both recentered around the mean log wage (the log of CHF 6,185). The posted-wage indicator, $T^w_{s(i)j}$, 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 $b^k_{s(i)j}$, which equals one if a benefit k is displayed as available at the firm posting vacancy j, interacted with a session-level indicator, $T^k_{s(i)}$, which equals one if session s belongs to the treatment arm that displays information about benefit k (see Section 5.1 for further discussion).[13] 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^{within}_w. It captures the semi-elasticity of job applications with respect to the posted wage. Quantitatively, $\eta^{within}_w/100$ measures the change in the job search outcomes associated with a one-percent increase in the posted wage. Due to the vacancy fixed effects, it is identified only from ads with experimental variation in wages: it compares the decisions of job seekers who saw the mean wage with those who saw the median wage for the same job ad.

Across-ad variation Our second empirical approach, based on *across-ad variation*, contrasts the job seeker behavior in the information treatment arms with that in the control group. Intuitively, we compare the difference in click rates for job seekers viewing a high- versus a low-wage job in the treatment group, and then subtract the difference for job seekers viewing the same ads in the control group. To abstract from the within-ad variation, we restrict the sample to the mean-wage treatment arms and the control group. The empirical counterpart of Equation (20), which implements the difference-in-differences (DiD) design, is

$$y_{s(i)j} = \gamma_w T^w_{s(i)j} + \eta^{across}_w \log(\widetilde{w_{s(i)j}}) \times T^w_{s(i)j} + \sum_{k=1}^{K} \eta^{across}_k b^k_{s(i)j} \times T^k_{s(i)} + \varphi_j + \chi_s + \pi_r + \varepsilon_{s(i)j} \qquad (9)$$

The outcome $y_{s(i)j}$, the recentered displayed wage $\log(\widetilde{w_{s(i)j}})$, and the wage-treatment indicator $T^w_{s(i)j}$ are defined as in Equation (8). Two differences are worth highlighting. First, because the across-ad sample excludes the median-wage arms, the displayed wage is always the mean wage. Second, $T^w_{s(i)j}$ now equals zero not only for ads without kununu wage information but also for all impressions in the control group. The new element in (9) is therefore the uninteracted term $\gamma_w T^w_{s(i)j}$. Because $T^w_{s(i)j}$ varies within an ad between treated and control impressions, it 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. As shown in Section 3, its sign is informative about whether job seekers on average under- or overestimate prevailing wages.

The vacancy fixed effects φ_j are central to identification in both designs. They absorb all observed and unobserved vacancy characteristics that do not vary across treatment arms—the vacancy text, the job title, the firm identity, and the job location. In the across-ad design, they ensure a causal interpretation of η^{across}_w by accounting for the possibility that ads that are more attractive in observed and unobserved dimensions might pay higher wages, thereby distinguishing our across-ad estimates from studies that have estimated

[13] Since we restrict the sample on vacancies with benefit information, the indicator $T^k_{s(i)}$ carries no vacancy j index. Thus, whether benefit information is displayed depends only on the treatment arm and is constant within a session. In contrast, $b^k_{s(i)j}$ varies between ads in a session because firms differ in the benefits they offer.

similar regressions using observational data. In the within-ad design, the vacancy fixed effects ensure that the mean- versus median-wage comparison is based on the same sample of ads.

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 absorb the baseline effect of providing benefit information, $T^k_{s(i)}$, and imply that the DiD design focuses on wage differences across ads that appeared in the same session and thus in the same treatment arm. We include them in our baseline specification for three reasons. First, they account for the search context, including the job seeker's characteristics and preferences, as well as the characteristics of other job ads displayed simultaneously, such as average wages or the overall attractiveness of competing ads. Second, they account for potential impacts of the treatment arms on job seekers' search behavior in a session, such as the total number of ads viewed or screened in detail. Third, they restrict attention to ads that job seekers most likely compared directly against each other, which 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. We show below that omitting the session fixed effects yields quantitatively similar but less precise estimates of γ_w and η_w^{across}, suggesting that they are not strictly necessary for identification but reduce noise.

4.2 Posted Wage Elasticities

Table 2 uses the analysis sample described in the design 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 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: an ad shown to a job seeker in the search results. 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, standard errors are clustered at the user-by-month level, the level at which treatment assignment varies.

The table provides several important insights. First, job seekers' views are higher for higher-paying jobs. In our baseline across-ad specification (Column 1), the semi-elasticity of ad views with respect to log posted wages is 0.015. Given that job seekers view 3.86% of the ads shown in their result lists, this implies an elasticity of ad views with respect to posted wages of $\eta_w^{across} = 0.393$. This elasticity directly measures the extent to which wage information shifts applications toward higher-paying vacancies. The estimate implies that full wage transparency would increase applications to vacancies in the top quintile of the wage distribution by approximately 30% relative to those in the bottom quintile.[14]

Second, the posted wage elasticities when exploiting within-ad variation are approximately twice as large as the corresponding across-ad elasticities, with implied elasticities η_w^{within} of 0.83 (Column 2) and 0.6 (Column 4). As highlighted by our conceptual framework in Section 3, the difference between the two elasticities speaks to the extent to which job seekers are informed about pay differences across jobs in the absence of posted wage information. In particular, the two baseline estimates suggest that θ_w equals $\theta_w = 1 - (\eta_w^{\text{across}}/\eta_w^{\text{within}}) = 1 - 0.015/0.033 \approx 0.54$, which implies that a 10% wage gap between two jobs

[14] Figure A7 shows that vacancies with wages in the top quintile are viewed 0.011 p.p. more likely in the treatment group relative to the vacancies in the bottom quintile. The average vacancy view rate is 3.8 percent.

in the kununu data translates into a 5.4% perceived gap among job seekers without wage information. The implied estimate of θ becomes slightly smaller if we restrict the across-ad estimate to ads in the central range of the wage distribution, where the within-ad elasticity is also identified.[15]

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, though, consistent with the role of the session fixed effects discussed in Section 4.1: by comparing only ads that were plausibly evaluated against one another, they absorb differences in click propensities across job seekers and search contexts without materially affecting the point estimate.

Fourth, the coefficients on the wage-shown indicator ($T^{w}_{s(i)j}$) imply that displaying a wage equal to the sample mean, relative to providing no wage information, strongly increases job seekers' propensity to engage with an ad. The positive estimates indicate that, on average, job seekers underestimate wages at firms for which wage information can be matched, and that providing salary information in job advertisements increases job seekers' interest in a position.

Fifth, the estimated wage elasticities are very similar when using ad actions rather than ad views 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 our main ad view estimates are biased upward by "curiosity clicks"—the concern 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 can directly test for such a reversal by estimating the effect on the probability that job seekers take an ad action conditional on viewing the full vacancy details shown in Appendix Table A10. The insignificant wage effects in this table are inconsistent with such a reversal.

Our estimates of the posted-wage elasticities η 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., 2024, 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 application volumes in the 0.5–1.0 range. A conceptual contribution of our paper is that we show that these posted-wage elasticities may underestimate job seekers' own-wage elasticity of job applications since they have prior information about wage differences across jobs or update their beliefs about their own pay only partly to posted wages.

Robustness. The estimates for η^{across}_{w} 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, plausibly because it compares decisions across observably similar ads rather than focusing solely on decisions taken on the same ad.[16] Column 3 shows similar wage estimates using a multinomial logit discrete-choice model, interpreting the

[15] Column 6 of Appendix TableA12 restricts the sample to ads in the middle three quintiles of the recentered log-wage distribution. Relating the estimated elasticity to the within-ad results yields an estimate of $\theta_w = 1 - 0.021/0.033 \approx 0.36$.

[16] The 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.

Table 2: Effect of posted wages on job seekers' interest in a job ad

	Ad views		**Ad views, wage-info ads**		**Ad actions**	
	(1)	(2)	(3)	(4)	(5)	(6)
	Across	Within	Across	Within	Across	Within
$\widetilde{\log(w_j^{\mathrm{k}})} \times T^w_{s(i)j}$	.015***	.033***	.013***	.026***	.0014***	.0037
	(.0016)	(.009)	(.0025)	(.0093)	(.00045)	(.003)
$T^w_{s(i)j}$	.0089***		.0071***		.0012***	
	(.00049)		(.00078)		(.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 on job seekers' interest in a job advertisement. The outcome is an indicator of ad view (Columns 1–4) and an indicator for taking any ad action—printing, sharing, saving, applying through the platform, or navigating to the employer's website (Columns 5–6). The unit of observation is an ad shown to a job seeker in the list of search results (ad impression). The sample corresponds to the analysis sample defined in Section 2.1.4. Columns 1, 3, and 5 report estimates from the across-ad specification that focuses on the control group and all mean-wage treatment arms. Columns 2, 4, and 6 report estimates from the within-ad specification (Equation (8)) that focuses on mean- and median wage treatment arms. Columns 3–4 restrict the sample to ads for which kununu wage information is available and replace the session fixed effects with ad and rank fixed effects. The posted-wage elasticity, 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. Standard errors are clustered at the user-by-month level.

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 multinomial 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, the wage estimates are unaffected by controlling for the wages job seekers already observed on prior impressions in the same session (Column 6). These prior wages shape clicking behavior, but by random assignment they are orthogonal to the wage displayed on the current ad.

In addition, Appendix Table A12 shows that the estimated wage effects are similar when we exclude short sessions, restrict to full-time jobs, or to job seekers who never switch treatment arms across sessions, limiting concerns that exposure to multiple arms attenuates the experimental estimates. While the results are also robust to excluding ads with posted wages substantially above or below the occupational median, the posted wage elasticity becomes slightly larger (around 0.53) if we focus on ads in the three middle quintiles of the wage distribution. Finally, the wage elasticity is also unchanged in the April treatment arm, in which the wage was displayed to the right of the benefit icons rather than to the left.

Directed Search. An implicit assumption of our framework is 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 settings, job seekers apply

for higher-paying jobs only up to the point where increased competition offsets the wage gain. Recent experimental evidence provides empirical support for this mechanism [e.g., Belot et al., 2022, He et al., 2023].

As a first test of this mechanism, Appendix Figure A7 reports the effect of wages on ad views non-parametrically, using indicators for each quintile of the occupational wage distribution. The wage effect rises approximately linearly across the lower quintiles but flattens above the fourth quintile. This is the pattern a competition channel would generate if higher wages no longer translate into higher interest at the top. The same pattern, however, is also consistent with random search if job seekers have reservation wages that cluster around the upper end of the wage distribution such that a high posted wage exceeds the reservation wage of nearly all job seekers.[17]

To assess the competition channel directly, we thus draw on the follow-up belief-elicitation survey, in which respondents reported the probability that they would be hired if they applied to a vacancy (Appendix Table A15). Consistent with the pattern in Belot et al. [2022], the negative effect on the log wage shows that job seekers in the control group view higher-paying jobs as harder to get. However, the positive estimate of the interaction term with the wage shows that job seekers who are informed that the job pays well report higher rather than lower hiring chances—an estimate that is not statistically significant. Consistent with Christensen et al. [2026], we thus find no evidence that wage information strengthens perceptions of competition for high-wage jobs, which in turn reduces concerns that our estimates are biased by directed search behavior.

Heterogeneity. Figure 3 illustrates heterogeneity in the posted wage elasticity of ad views across occupations. Panel (a) reports estimates by broad occupational group (ISCO 1-digit occupations). The figure suggests that the posted wage elasticity declines with 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) examines heterogeneity by terciles of the occupational wage level. 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 higher in occupations with higher female employment shares. The higher posted wage elasticity in female-dominated occupations is consistent with our survey-experimental evidence that women update wage beliefs more strongly to wage information than men.[18] Both differences are statistically significant: the posted wage elasticity is significantly higher in the lowest-paying than in the highest-paying tercile of occupations, and significantly higher in the most female-dominated than in the least female-dominated occupations. The gender-related heterogeneity aligns with Jalal [2026], who finds that pay disclosure raises women's applications by more than men's, and stands in contrast to studies reporting little heterogeneity in the wage elasticity of applications or recruitment by gender [Balgova et al., 2025, Dal Bó et al., 2013].

[17] In addition, the plateau could arise from diminishing marginal value of further wage increases at the top or from skepticism about implausibly high posted wages.

[18] Appendix 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: Posted elasticity of ad views: Heterogeneity by occupation



(a) By occupational group

(b) By wage level and female share

Notes: The figure shows heterogeneity of the posted wage elasticity of ad views by occupation. Panel (a) shows estimates by broad occupation (ISCO one-digit occupation). Panel (b) shows estimates by tercile 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. Wage elasticities are calculated by dividing the wage coefficient by the average ad view rate in the control group. The occupational wage level is the median wage in the ISCO 3-digit occupation, and the female employment share is the share of women in employment in the ISCO 3-digit occupation, both computed from the pooled Swiss Earnings Structure Survey (Lohnstrukturerhebung) waves 2018, 2020, and 2022.

4.3 Beliefs and Prior Updating

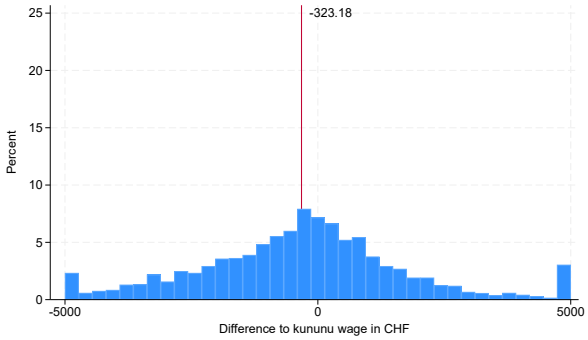
Our theoretical framework implies that the posted-wage elasticities η combine job seekers' own-wage elasticity of job applications β_w, their misperception of pay differences across jobs in the absence of wage information $1-\theta_w$, and their belief updating when given wage information α_w. Comparing the across-ad and within-ad specifications, we estimated θ_w to lie between 0.36 (middle three quintiles) and 0.54 (full sample). In this section, we use the belief-elicitation survey to estimate α_w and to provide an independent estimate of θ_w.

The key conceptual question when quantifying misperceptions in our field experiment is not whether job seekers can tell higher- from lower-paying jobs *across the labor market as a whole*. Rather, it is whether job seekers 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*. We mirror this setting by personalizing the ads we show respondents to their typical search region and occupation. For the same reason, our preferred regression specification compares the treatment and control groups only among job seekers within the same occupation-region cell.

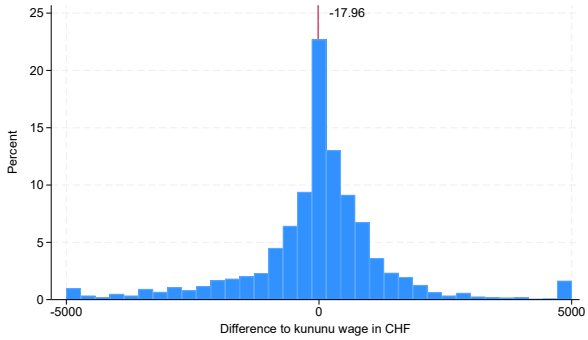
We start with a descriptive analysis of wage beliefs. Panels (a) and (b) of Figure 4 plot the distribution of the gap between a respondent's expected own pay in a job and the job's matched kununu wage, winsorized at ±5,000 CHF. We use the analysis sample described in Section 2.2. In the control group (Panel (a)), priors are widely dispersed around the benchmark and, on average, fall below it by 323 CHF or 4.8% of the mean full-time kununu wage. Panel (b) shows the same distribution for respondents who were informed about the kununu wage. Compared to the control group (Panel (a)), 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 on belief updating by plotting job seekers' own-wage beliefs against the kununu wage, separately for the control and wage-treatment groups. The binscatter for the control group indicates that job seekers without wage information can, on average,

distinguish between high- and low-paying jobs. Yet, they tend to underestimate wage dispersion: their beliefs lie below the kununu wage for high-wage vacancies and above it for low-wage ones. Providing wage information (red line) steepens the relationship between the expected and the kununu wage and shifts beliefs toward the benchmark wage. However, the slope remains below the slope of one that full updating would imply.

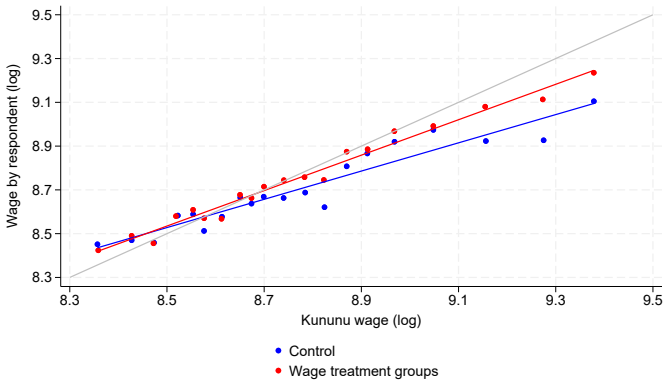
Figure 4: Descriptive evidence on wage misperceptions and belief updating from the belief-elicitation survey



(a) Wage beliefs: Control group



(b) Wage beliefs: Wage treatment groups



(c) Wage beliefs versus kununu wages

Notes: Panel (a) shows wage information deficits in the control group, measured as the difference between job seekers' wage priors and the job-level kununu wage. Panel (b) shows the difference between the wage beliefs of job seekers in the wage treatment groups and the job-level kununu wage. Panel (c) shows a binned scatter plot of the own-wage beliefs and the kununu wage, separately for the control group (in blue) and the wage treatment groups (in red). Wages reflect monthly full-time wages.

Table 3 formalizes the visual relationship in Panel (c) of Figure 4 by reporting estimates from the following regression:

$$\widetilde{\log(w_{ij}^{\text{post}})} = \theta_w \widetilde{\log(w_j)} + \gamma_w T_i^w + \sigma_w \widetilde{\log(w_j)} \times T_i^w + \varepsilon_{ij}. \tag{10}$$

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 recentered 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) for comparability.

To see how the regression coefficients map into the parameters of our conceptual framework, note that

θ_w captures the slope of wage beliefs on actual wages in the control group. Under a binary regressor, this is

$$\theta_w = \frac{E\left[\log w_{ij}^{\text{post}} \mid T_i = 0, \log w_2\right] - E\left[\log w_{ij}^{\text{post}} \mid T_i = 0, \log w_1\right]}{\log w_2 - \log w_1} = \frac{E\left[\log w_2^{\text{prior}}\right] - E\left[\log w_1^{\text{prior}}\right]}{\log w_2 - \log w_1}, \quad (11)$$

which is exactly the prior-knowledge parameter θ_w from Section 3. The coefficient γ_w 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 wage information steepens the belief–wage relationship. Using the belief-updating rule from Equation (2), one can show that

$$\sigma_w = \alpha_w(1 - \theta_w). \quad (12)$$

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 $\hat{\alpha}_w = \hat{\sigma}_w/(1 - \hat{\theta}_w)$, measuring the share of the initial misperception that is corrected when wage information is provided.

Table 3 reports the results from estimating Equation (10). We begin with θ_w. The estimates from our preferred specification with occupation-region fixed effects, in Column (2), indicate that control-group job seekers are relatively good at reading relative pay differences from job-ad previews: a one-log-point increase in the kununu wage is associated with an increase of $\hat{\theta}_w = 0.46$ in expected pay beliefs. Reweighting the survey sample to match the occupational distribution of field-experiment impressions—both in the full sample and among impressions that displayed wage information—leaves the estimates essentially unchanged (see Appendix Table A17). 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 within their search scope.[19]

Next, we turn to the effect of providing wage information on own-wage beliefs. Table 3 shows that the treatment shifts both the level and the slope of wage beliefs. The coefficient on the wage-information indicator T_i^w implies that job seekers underestimate full-time monthly wages by 5.2% on average (Column (2)). This is smaller than the 18% underestimation of starting salaries that Andresen et al. [2026] document in a comparable survey of Norwegian university students, but it is consistent with our field-experimental finding that posting an average wage increases job seekers' interest in job ads (cf. Table 2).

Finally, we turn to the interaction term, which measures whether job seekers update their wage beliefs in response to the kununu wage. 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 $\hat{\theta}_w + \hat{\sigma}_w = 0.61$. The implied belief-updating parameter – $\hat{\alpha}_w = \hat{\sigma}_w/(1 - \hat{\theta}_w) \approx 0.29$ – means that wage information closes about 29% of the gap between job seekers' priors and the kununu benchmark.

[19] In 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 perceived on actual wage gaps of 0.54, 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.

The result that workers update wage beliefs only partially to wage information aligns with a recent survey and field-experimental literature that documents 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. Indeed, 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 greater belief updating. 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).

There are many potential explanations for incomplete updating: an average wage may not reflect the starting wage a job seeker may expect to be offered, or job seekers may pay only limited attention to wage information. A related mechanism is that job seekers may 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 this notion, 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 show that posting benchmark wages does not eliminate the strong link between past pay and wage beliefs.

Table 3: Belief-elicitation survey: Wage misperceptions and belief updating

	(1)	(2)	(3)	(4)	(5)
$\log(\widetilde{w_j})$	0.645***	0.456***	0.456***		0.406***
	(0.0252)	(0.0209)	(0.0207)		(0.0224)
T_i^w	0.0447***	0.0450***	0.0469***	0.0462***	0.0206
	(0.0114)	(0.0107)	(0.0105)	(0.0107)	(0.0143)
$\log(\widetilde{w_j}) \times T_i^w$	0.164***	0.156***	0.156***	0.155***	0.252***
	(0.0380)	(0.0363)	(0.0355)	(0.0360)	(0.0442)
$\log(\widetilde{lastwage_i})$			0.118***	0.118***	
			(0.0172)	(0.0176)	
Occupation*Region FE		Yes	Yes		Yes
Job FE				Yes	
Control + kununu info sample					Yes
$\hat{\alpha}_w$	0.463	0.287	0.287		0.424
Observations	11872	11872	11872	11848	6880
R-squared	0.283	0.327	0.346	0.471	0.315

Notes: The table shows regression estimates of Equation 10 based on data from the belief-elicitation survey. The dependent variable is the respondents' expected own log wage in a job. $\log(\widetilde{w_j})$ measures the average kununu wage in job title and firm posting vacancy j. T_i^w is a dummy variable equal to one if job seeker i was assigned to a treatment arm with wage information. $\log(\widetilde{lastwage_i})$ measures the most recent monthly full-time wage (for unemployed job seekers) or the current wage (for employed job seekers). Log-wages are recentered around the mean log wage in the field experiment. The estimate of the belief updating parameter $\hat{\alpha}_w$ is calculated as $\hat{\alpha}_w = \hat{\sigma}_w/(1-\hat{\theta}_w)$. Standard errors (in parentheses) are clustered by respondent.

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 and θ with the posted-wage elasticities, $\hat{\eta}_w^{\text{across}}$ and $\hat{\eta}_w^{\text{within}}$, to quantify job seekers' implied 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 $\hat{\beta}_w = 2.91$. 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 two- to three-times as large as these elasticities. Our within-ad estimate of 2.91 also aligns closely with complementary evidence from Christensen et al. [2026], who estimate an own-wage application elasticity of approximately 3.

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	0.836
	(0.040)	(0.226)
Belief-update factor	0.156	0.287
$(\hat{\sigma}_w \ / \ \hat{\alpha}_w)$	(0.036)	(0.059)
Own-wage elasticity $\hat{\beta}_w$	**2.51**	**2.91**
	(0.64)	(0.99)
95% CI	[1.53, 4.13]	[1.49, 5.67]

Notes: The table reports the implied own-wage elasticity of job applications (proxied by ad views), $\hat{\beta}_w$, computed by adjusting the posted-wage elasticity for the degree of prior misperception and the strength of belief updating identified by the belief-elicitation survey. Column 1 uses the across-ad posted-wage elasticity, $\hat{\eta}_w^{\text{across}}$, from Column 1 of Table 2 and the belief-update factor $\hat{\sigma}_w = \hat{\alpha}_w(1-\hat{\theta}_w)$ from Table 3, Column 2; the formula is $\hat{\beta}_w^{\text{across}} = \hat{\eta}_w^{\text{across}}/\hat{\sigma}_w$. Column 2 uses the within-ad posted-wage elasticity, $\hat{\eta}_w^{\text{within}}$, from Column 2 of Table 2 and the belief-updating parameter $\hat{\alpha}_w = \hat{\sigma}_w/(1-\hat{\theta}_w)$ from the same belief regression; the formula is $\hat{\beta}_w^{\text{within}} = \hat{\eta}_w^{\text{within}}/\hat{\alpha}_w$. Standard errors in parentheses are computed using the delta method, treating the field experiment and the wage-belief survey as independent sources of variation. The 95% confidence intervals are constructed on the log scale and back-transformed.

The findings in this section have two implications for the extent and sources of Swiss employers' labor market power. First, our estimates of the own-wage elasticity of job applications β_w, on the order of three, suggest that job-seeker preferences are an important source of firms' wage-setting power.[20] Second, our results suggest that incomplete information constitutes an additional source of labor market power. Even if workers value wages highly, firms may retain wage-setting power when workers do not fully perceive wage differences across jobs. Our estimates of θ suggest that workers underestimate pay differences by roughly half. While the canonical monopsony literature largely abstracts from worker-level misperceptions [Azar and Marinescu, 2024], a small but growing literature has begun to study the role of imperfect information

[20] In standard models of monopsony with differentiated jobs, workers trade off wages against non-wage amenities and idiosyncratic job preferences [Azar and Marinescu, 2024], and firms' wage-setting power depends on the elasticity of labor supply to the firm—equivalently, the responsiveness of applications to wages.

on wages in shaping labor market power [Cheremukhin and Restrepo-Echavarria, 2025, Christensen et al., 2026]. Our findings provide experimental evidence that incomplete information is quantitatively important.

5 Job Benefits

We next turn to job seekers' valuation of benefits when searching for jobs. As with wages, the analysis proceeds in three steps. Section 5.1 uses the platform 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 derives estimates of job seekers' willingness to pay (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 $T_{s(i)}^k$—whether job seeker i was informed about benefit k in session s—and the firm-benefit indicators $b_{s(i)j}^k$, which equal 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 search fixed effects, and 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—work from home, flexibility, and childcare—we additionally report a within-ad estimate, η_k^{within}, based on Equation 8. This specification restricts the sample to job seekers who received benefit information 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. Appendix Table A13 shows that the estimated benefit effects are highly stable across a range of alternative regression specifications.[21]

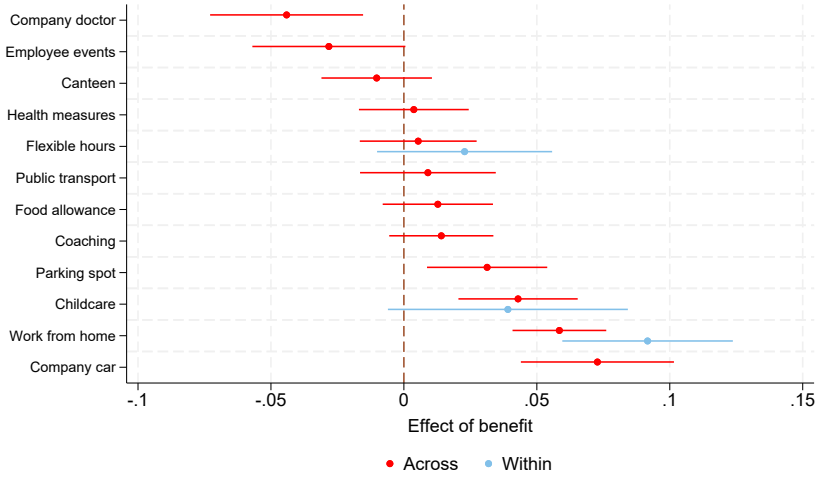
The estimates suggest that some job benefits play an important role in job search. The largest effect is for the provision of a company car, which raises job seekers' probability of viewing an ad by 7.3%. The opportunity to work from home (+6.0%), firm-sponsored childcare (+4.4%), and a parking spot (+3.1%) also generate economically and statistically significant increases in ad views. 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, and for the company doctor this reduction is

[21] Appendix Table A13 reports the across-ad benefit estimates for four alternative specifications: dropping session fixed effects, using a fixed-effects Poisson model, excluding very short sessions, and restricting the sample to job seekers always assigned to the same treatment arm.

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, suggesting that job seekers have relatively little prior knowledge about which firms in their search scope offer which benefits. While this conclusion is tentative given the limited precision of the within-ad estimates, it is confirmed by the direct estimates of θ_b derived from the belief-elicitation survey reported below.

Figure 5: Effect of job benefits on job seekers' views of job ads



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 of whether an ad was viewed. 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 estimating Equation (8) for the three benefits with experimentally varied availability thresholds (20% vs. 50%). All specifications include ad, rank, and session fixed effects. Whiskers denote 95% confidence intervals; standard errors are clustered at the user-by-month level. Coefficients are scaled by the control-group mean of the ad-view rate (0.039), so values can be read as percentage changes relative to that baseline. The corresponding coefficient estimates are reported in Table A13, Column (1).

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[22], similar across terciles of the occupational wage distribution. These results run counter to the notion that lower-paid workers react less to job benefits than higher-paid workers.

Panel (b) reports estimates split by the female employment share in the occupation. The clearest evidence for heterogeneity arises for the opportunity to work from home, which has a significantly larger effect on ad views in female-dominated occupations. This pattern is consistent with prior work suggesting that women may place higher value on workplace flexibility [Escudero et al., 2024, Mas and Pallais, 2017, Wiswall

[22]The clearest exception is a canteen, whose effect on ad views is negative and statistically significant in high-paying occupations, suggesting that job seekers in these occupations perceive it as a disamenity. The company car has its largest effect in middle-paying occupations, and childcare a somewhat larger effect in lower-paying occupations, though these and the remaining differences across terciles are mostly small and statistically insignificant.

and Zafar, 2018]. However, most benefits show little systematic evidence for gender-related heterogeneity, including flexible working hours and childcare.

Panel (c) splits the sample by how commonly each benefit is offered within an occupation.[23] 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 a firm-level amenity by construction, work-from-home may be available only to white-collar workers. Indeed, the work-from-home effect is substantially 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 weaker, or the effect is larger when the benefit is uncommon. We read these results as broadly confirming our approach of treating benefits as firm-level amenities.

5.2 Misperceptions and Belief Updating

The conceptual framework implies that the across-ad benefit estimates 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). Using the belief-elicitation survey, we estimate θ_b^k and α_b^k based on the following regression:

$$b_{ij}^{\text{post},k} = \theta_b^k b_j^k + \gamma_b^k T_i^b + \sigma_b^k b_j^k \times T_i^b + \varepsilon_{ij}. \tag{13}$$

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 kununu indicator of benefit availability. T_i^b equals one if 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. The belief-updating parameter is $\hat{\alpha}_b^k = \hat{\sigma}_b^k/(1-\hat{\theta}_b^k)$. We estimate Equation (13) with and without controls for the wage-information treatment and the kununu wage.

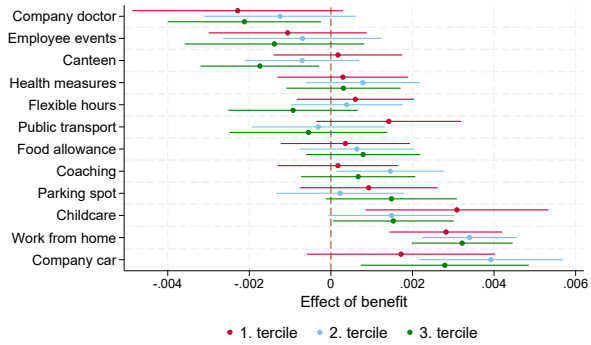
Table 5 reports our baseline estimates of Equation (13) separately for the three benefits. We focus on θ_b first. Column 1 shows that respondents in the control group rate the probability that a firm offers flexible working hours about 7.7 percentage points higher ($\hat{\theta}_b = 0.077$) when the 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 $\hat{\theta}_b = 0.148$, and the estimate for childcare (Column 5) is $\hat{\theta}_b = 0.138$. These values are several times smaller than the wage-misperception 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 consistent with the indirect evidence in Figure 5 that the across-ad and within-ad benefit estimates are similar.

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

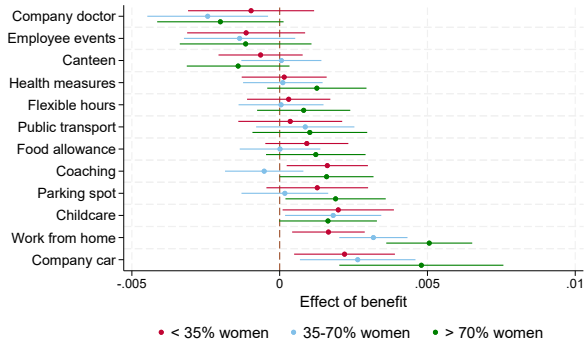
[23] Ideally, we would split occupations into those where a benefit (e.g., work-from-home) is essentially absent and those where a non-negligible share of workers receive it. However, the employer review data do not record reviewers' occupations, so we cannot directly observe the within-firm distribution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

Figure 6: Heterogeneity of the effect of job benefits on job seekers' views

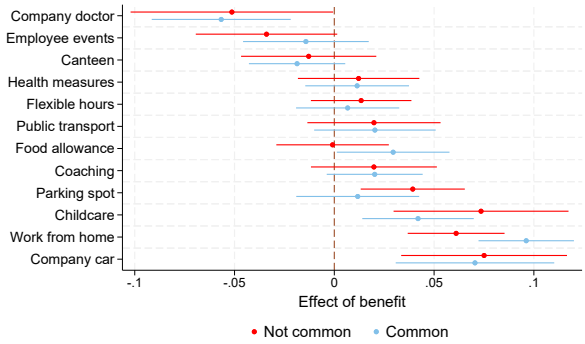
(a) By wage level of occupation



(b) By female share in occupation



(c) By occupational prevalence of benefit



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. Coefficients are scaled by the control-group mean of the ad-view rate. Specifications include ad, rank, and session fixed effects. Whiskers denote 95% confidence intervals, and standard errors are clustered at the user-by-month level. Panel (a) splits the across-ad estimates by tercile of the occupational wage level (median wage in the ISCO 3-digit occupation), and Panel (b) splits them by the female employment share in the ISCO 3-digit 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: we observe the share of kununu reviewers confirming a benefit among firms posting vacancies in each occupation, and then average this share across vacancies in an ISCO 2-digit occupation. We classify the occupation as "common" (blue) or "not common" (red) relative to the unweighted mean across occupations.

Table 5: Belief-elicitation survey: 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	-0.0246	-0.0851***	-0.0836***	-0.148***	-0.156***	-0.110***	-0.110***
	(0.0245)	(0.0258)	(0.0157)	(0.0165)	(0.0184)	(0.0204)	(0.0142)	(0.0152)
flx_j	0.0767***	0.0688***						
	(0.0156)	(0.0157)						
$flx_j \times T_i^b$	0.215***	0.214***						
	(0.0266)	(0.0268)						
wfh_j			0.148***	0.125***				
			(0.0106)	(0.0107)				
$wfh_j \times T_i^b$			0.240***	0.245***				
			(0.0199)	(0.0201)				
$chld_j$					0.138***	0.130***		
					(0.0125)	(0.0125)		
$chld_j \times T_i^b$					0.286***	0.288***		
					(0.0244)	(0.0244)		
b_j							0.0892***	0.0796***
							(0.00864)	(0.00868)
$b_j \times T_i^b$							0.275***	0.276***
							(0.0182)	(0.0183)
$\widetilde{\log(w_j)}$		0.206***		0.327***		0.182***		0.248***
		(0.0314)		(0.0289)		(0.0307)		(0.0208)
T_i^w		-0.0156		-0.00698		0.0207		-0.000624
		(0.0169)		(0.0126)		(0.0192)		(0.0117)
$\widetilde{\log(w_j)} \times T_i^w$		-0.0159		-0.0934**		-0.0251		-0.0488
		(0.0467)		(0.0367)		(0.0538)		(0.0323)
Occ*Region FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
$\hat{\alpha}_b$	0.233	0.229	0.282	0.280	0.332	0.331	0.302	0.300
Observations	11872	11872	11872	11872	11872	11872	35616	35616
R-squared	0.119	0.124	0.278	0.288	0.0976	0.101	0.121	0.127

Notes: The table shows regression estimates of Equation 13, based on data from the belief-elicitation survey. The dependent variables are respondents' beliefs about the availability of flexible working hours (Columns 1 and 2), work from home (Columns 3 and 4), and firm-sponsored childcare (Columns 5 and 6). In columns 7 and 8, we pool observations for the three benefits in a single regression. T_i^w and T_i^b are dummy variables equal to one if job seeker i was assigned to a treatment arm with wage information or benefit information, respectively. flx_j, wfh_j, and $chld_j$ indicate the availability of flexible working hours, work from home, and firm-sponsored childcare, according to kununu data in the firm posting vacancy j. $\widetilde{\log(w_j)}$ measures the average kununu wage for the job title and firm posting vacancy j in logs, recentered around the mean log wage in the field experiment. The estimate of the belief updating parameter $\hat{\alpha}_b$ is calculated as $\hat{\alpha}_b = \hat{\sigma}_b/(1-\hat{\theta}_b)$. Standard errors (in parentheses) are clustered by respondent.

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

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—willingness-to-pay (WTP) for benefit k. We maintain the assumption throughout—broadly supported using the comparable estimates of $\hat{\theta}_b$ and $\hat{\alpha}_b$ across benefits 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

$$\text{WTP}^k \;=\; \frac{\beta_b^k}{\beta_w} \;=\; \frac{\eta_k^{\text{across}}}{\eta_w^{\text{across}}} \cdot \frac{\alpha_w\,(1-\theta_w)}{\alpha_b\,(1-\theta_b)}. \tag{14}$$

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. The highest-valued benefit is the provision of a company car, with a willingness-to-pay 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]

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). According to responses to the belief-elicitation survey, this estimate should be interpreted 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].

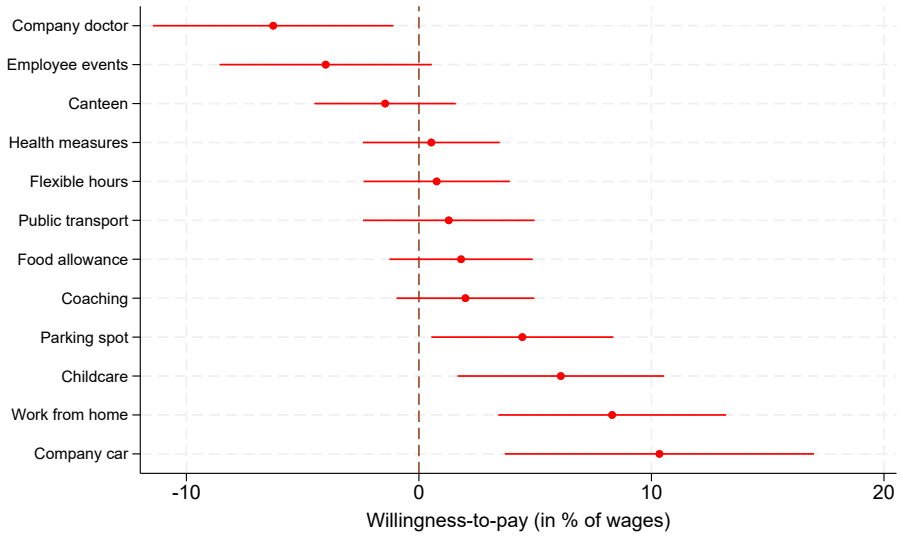
The third-most-valued benefit is firm-sponsored childcare, with a WTP of 6.1% of the wage. The belief-elicitation survey responses suggest that this estimate should be interpreted as an approximately 50% childcare subsidy rather than fully employer-financed or on-site care.[26] A back-of-the-envelope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

[24] All-inclusive leasing rates, including insurance, tax, and maintenance, amount to approximately CHF 500–700 per month for compact-to-midsize models such as a VW Golf or Skoda Octavia.

[25] In 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.

[26] We presented survey respondents the kununu childcare icon and asked what share of the costs of external childcare they

Figure 7: 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 $\widehat{\text{WTP}}^k = (\hat{\eta}_k^{\text{across}}/\hat{\eta}_w^{\text{across}}) \cdot (\hat{\sigma}_w/\hat{\sigma}_b)$ per Equation (14). 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), combining the joint covariance of $(\hat{\eta}_k^{\text{across}}, \hat{\eta}_w^{\text{across}})$ within the field experiment and the joint covariance of $(\hat{\sigma}_w, \hat{\sigma}_b)$ within the belief survey, treating the two experiments as independent because they are run on different samples.

gests that this WTP is approximately consistent with the cost savings implied by a 50% subsidy on childcare costs in Switzerland.[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 hours. This null result is consistent with Mas and Pallais [2017], who find that a majority of workers do not value schedule flexibility, but at odds with the sizable WTP in Wiswall and Zafar [2018] and Maestas et al. [2023]. A likely reason 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]

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%. A further 15.9% selected "Other," and their free-text responses frequently described the icon as indicating a company-run, on-site daycare.

[27] The OECD reports that the typical net out-of-pocket cost for a two-earner couple for two children in full-time center-based childcare, after subsidies and tax rebates, is around 35% of average earnings [OECD, 2022]. Combining the net cost of 35% of average earnings with the perceived 50% subsidy rate implies a benefit worth 17.5% of earnings. Among the survey respondents, 32.9% are between 25 and 40 years old. We treat this as an approximation to the share of job seekers for whom the benefit is potentially relevant. Multiplying the 17.5% average valuation by 32.9% implies that the benefit is worth 5.8% of the wage, close to our WTP estimate of 6.1%.

[28] When 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."

6 Are Wages a Signal for Other Job Attributes?

In many settings, wage information is more readily available than information on benefits. For example, many US states and the 2026 EU Pay Transparency Directive mandate increased transparency regarding salaries, but not for benefits. In addition, many job board platforms have started displaying wage estimates on vacancies. Do job seekers in such settings use the available wage information to update their perceptions about other job aspects? This possibility has long been recognized as a potential source of bias in the stated-preference literature and in settings such as ours [Adrjan 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:

$$\log w_{ij}^{\text{post}} = \alpha_w \log w_{ij} + (1 - \alpha_w) \log w_{ij}^{\text{prior}}, \tag{15}$$

$$b_{ij}^{\text{post}} = \alpha_{wb}(\log w_{ij} - \log w_{ij}^{\text{prior}}) + (1 - \alpha_{wb}) b_{ij}^{\text{prior}}. \tag{16}$$

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 ($\log w_{ij}$). In contrast to the previous framework, the belief updating for benefits is now driven by wage information rather than the benefit information. The parameter α_{wb} captures belief spillovers from wage information to beliefs about benefits. An α_{wb} equal to zero implies that job seekers do not use the wage information shown to update their beliefs on the availability of a certain benefit. In contrast, a positive α_{wb} implies that job seekers expect a higher likelihood of the benefit being available when the displayed wage exceeds what they expected ($\log w_{ij} - \log w_{ij}^{\text{prior}} > 0$). A negative α_{wb} implies that a higher-than-expected wage lowers the expected likelihood that the benefit will be available.

Our experimental setup allows us to quantify the extent of belief spillovers, α_{wb}, by leveraging the fact that, in both the field and survey experiments, we had a wage-only information treatment arm. In the survey experiment, the wage-only group was surveyed about their beliefs on the availability of flexible working hours, work from home, firm-sponsored childcare, and a company car. In addition, a random subset of respondents was asked about their beliefs regarding work atmosphere and the probability of getting the job. Thus, we can directly measure α_{wb} by running the following regression:

$$b_{ij}^{\text{post},k} = \theta_w^k \widetilde{\log(w_j)} + \gamma_w^k T_i^w + \sigma_w^k \widetilde{\log(w_j)} \times T_i^w + \gamma^k b_j^k + \epsilon_{ij} \tag{17}$$

where the notation is as before. The coefficient σ_w^k captures the degree of spillovers from wage information to beliefs about the provision of benefit k, corresponding to α_{wb} in our conceptual framework.

Table 6 presents the results of this regression equation for different benefits, and for different fixed effect specifications. All point estimates are close to zero, and for all non-wage job attributes, we cannot reject the null hypothesis of zero cross-belief updating at conventional confidence levels. This is true even when we pool all benefits to maximize statistical precision in detecting spillovers (Column 9). Thus, the wage level shown to individuals did not meaningfully change their beliefs about the availability of benefits. In addition, we find no evidence that the posted wage affected respondents' expected probability of receiving

Table 6: Spillovers of wage info on benefits belief updating

	(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}
$\widetilde{\log(w_j)}$	0.236***		0.289***		0.230***		0.253***		0.263***	
	(0.0368)		(0.0345)		(0.0353)		(0.0314)		(0.0243)	
T_i^w	-0.0180	-0.0193	-0.00916	-0.0105	0.0385	0.0381	0.0317*	0.0321*	0.00379	-0.0900
	(0.0225)	(0.0232)	(0.0167)	(0.0172)	(0.0282)	(0.0292)	(0.0170)	(0.0176)	(0.0161)	(0.122)
$\widetilde{\log(w_j)} \times T_i^w$	-0.0218	-0.00418	-0.0186	0.00309	-0.0331	-0.0305	0.0644	0.0555	-0.0247	0.513
	(0.0673)	(0.0689)	(0.0539)	(0.0531)	(0.0867)	(0.0913)	(0.0587)	(0.0598)	(0.0494)	(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									0.0739***	
									(0.00866)	
Occ*Region FE	Yes		Yes		Yes		Yes		Yes	
Job FE		Yes		Yes		Yes		Yes		Yes
Observations	6960	6917	6960	6917	6960	6917	6960	6917	20880	2294
R-squared	0.0979	0.215	0.288	0.469	0.0733	0.128	0.189	0.306	0.0868	0.213

Notes: The table shows results of a regression including data from the control and wage-only treatment arms in the belief-elicitation survey. The dependent variable is the respondents' belief about the availability of flexible working hours in columns 1 and 2, of work from home in columns 3 and 4, of firm-sponsored childcare in columns 5 and 6, and of a company car in columns 7 and 8, in vacancy j. In column 9, we pool observations for the three benefits in a single regression. The dependent variable in the last column is the respondents' belief about the workplace atmosphere in the firm posting the vacancy (measured on a scale from 0 to 7). $\widetilde{\log(w_j)}$ measures the average kununu wage for the job title and firm posting vacancy j, in logs. T_i^w is a dummy variable equal to one if job seeker i was assigned to a treatment arm with wage information. flx_j, wfh_j, $chld_j$, and car_j indicate the availability of flexible working hours, work from home, firm-sponsored childcare, and a company car, according to kununu data, in the firm posting vacancy j. All log wages are recentered around their mean log wage. Standard errors (in parentheses) are clustered by respondent.

a job offer or expected wage trajectories (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.[29]

Our field experiment provides a second, independent test of belief spillovers. If job seekers in the wage-only arm used the wage to infer the availability of benefits, the wage could also influence beliefs about benefit availability. Formally, the within-ad elasticity in the wage-only arms equals $\alpha_w\beta_w + \alpha_{wb}\beta_b$, and therefore coincides with the posted wage elasticity in the wage-and-benefit arms only in the absence of belief spillovers ($\alpha_{wb} = 0$) or if benefits are not valued ($\beta_b = 0$) (see Appendix F for a formal derivation). Appendix Table A20 shows that we cannot reject the null hypothesis $\eta_w = \eta_{wb}$ at conventional significance levels in any specification. These results, too, suggest limited spillovers from wages to benefits.

Although we cannot exclude the possibility that wages shape perceptions of other, unobserved job attributes, both tests consistently fail to detect belief spillovers from wages to the job attributes included in our experiment. These findings are consistent with the limited evidence for belief spillovers in Christensen et al.

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

[2026]. They contrast, however, with the belief spillovers from posted wages in hypothetical ads documented by Adrjan et al. [2026]. One difference between the settings is that in their discrete choice experiment, where the employer's name is not shown, changes in posted wages or benefits may shift beliefs about the employer behind the ad. In our experiments, job seekers evaluate real vacancies from named firms, holding the employer fixed while varying the information.

7 Implications for Job Inequality

A central question in labor economics 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. 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 this section, we combine our willingness-to-pay estimates from Section 5.3 with data on wages and benefit availability across jobs to quantify the extent to which non-wage benefits affect inequality in the Swiss labor market.

To move from wage inequality to job-value inequality, we construct a money-metric measure of total compensation that aggregates a job's pay and the monetized value of all benefits offered by the firm. Specifically, for each job vacancy j on the Jobchannel websites that we were able to match to wage and benefit information, we define total job value as:

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

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, recovered from the revealed-preference analysis in Section 5.3. The summation includes all twelve benefits studied in the paper, including the small number with negative estimated valuations.

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 posted wage and of total job value across job advertisements in the matched experimental sample. Total job value equals the wage multiplied by one plus the summed value of the benefits the job offers, where each of the twelve benefits is valued at its estimated willingness to pay as a share of the wage (including the few benefits with negative estimated valuations), and a benefit is counted as available when at least 20% of a firm's reviewers confirm it. P90/P10 and P50/P10 are the ratios of the corresponding percentiles, and Gini is the Gini coefficient.

Table 7 compares the distribution of wages and total job value across the 36,504 job vacancies in the matched experimental sample. 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, and more importantly for inequality, the dispersion of total job value

exceeds the dispersion of wages by every measure reported. For instance, the standard deviation increases from CHF 2,109 to CHF 2,456—a rise of approximately 16 percent. The Gini coefficient rises from 0.172 to 0.182, an increase of roughly 6 percent.

Why is job value inequality higher than wage inequality? As we show in Appendix Figure A11, higher-paying jobs offer more benefits. These results are consistent with a growing body of evidence that amenities and wages tend to co-move across jobs and firms, so that accounting for non-wage compensation amplifies rather than compresses measured inequality [e.g., Caldwell et al., 2026, Costas-Fernández et al., 2026, Lamadon et al., 2022, Ouimet and Tate, 2023, Sockin, 2024, Sorkin, 2018]. The magnitude of our estimates is consistent with Maestas et al. [2023], who also value a set of observed amenities and find that the 90–10 differential in total compensation is about 11% larger than the corresponding wage differential. In contrast, structural approaches that infer amenities from job transitions or offer data 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 complementary belief-elicitation survey to 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.

We are able to provide several key insights.

First, revealing wage information has a meaningful effect on job search behavior: a 1% increase in the posted wage raises the likelihood that a job seeker views or clicks on an advertisement by approximately 0.4%. However, this posted-wage elasticity substantially understates job seekers' preferences for wages relative to other job aspects.

Using the survey experiment, we document that job seekers are already partially informed about wage differences across vacancies and update their own-wage expectations only partially in response to posted wages. This implies that job seekers' elasticity of applications with respect to their own expected wages is between 2.5 and 3.

Second, we estimate job seekers' willingness to pay for job benefits using a revealed-preference approach. Among the twelve benefits studied, a company car and work from home are the highest-valued benefits. Job seekers are willing to forgo 10% of their wage for a company car and 8% of their wages to work from home. The estimated willingness to pay is approximately 6% of wages for firm-sponsored childcare and 4% for a parking spot. Overall, job benefits constitute an important component of total compensation. In our sample, the average job offers benefits whose value amounts to approximately 11% of wages.

Third, we find no evidence of economically meaningful belief spillovers: wage information does not systematically shift job seekers' beliefs about job benefits and other non-wage job attributes such as the probability of receiving a job offer or workplace atmosphere.

Fourth, we document important implications for inequality. Because higher-paying firms tend to offer more and better benefits, inequality in total job value—defined as the sum of wages and the monetized value of benefits—is significantly larger than inequality in wages alone.

References

Girum Abebe, A. Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, and Simon Quinn. 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. Technical report, University of Bonn and University of Cologne, Germany, 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. arXiv:2407.13204, 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. URL `https://www.iza.org/publications/dp18220`.

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 Haegele, and Jörg Heining. Why workers stay: Pay, beliefs, and attachment. NBER Working Paper 33445, National Bureau of Economic Research, 2025.

Sydnee Caldwell, Ingrid Haegele, and Jörg Heining. Firm pay, amenities, and inequality. Working paper, UC Berkeley and LMU Munich, 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? Working Paper; February 12, 2026, 2026.

Julián Costas-Fernández, Sebastian Findeisen, Anna Raute, and Uta Schönberg. Family-friendly workplace policies. Working paper, 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 Thákral, and Linh T. Tô. Wage differentials and the price of workplace flexibility. Working paper, Boston University, January 2025. URL `https://linh.to/files/papers/amenity.pdf`.

Verónica Escudero, Hannah Liepmann, and Damián Vergara. Directed search, wages, and non-wage amenities: Evidence from an online job board. IZA Discussion Paper 17211, Institute of Labor Economics (IZA), Aug 2024.

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. Technical report, National Bureau of Economic Research, 2023.

Jason Sockin. Show me the amenity: Are higher-paying firms better all around? 2024.

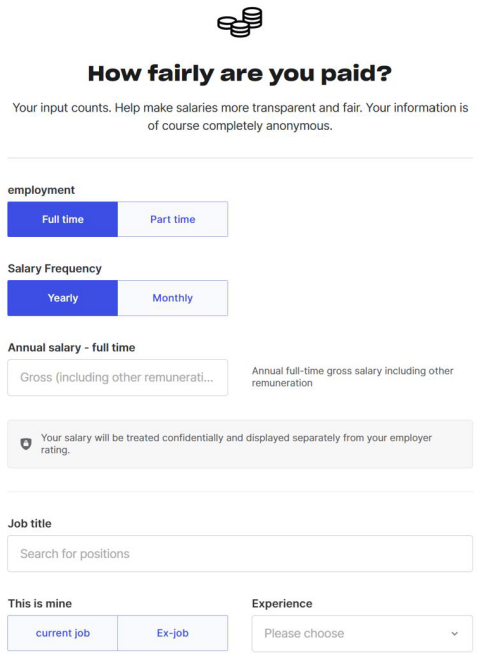
Isaac Sorkin. Ranking firms using revealed preference. *The quarterly journal of economics*, 133(3):1331–1393, 2018.

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.

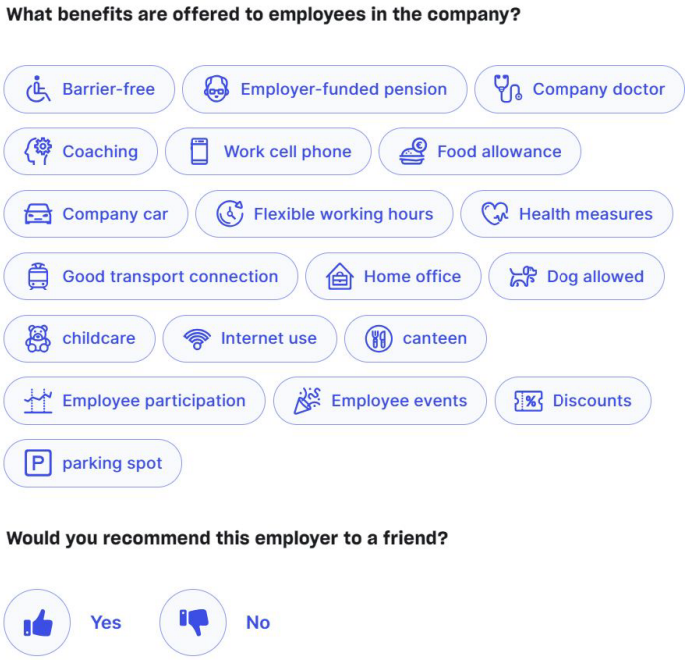
A Further Figures

Figure A1: Snapshots of kununu Employer Review Forms

(a) Wage Reviews

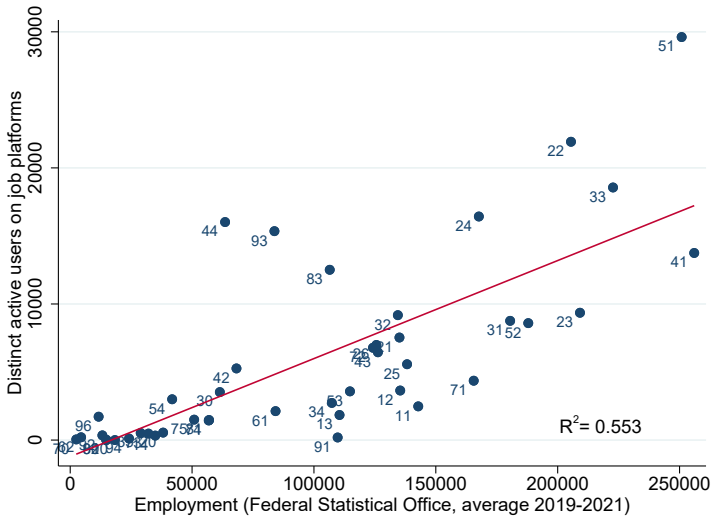


(b) Benefits Review



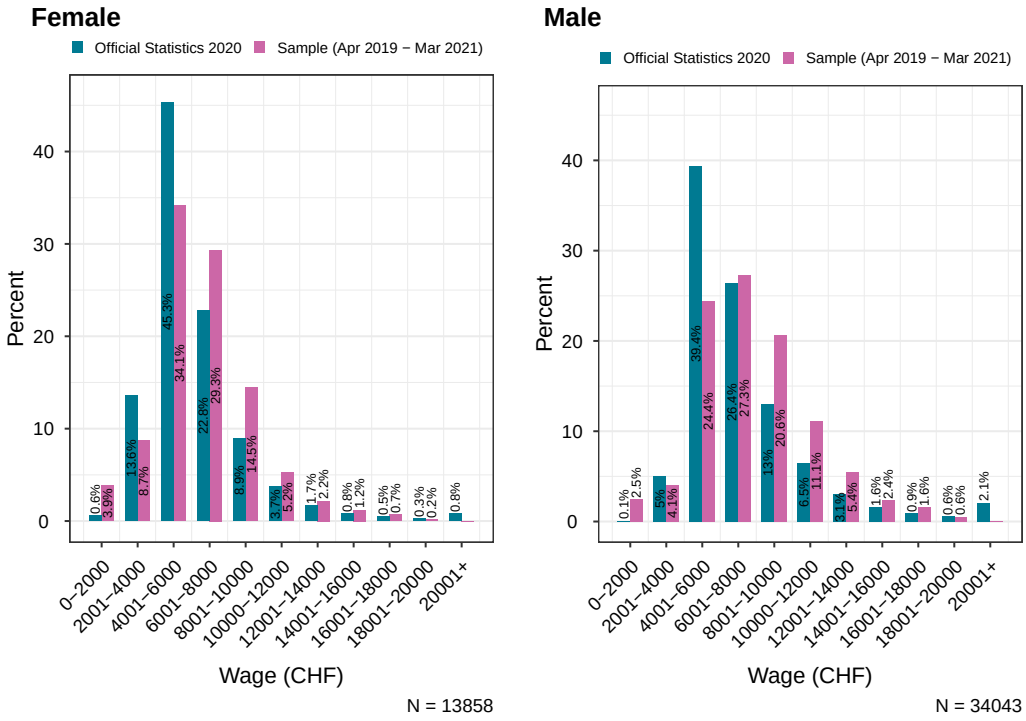
Notes: The figure shows snapshots of the kununu wage and benefit review pages. Workers can choose either a monthly or annual pay frequency to report their wages. Since one or two additional monthly payments are common in Switzerland, reviewers who select the monthly payment are asked about the number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

Figure A2: Users in Experiment vs. Employment in Switzerland by ISCO 2-Digit Occupations



Notes: The figure plots the number of users in the analysis sample against employment in Switzerland by ISCO 2-digit 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 net monthly wages of full-time employees in 2020 according to the Swiss Earnings Structure Survey (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: Screenshot of Result List for Treatment Group: Desktop Users

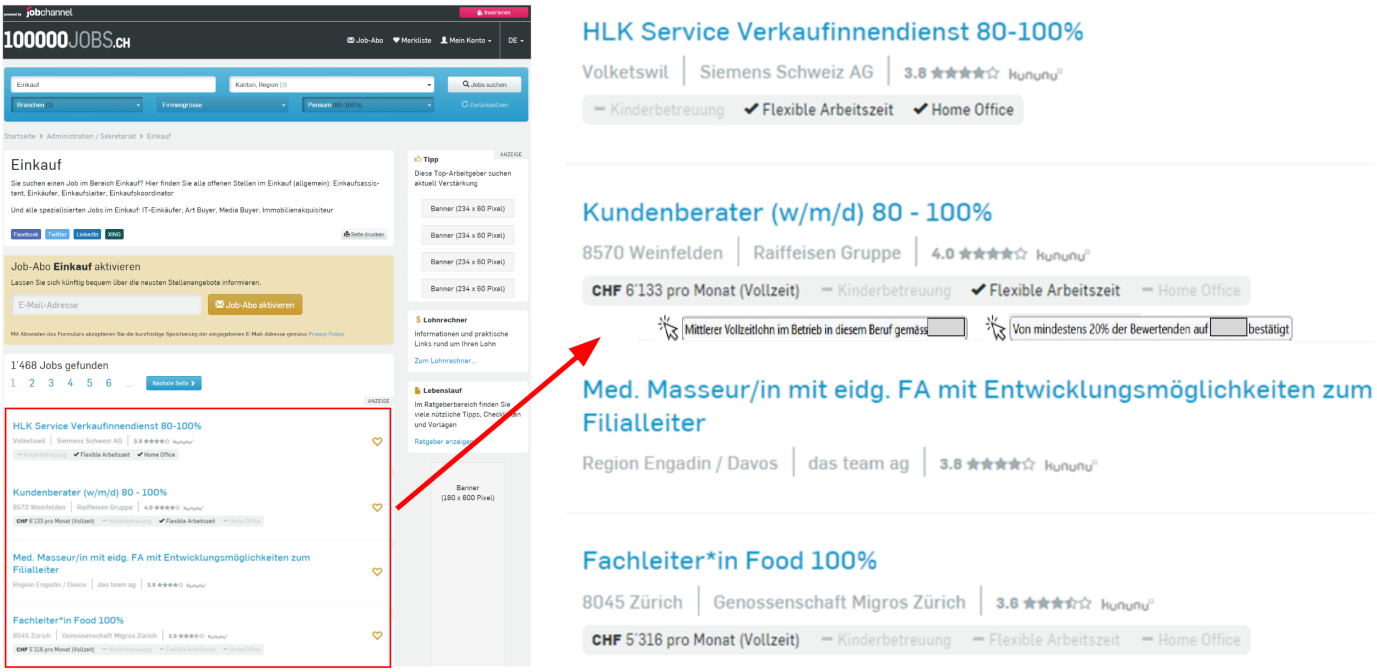


Figure A5: Screenshot of Single Ad View for Treatment Group: Desktop Users

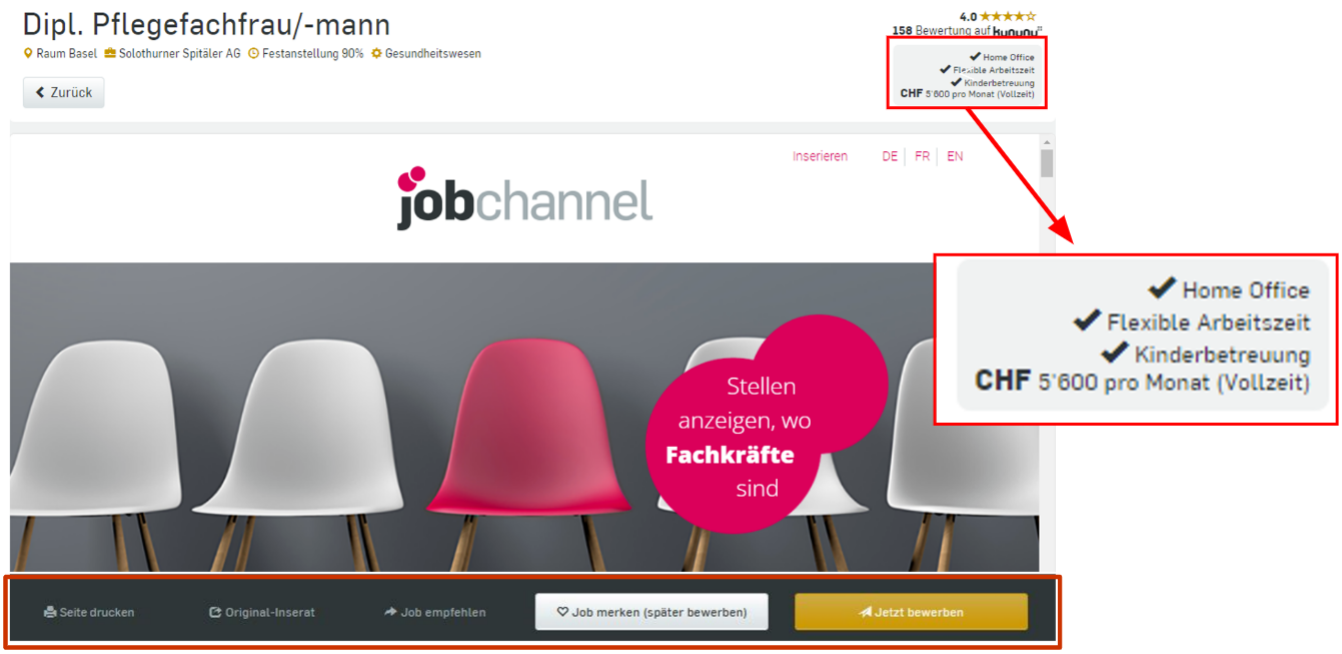
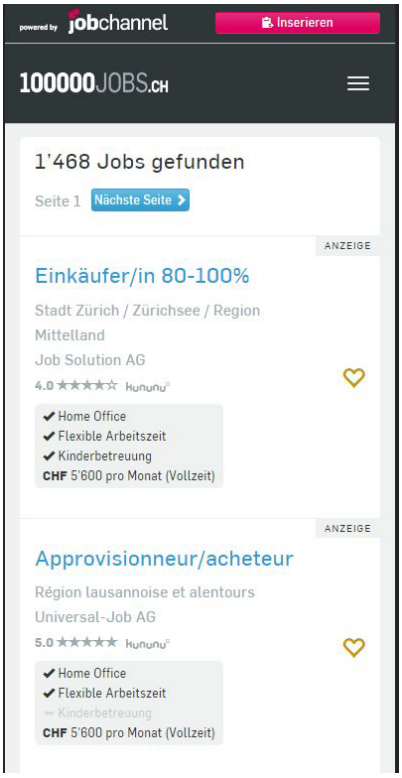


Figure A6: Screenshots of the Experiment for Treatment Group: Mobile Users

(a) Result List



(b) Single Ad View

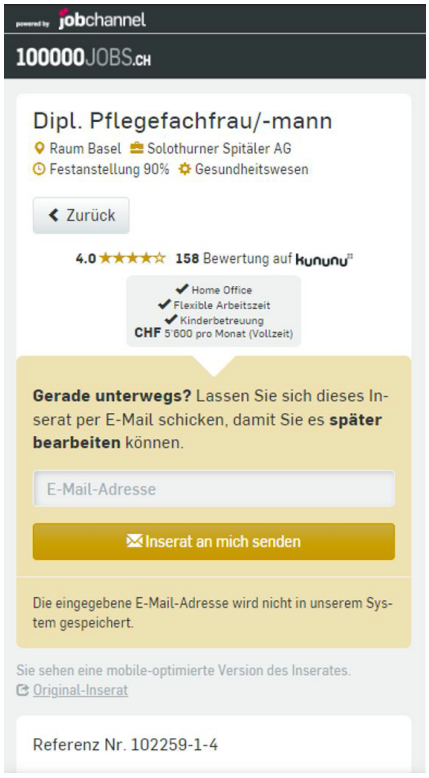
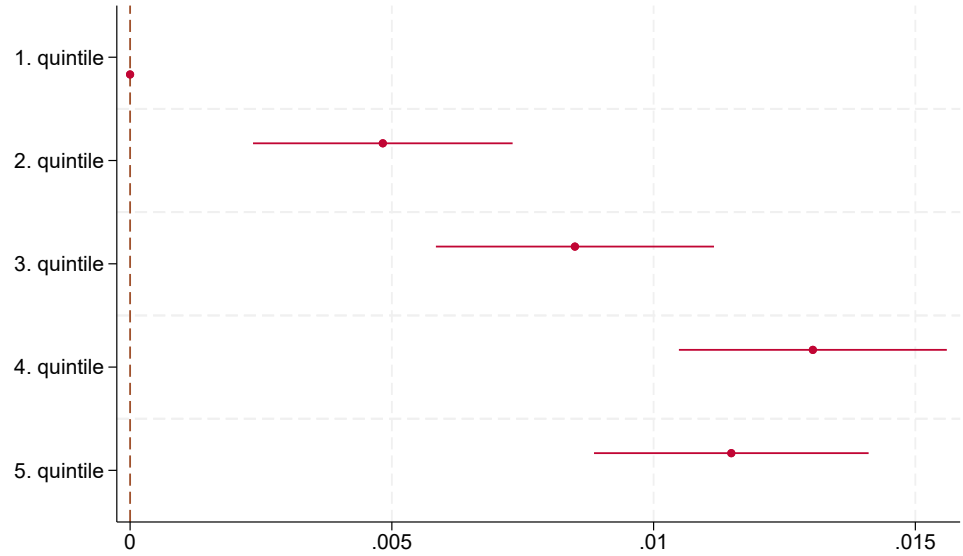
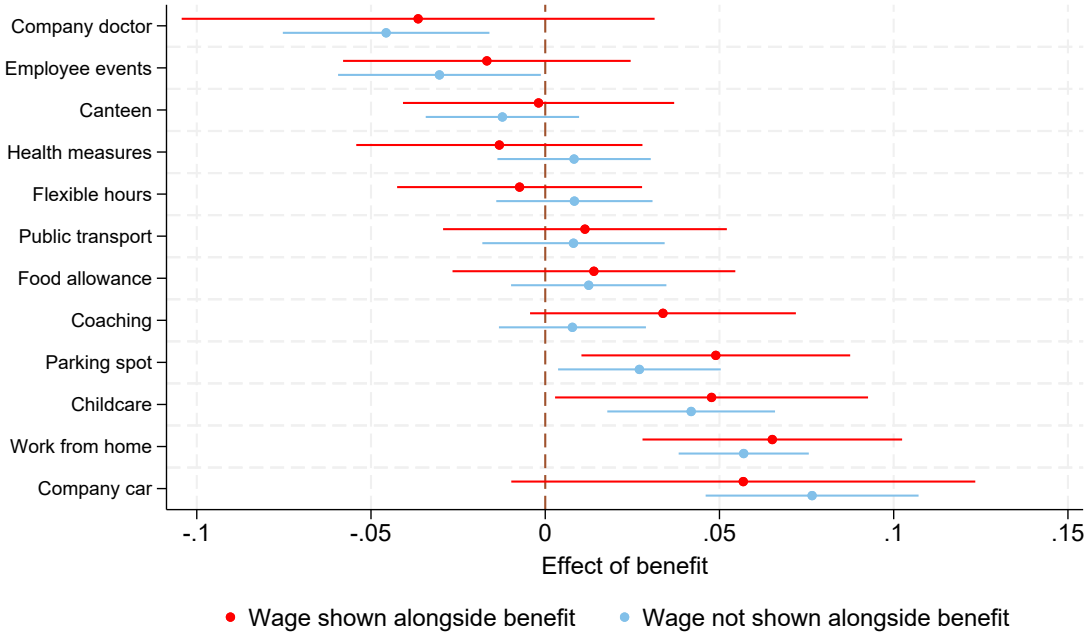


Figure A7: Effect of the posted wage on ad views by quintiles of wages



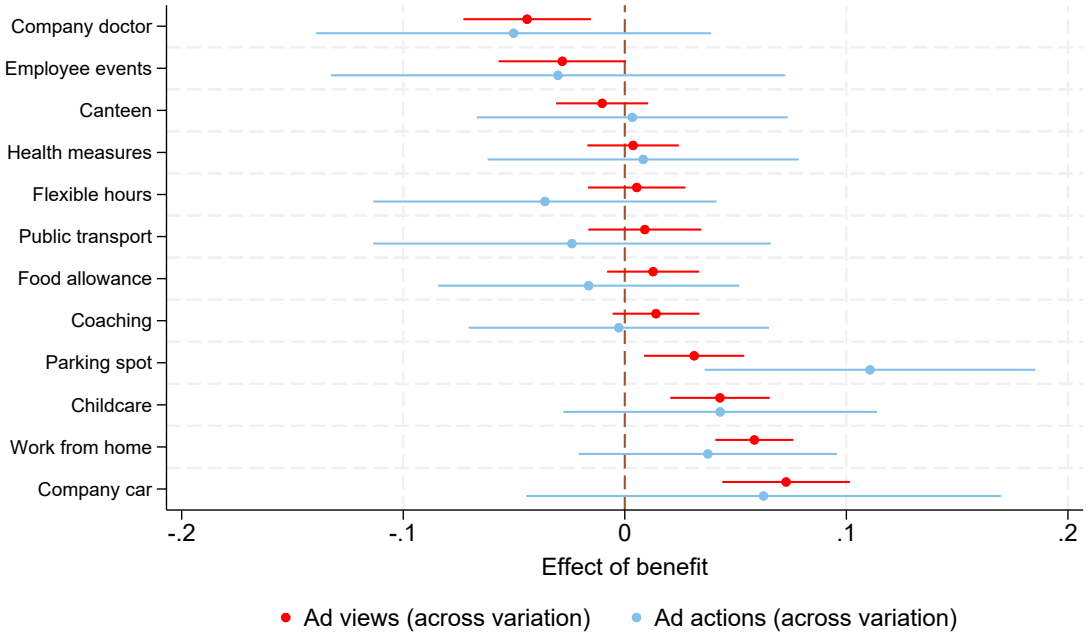
Notes: The figure shows the effect of wages on the likelihood of viewing an ad by quintiles of the (recentered) log wage. The regression model, sample, and dependent variable are the same as in Column 1 of Table 2 except for the wage variables.

Figure A8: Effect of job benefits on job seekers' views of job ads, by whether the wage was shown alongside the benefit



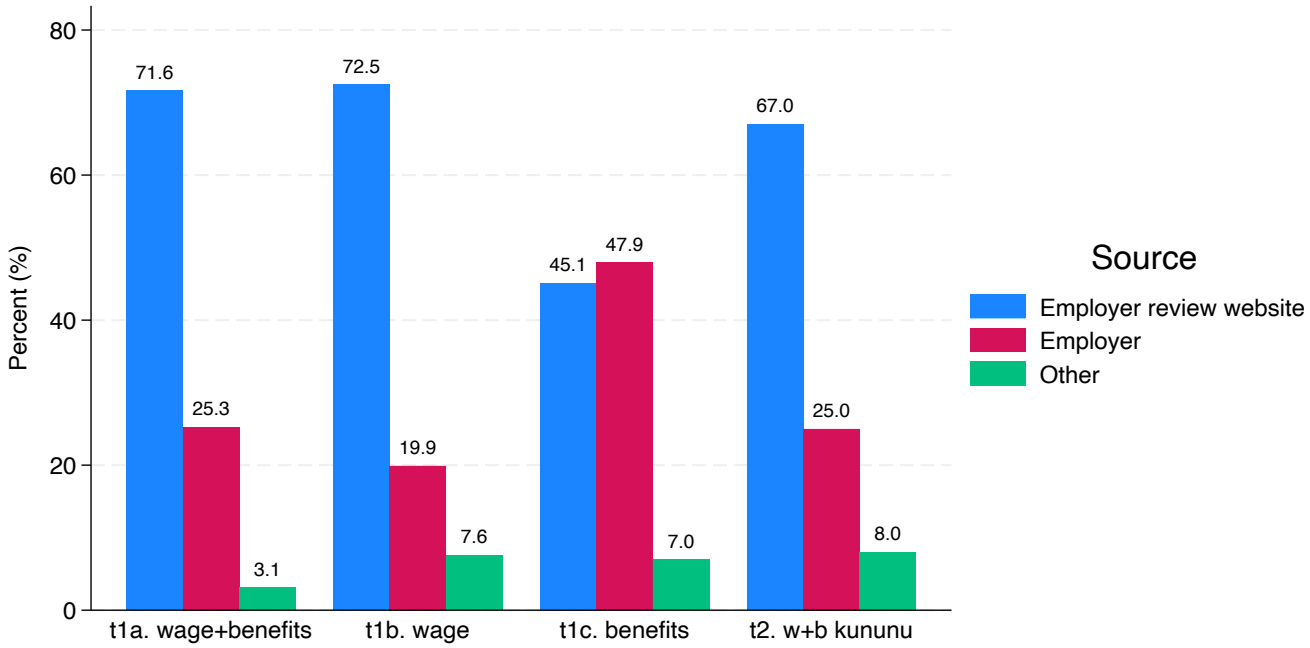
Notes: The figure shows reduced-form estimates of the effect of each benefit on the probability of viewing a job ad, separately for impressions in which the wage was displayed alongside the benefit (in red) and impressions in which the wage was not displayed (in blue). The latter cases arise within the wage-and-benefits treatment arms when the ad lacks kununu wage information, so that benefits are still displayed but the wage is not. The two cell effects are estimated jointly in a single across-ad regression on the same sample as Figure 5 (excluding the 50%-threshold treatment arms). For each benefit k, the regression includes the interaction $b^k_{s(i)j} \times T^k_{s(i)} \times T^w_{s(i)j}$ (red series) and $b^k_{s(i)j} \times T^k_{s(i)} \times (1 - T^w_{s(i)j})$ (blue series). All coefficients are normalized by the mean ad-view rate in the control group.

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. The figure shows that the ranking of benefits and the relative magnitudes of the effects are similar across the two engagement margins.

Figure A10: Belief about information source by treatment arm

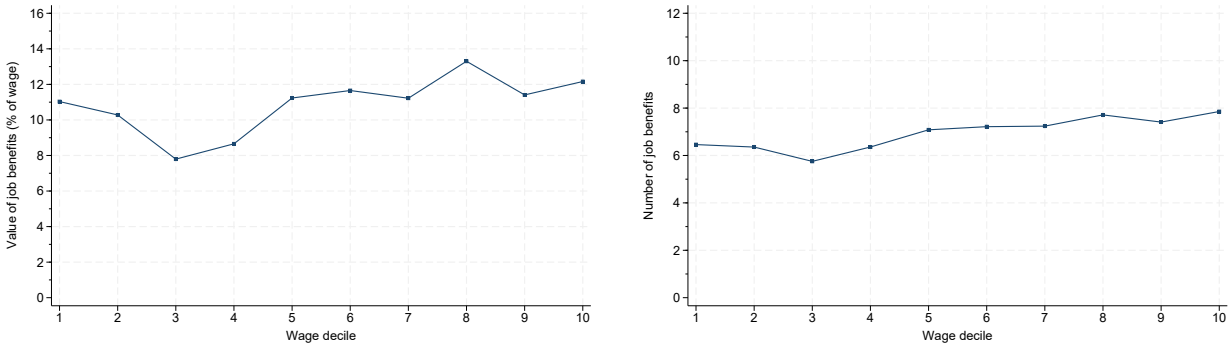


Notes: The figure displays the percentage of respondents who attributed the information to each source category, by treatment arm.

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 posted (kununu firm–occupation) wage. 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. A benefit is counted as available when at least 20% of a firm's reviewers confirm it.

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: Correspondence between work-from-home information from vacancies' text and kununu

Vacancy text	Kununu	Number of job ads	FNR
Review threshold: 20%			
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
Review threshold: 50%			
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
Likely telecommutable, Thr: 20%			
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 A3: Balance Tests for User's characteristics: March

	0	1	2	3	4	5	6	7	p-value
Share in Zurich	0.270	0.268	0.269	0.265	0.269	0.273	0.270	0.272	0.884
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Share in Switzerland	0.723	0.710	0.716	0.714	0.714	0.718	0.712	0.715	0.532
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Share in Germany/France/Italy/Austria	0.076	0.078	0.074	0.081	0.075	0.078	0.083	0.076	0.126
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	
mobile	0.559	0.566	0.562	0.563	0.569	0.555	0.561	0.557	0.323
	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	
desktop	0.424	0.416	0.418	0.419	0.413	0.425	0.417	0.426	0.314
	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Language: DE	0.841	0.836	0.841	0.836	0.836	0.843	0.839	0.847	0.226
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	
Hour of the first session: 8-14	0.369	0.365	0.374	0.372	0.370	0.366	0.368	0.379	0.398
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Hour of the first session: 14-20	0.339	0.342	0.332	0.336	0.340	0.334	0.345	0.335	0.433
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Share of users active in February	0.078	0.081	0.076	0.076	0.079	0.079	0.079	0.081	0.707
	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)	
For users active in Feb:									
Prob. of opening an ad after impression in Feb	0.065	0.065	0.069	0.063	0.061	0.064	0.067	0.065	0.687
(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)			
Prob. of action on an ad after view in Feb	0.121	0.131	0.123	0.129	0.130	0.130	0.116	0.126	0.433
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	
Number of Users	13015	13418	13565	13324	13346	13280	13381	13263	

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 User's characteristics: April

	0	1	2	3	4	5	6	7	p-value
Share in Zurich	0.289	0.281	0.281	0.283	0.283	0.288	0.273	0.284	0.247
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Share in Switzerland	0.716	0.716	0.712	0.712	0.717	0.720	0.706	0.713	0.411
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Share in Germany/France/Italy/Austria	0.078	0.081	0.076	0.078	0.078	0.079	0.081	0.077	0.896
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	
mobile	0.555	0.556	0.561	0.560	0.557	0.552	0.560	0.549	0.550
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	
desktop	0.424	0.424	0.421	0.421	0.424	0.429	0.421	0.432	0.680
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	
Language: DE	0.840	0.841	0.842	0.848	0.849	0.847	0.839	0.840	0.241
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	
Hour of the first session: 8-14	0.372	0.378	0.372	0.370	0.381	0.379	0.379	0.372	0.497
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	
Hour of the first session: 14-20	0.345	0.343	0.346	0.345	0.340	0.339	0.335	0.346	0.588
	(0.004)	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	
Share of users active in February	0.046	0.041	0.046	0.045	0.042	0.046	0.045	0.044	0.392
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	
For users active in Feb:									
Prob. of action on an ad after view in Feb	0.083	0.066	0.085	0.084	0.083	0.067	0.075	0.075	0.029
	(0.006)	(0.005)	(0.005)	(0.005)	(0.006)	(0.003)	(0.005)	(0.005)	
Prob. of action on an ad after view in Feb	0.116	0.114	0.135	0.124	0.121	0.124	0.116	0.115	0.896
	(0.010)	(0.011)	(0.011)	(0.010)	(0.011)	(0.011)	(0.010)	(0.010)	
Number of Users	11159	11044	11098	11172	11188	11330	11109	11089	

Notes: The table reports balance tests for user characteristics across the experiment arms of the April wave. The unit of observation is a user (cookie identifier) in the final analysis sample who was active in April; 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. The variables are defined as in Table A3.

Table A5: Balance Tests for User's characteristics: May

	0	1	2	3	4	5	6	7	p-value
Share in Zurich	0.270	0.268	0.269	0.265	0.269	0.273	0.270	0.272	0.884
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Share in Switzerland	0.723	0.710	0.716	0.714	0.714	0.718	0.712	0.715	0.532
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Share in Germany/France/Italy/Austria	0.076	0.078	0.074	0.081	0.075	0.078	0.083	0.076	0.126
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	
Device: mobile	0.559	0.566	0.562	0.563	0.569	0.555	0.561	0.557	0.323
	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	
Device: desktop	0.424	0.416	0.418	0.419	0.413	0.425	0.417	0.426	0.314
	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Language: DE	0.841	0.836	0.841	0.836	0.836	0.843	0.839	0.847	0.226
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	
Hour of the first session: 8-14	0.369	0.365	0.374	0.372	0.370	0.366	0.368	0.379	0.398
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Hour of the first session: 14-20	0.339	0.342	0.332	0.336	0.340	0.334	0.345	0.335	0.433
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Share of users active in February	0.032	0.032	0.030	0.030	0.033	0.034	0.030	0.029	0.291
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	
For users active in Feb:									
Prob. of opening an ad after impression in Feb	0.073	0.071	0.065	0.069	0.079	0.077	0.070	0.073	0.643
	(0.005)	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	
Prob. of action on an ad after view in Feb	0.127	0.120	0.120	0.113	0.142	0.114	0.121	0.092	0.291
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.011)	(0.012)	(0.010)	
Number of Users	12297	11936	12453	12365	12251	12485	12242	12175	

Notes: The table reports balance tests for user characteristics across the experiment arms of the May wave. The unit of observation is a user (cookie identifier) in the final analysis sample who was active in April; 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. The variables are defined as in Table A3.

Table A6: Summary of Session and User Characteristics

	Full Sample	Analysis Sample (Matched Sample)
A. Session Statistics		
Avg session length (min)	8.7	8.91
	(14.54)	(14.65)
Avg number of searches	1.76	1.78
	(1.37)	(1.39)
Avg total impressions per session	57.43	59.82
	(96.74)	(98.35)
Avg in-sample impressions per session	50.87	31.01
	(87.54)	(56.28)
Avg total views per session	2.51	2.6
	(4.38)	(4.47)
Avg in-sample views per session	2.05	1.24
	(3.48)	(2.27)
Avg total actions per session	0.26	0.27
	(1.27)	(1.31)
Avg in-sample actions per session	0.21	0.13
	(1.12)	(0.69)
Region (%)		
Zurich	32.4	31.73
Switzerland	78.82	77.09
Germany/France/Italy/Austria	7.19	6.98
Device properties (%)		
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
Assigned occupation (%):		
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
Number of sessions	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
Number of users	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 and 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 (ISCO one-digit) 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 A7: Descriptive statistics on wages and benefits in experimental and survey sample

	Main sample	With wage information	Survey sample
Wages			
Wage	6,614	6,614	6,525
	(11)	(11)	(75)
Benefits			
Flexible hours	.83	.91	.82
	(.00091)	(.0016)	(.016)
Work from home	.58	.66	.56
	(.0012)	(.0027)	(.021)
Childcare	.14	.15	.21
	(.00083)	(.002)	(.017)
Parking spot	.72	.77	.8
	(.0012)	(.0025)	(.017)
Public transport	.8	.89	.83
	(.0011)	(.0019)	(.016)
Company car	.19	.21	.21
	(.001)	(.0025)	(.018)
Canteen	.56	.64	.56
	(.0013)	(.0029)	(.021)
Food allowance	.37	.43	.4
	(.0013)	(.003)	(.021)
Coaching	.45	.58	.49
	(.0013)	(.003)	(.022)
Health measures	.5	.57	.5
	(.0013)	(.003)	(.022)
Company doctor	.15	.11	.14
	(.00093)	(.0019)	(.015)
Employee events	.86	.95	.92
	(.00091)	(.0014)	(.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), of the analysis-sample ads with kununu wage information (Column 2), and of the 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, averages are computed 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 above 2,000 job ads only have impressions 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 (kununu)	p-value
Occupational field (%):								
Sales and customer advisory	11.98	12.26	12.59	11.28	9.48	11.93	15.42	0.454
		(32.81)	(1.29)	(2.27)	(2.02)	(2.20)	(2.55)	
Administrative and specialised secretaries, clerks	5.71	17.99	18.36	15.90	18.48	17.43	18.91	0.927
		(38.42)	(1.51)	(2.62)	(2.68)	(2.57)	(2.77)	
Cleaning and facility services	8.19	4.11	2.58	7.69	4.27	5.05	4.48	0.066
		(19.86)	(0.62)	(1.91)	(1.39)	(1.49)	(1.46)	
Insurance and pension advisory, field sales	2.52	1.82	1.82	3.08	1.42	0.92	1.99	0.573
		(13.37)	(0.52)	(1.24)	(0.82)	(0.65)	(0.99)	
Truck and bus drivers	7.39	0.67	1.06	0.00	0.95	0.46	0.00	0.018
		(8.18)	(0.40)	(0.00)	(0.67)	(0.46)	(0.00)	
Logistics and transport workers	3.43	6.81	6.07	7.69	7.58	9.17	4.98	0.436
		(25.19)	(0.93)	(1.91)	(1.83)	(1.96)	(1.54)	
Customer information clerks	4.02	3.37	4.55	3.08	2.37	1.38	2.99	0.088
		(18.05)	(0.81)	(1.24)	(1.05)	(0.79)	(1.20)	
Marketing, PR and real estate professionals	3.64	9.16	10.17	8.72	7.58	8.72	8.46	0.792
		(28.86)	(1.18)	(2.02)	(1.83)	(1.91)	(1.97)	
Health professionals	3.18	10.11	9.86	9.23	9.00	12.84	9.95	0.740
		(30.15)	(1.16)	(2.08)	(1.97)	(2.27)	(2.11)	
Cooks, kitchen staff and service staff	5.06	4.51	3.49	6.15	6.64	4.59	3.98	0.359
		(20.77)	(0.72)	(1.72)	(1.72)	(1.42)	(1.38)	
Finance professionals	1.33	6.54	7.59	7.69	5.69	2.75	6.97	0.018
		(24.73)	(1.03)	(1.91)	(1.60)	(1.11)	(1.80)	
Teachers	3.16	3.30	2.28	4.10	2.37	5.96	3.98	0.174
		(17.87)	(0.58)	(1.42)	(1.05)	(1.61)	(1.38)	
Technical and engineering professionals	5.41	12.26	12.59	8.72	16.11	10.09	12.94	0.166
		(32.81)	(1.29)	(2.02)	(2.54)	(2.04)	(2.37)	
Machinery mechanics and repairers	2.66	1.68	2.12	2.05	0.47	0.92	1.99	0.162
		(12.87)	(0.56)	(1.02)	(0.47)	(0.65)	(0.99)	
Software and applications developers and analysts	3.40	5.39	4.86	4.62	7.58	7.80	2.99	0.123
		(22.59)	(0.84)	(1.51)	(1.83)	(1.82)	(1.20)	
N	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; 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"	T1a: Wage + Benefits	T1b: Wage Only	T1c: Benefits Only	T2: Wage + Benefits (kununu)	p-value
Demographics:							
Age	43.163	43.272	43.206	42.716	43.129	43.274	0.987
	(11.982)	(0.462)	(0.845)	(0.877)	(0.831)	(0.821)	
Female	0.504	0.478	0.518	0.474	0.555	0.552	0.143
	(0.500)	(0.019)	(0.036)	(0.034)	(0.034)	(0.035)	
Lives in Switzerland	0.929	0.932	0.903	0.948	0.936	0.920	0.495
	(0.257)	(0.010)	(0.021)	(0.015)	(0.017)	(0.019)	
Has a child aged 0–6	0.127	0.138	0.092	0.104	0.147	0.124	0.268
	(0.333)	(0.013)	(0.021)	(0.021)	(0.024)	(0.023)	
Has a child aged 7–18	0.244	0.241	0.251	0.251	0.243	0.239	0.997
	(0.430)	(0.017)	(0.031)	(0.030)	(0.029)	(0.030)	
Secondary education	0.410	0.405	0.431	0.398	0.408	0.418	0.965
	(0.492)	(0.019)	(0.036)	(0.034)	(0.033)	(0.035)	
Tertiary education	0.522	0.543	0.482	0.502	0.500	0.532	0.514
	(0.500)	(0.019)	(0.036)	(0.034)	(0.034)	(0.035)	
Job history:							
Duration of search (weeks)	17.090	17.839	14.585	16.810	16.821	17.647	0.452
	(23.321)	(0.946)	(1.469)	(1.509)	(1.663)	(1.602)	
Work experience in Switzerland	20.306	20.438	18.951	20.467	20.525	20.777	0.680
	(14.009)	(0.537)	(0.973)	(1.010)	(0.985)	(0.979)	
Employed	0.441	0.448	0.462	0.393	0.459	0.433	0.596
	(0.497)	(0.019)	(0.036)	(0.034)	(0.034)	(0.035)	
Monthly full-time wage (CHF)	7,122	7,317	6,829	7,039	7,026	6,960	0.355
	(3,499)	(141)	(220)	(243)	(246)	(235)	
Full-time equivalents (%)	83.951	85.404	82.923	82.332	81.917	84.095	0.250
	(23.305)	(0.865)	(1.732)	(1.678)	(1.756)	(1.526)	
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 he 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.

Table A10: Effect of the wage on ad actions conditional on viewing the ad

	(1) Across	(2) Within
$\widetilde{\log(w_j^{\mathrm{k}})} \times T^w_{s(i)j}$	.014	.068
	(.014)	(.077)
$T^w_{s(i)j}$	.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 on job seekers' probability to take any ad action—printing, saving, sharing, or navigating to the employer's website, conditional that the ad has been viewed. Column 1 report estimates from the across-ad specification, which exploits variation in posted wages across treatment arms. Column 2 report 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)
$\widetilde{\log(w_j^k)} \times T^w_{s(i)j}$	.015***	.017***	.03***	.016***	.013***	.015***
	(.0016)	(.00091)	(.0034)	(.0016)	(.0025)	(.0016)
$T^w_{s(i)j}$	.0089***	.0089***	.016***	.0075***	.0062***	.0083***
	(.00049)	(.00029)	(.00092)	(.00086)	(.00088)	(.00049)
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	.433	.458	.409	.331	.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	10,444,044	6,257,436	10,575,688	10,587,447	10,565,784

Notes: This table reports the results of specification checks for the effect of posted wages on job seekers' interest in a job advertisement. The outcome is an indicator of viewing an ad. The unit of observation is an ad shown to a job seeker in the search results (an ad impression). Column 1 reports the baseline specification, estimated using a linear probability model (LPM) with ad, rank, and session fixed effects. Column 2 removes ad fixed effects and instead controls for all ad characteristics visible in the search results list, including detailed occupation controls. Column 3 reports estimates from a Poisson model. Column 4 reports estimates from an LPM with user-by-month instead of session fixed effects. Column 5 reports estimates from an LPM with experiment arm FE instead of session fixed effects. Column 6 augments the baseline specification with three within-session controls that summarize what the job seeker has already learned about wages from prior impressions in the same search session: an indicator for whether the job seeker has been shown a wage in a prior impression, the log wage of the most recent prior wage-shown impression, and the running average log wage across all prior wage-shown impressions. The two continuous controls use the same recentered log wage $\widetilde{\log(w_j^k)}$ as the main wage variable, expressed as deviations from the mean log wage (the log of CHF 6,185). All three controls are constructed in event-time order and reflect only impressions that precede the current one (with zeros for impressions before any prior wage was shown, flagged by the indicator). All specifications exploit variation in posted wages across treatment arms. The posted-wage elasticity, reported at the bottom of the table, is calculated by dividing the log-wage coefficient 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) Baseline	(2) No short sessions	(3) Fulltime	(4) Same treatment	(5) No outliers	(6) Wage quintiles 2–4	(7) April: reordered icons
$\widetilde{\log(w_j^{\mathrm{k}})} \times T^w_{s(i)j}$	.015***	.016***	.017***	.016***	.015***	.021***	.016***
	(.0016)	(.0016)	(.0024)	(.0017)	(.0018)	(.0034)	(.0039)
$T^w_{s(i)j}$	.0089***	.0091***	.009***	.0091***	.0086***	.011***	.0065***
	(.00049)	(.0005)	(.00077)	(.00055)	(.0005)	(.00061)	(.0013)
Control mean	.0386	.0392	.0388	.0379	.0386	.0386	.0388
Posted wage elasticity	.393	.397	.434	.435	.377	.536	.405
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	10,565,784	10,186,919	4,328,518	7,951,644	10,396,098	9,835,871	916,992

Notes: This table reports robustness checks for the effect of posted wages 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 (an ad shown to a job seeker in the search results list). Column 1 reports the baseline specification. 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 occupation median (ISCO-2). Column 6 excludes matched ads in the bottom and top quintile of the recentered log-wage distribution (Appendix Figure A7)—focusing on the slope in the middle of the wage distribution and disregarding outliers at either end—and keeps all other impressions, including unmatched ads, from both the treatment and control arms. Column 7 restricts the sample to April impressions from the control group and a treatment arm in which the wage was shown after rather than before the benefit icons, providing a check that the placement (rather than the content) of the wage information does not drive the estimated effect. All specifications exploit variation in posted wages across treatment arms. The posted-wage elasticity, reported at the bottom of the table, is calculated by dividing the log-wage coefficient 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) Baseline	(2) No session FE	(3) Poisson	(4) No short sessions	(5) Same treatment arm
$flx_j^k \times T_{s(i)}^b$	.00021	-.00033	.00038	.00028	-.0001
	(.00043)	(.00037)	(.00073)	(.00044)	(.0005)
$ho_j^k \times T_{s(i)}^b$	.0023***	.0022***	.0038***	.0024***	.0028***
	(.00035)	(.00043)	(.00058)	(.00035)	(.0004)
$chld_j^k \times T_{s(i)}^b$	.0017***	.00086*	.0035***	.0017***	.0022***
	(.00044)	(.00046)	(.00077)	(.00045)	(.0005)
$prk_j^k \times T_{s(i)}^b$	.0012***	.0013***	.0019**	.0012***	.0012**
	(.00045)	(.00045)	(.00074)	(.00045)	(.0005)
$pub_j^k \times T_{s(i)}^b$	.00035	.00061	.0006	.00035	.00055
	(.0005)	(.00047)	(.00088)	(.00051)	(.00057)
$car_j^k \times T_{s(i)}^b$	.0028***	.0029***	.0047***	.0029***	.0029***
	(.00057)	(.00067)	(.00088)	(.00058)	(.00064)
$cant_j^k \times T_{s(i)}^b$	-.00039	.00035	-.000099	-.00045	-.00065
	(.00041)	(.00044)	(.00069)	(.00042)	(.00047)
$food_j^k \times T_{s(i)}^b$	.00049	-.00037	.00091	.0005	.00064
	(.00041)	(.00052)	(.00069)	(.00042)	(.00047)
$coach_j^k \times T_{s(i)}^b$	.00054	-.00011	.001	.00056	.00049
	(.00039)	(.00046)	(.00065)	(.00039)	(.00044)
$hlth_j^k \times T_{s(i)}^b$	.00014	-.00031	.00035	.00017	-.000053
	(.00041)	(.00049)	(.00069)	(.00041)	(.00047)
$doc_j^k \times T_{s(i)}^b$	-.0017***	-.0017***	-.0023**	-.0017***	-.0018***
	(.00057)	(.00058)	(.001)	(.00058)	(.00063)
$evnt_j^k \times T_{s(i)}^b$	-.0011*	.00052	-.0015	-.0011*	-.0011
	(.00057)	(.00048)	(.00097)	(.00058)	(.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	0.0120	0.00885
	(0.0231)	(0.0183)	(0.0192)
Lives in Switzerland	0.0219	-0.00874	-0.00118
	(0.0604)	(0.0595)	(0.0610)
Primary education	-0.0137	0.0803	0.0857
	(0.108)	(0.108)	(0.114)
Secondary education	0.0944	0.0898	0.0906
	(0.0775)	(0.0719)	(0.0747)
Tertiary education	0.213***	0.128*	0.129*
	(0.0804)	(0.0741)	(0.0770)
Employed	-0.0168	-0.000553	-0.00301
	(0.0206)	(0.0172)	(0.0179)
log(lastWage)	0.185***	0.114***	0.113***
	(0.0404)	(0.0339)	(0.0358)
Last job FTE	0.0000909	-0.0000826	-0.0000540
	(0.000649)	(0.000534)	(0.000557)
Work experience in Switzerland	0.000167	0.00102	0.000947
	(0.000977)	(0.000956)	(0.00100)
Has a child under 6	-0.00306	0.0332	0.0317
	(0.0295)	(0.0248)	(0.0258)
Has a child 7-18 yo	0.0391	0.0427*	0.0443*
	(0.0254)	(0.0228)	(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 the heterogeneity in elicited beliefs about respondents' own pay in the control group. The dependent variable is respondents' expected own log wage in a job. $log(lastWage)$ 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) $\Pr(\text{Offer}_{ij})$	(2) $\Pr(\text{Offer}_{ij})$	(3) $\Pr(\text{Offer}_{ij})$	(4) $\log(\text{Wage5yr}_{ij})$	(5) $\log(\text{Wage5yr}_{ij})$	(6) $\log(\text{Wage5yr}_{ij})$
$\widetilde{\log(w_j)}$	-0.0968***		-0.111***	0.483***		0.485***
	(0.0347)		(0.0346)	(0.0373)		(0.0372)
T_i^w	-0.0216	-0.0216	-0.0144	0.0440**	0.0446**	0.0429**
	(0.0216)	(0.0228)	(0.0237)	(0.0203)	(0.0213)	(0.0218)
$\widetilde{\log(w_j)} \times T_i^w$	0.0426	0.0590	0.0467	0.0101	0.00264	0.0101
	(0.0629)	(0.0689)	(0.0626)	(0.0707)	(0.0775)	(0.0692)
T_i^b			-0.00376			-0.00547
			(0.0321)			(0.0269)
wfh_j			0.0163			-0.00908
			(0.0100)			(0.0106)
$wfh_j \times T_i^b$			-0.0109			-0.00643
			(0.0183)			(0.0175)
flx_j			0.0220			-0.0238**
			(0.0138)			(0.0117)
$flx_j \times T_i^b$			-0.0114			0.0173
			(0.0255)			(0.0236)
$chld_j$			0.00870			0.0324***
			(0.0106)			(0.0108)
$chld_j \times T_i^b$			0.00722			-0.0102
			(0.0192)			(0.0197)
T_i^w						0
						(.)
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' belief about the competitiveness of a job vacancy and their future wages. The dependent variable in columns 1 through 3 is the respondents' perceived likelihood of getting the job and in columns 4 through 6 is respondents' expected own log wage in the job after 5 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 info on wage, company car, and atmosphere belief updating

	(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	0.188***	0	0
	(0.0208)	(.)	(0.0284)	(.)	(.)
T_i^b	-0.0206	-0.0217	-0.00262	0.00690	0.131
	(0.0253)	(0.0255)	(0.0330)	(0.0344)	(0.202)
flx_j	-0.0304***		-0.0572***		
	(0.00927)		(0.0150)		
$flx_j \times T_i^b$	-0.00803	-0.0187	-0.0226	-0.0376	-0.230
	(0.0219)	(0.0225)	(0.0305)	(0.0313)	(0.187)
wfh_j	0.000255		0.0713***		
	(0.00873)		(0.0107)		
$wfh_j \times T_i^b$	-0.0178	0.00132	-0.0121	-0.00192	-0.0955
	(0.0188)	(0.0171)	(0.0197)	(0.0203)	(0.111)
$chld_j$	0.0223**		-0.0683***		
	(0.00882)		(0.0110)		
$chld_j \times T_i^b$	0.0149	0.0125	0.0687***	0.0614***	0.0889
	(0.0199)	(0.0192)	(0.0221)	(0.0224)	(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 displayed the availability of flexible working hours, work from home, and firm-sponsored childcare, but not of 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***	0.413***	0.497***	0.462***	0.439***
	(0.0209)	(0.0425)	(0.0644)	(0.0241)	(0.0267)
T_i^w	0.0450***	0.0465**	0.0688***	0.0603***	0.0591***
	(0.0107)	(0.0212)	(0.0232)	(0.0139)	(0.0145)
$\widetilde{\log(w_j)} \times T_i^w$	0.156***	0.157**	0.116	0.161***	0.223***
	(0.0363)	(0.0794)	(0.0986)	(0.0430)	(0.0457)
Occ*Region FE	Yes	Yes	Yes	Yes	Yes
$\hat{\alpha}_w$	0.287	0.267	0.230	0.299	0.399
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 the 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***	-0.0992***	-0.187***	-0.107***	-0.114***
	(0.0142)	(0.0233)	(0.0558)	(0.0187)	(0.0201)
ben_j^k	0.0892***	0.0781***	0.00806	0.0589***	0.0582***
	(0.00864)	(0.0145)	(0.0398)	(0.0141)	(0.0164)
$ben_j^k \times T_i^b$	0.275***	0.250***	0.367***	0.275***	0.273***
	(0.0182)	(0.0287)	(0.0482)	(0.0238)	(0.0261)
Occ*Region FE	Yes	Yes	Yes	Yes	Yes
$\hat{\alpha}_b$	0.302	0.271	0.370	0.293	0.290
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 the benefits belief updating. We pool all respondents' beliefs about the availability of the three benefits in one regression. The variables are defined as in Table 5. Column 1 reports the estimates in the main survey sample (column 7 in Table 5). 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 where 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 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***	0.478***	0.412***	0.420***	0.466***
	(0.0301)	(0.0298)	(0.0381)	(0.0363)	(0.0344)
T_i^w	0.0570***	0.0309*	0.106***	0.0443***	-0.0204
	(0.0149)	(0.0162)	(0.0204)	(0.0159)	(0.0197)
$\widetilde{\log(w_j)} \times T_i^w$	0.230***	0.108**	0.146**	0.279***	0.202***
	(0.0496)	(0.0524)	(0.0672)	(0.0597)	(0.0599)
Occupation*Region FE	Yes	Yes	Yes	Yes	Yes
$\hat{\alpha}_w$	0.398	0.207	0.248	0.481	0.379
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 through 5 report estimates for respondents based on their tercile of reported last/current wage.

C Data Cleaning and Sample Construction

This appendix describes how the analysis sample of ad impressions is constructed 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. *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.

3. *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.

4. *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.

5. *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).

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 three 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 the benefit and wage reviews is described in the main text.

D Derivations of the Conceptual Framework

This section presents the derivations of decompositions of the reduced form elasticities in the structural parameters.

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

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

The second line substitutes the belief-updating rules from equations (2) and (3) into the application 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-vacancy estimate of the wage 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-vacancy 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:

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

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 $\left(\log w_j - \overline{\log w}\right)$ 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 in case of a positive γ_w. 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.

Without loss of generality, let us consider the case of two vacancies, one paying $\log w_1$ and the other

paying $\log w_2 > \log w_1$. Using the OLS formula, η_w^{across} can be written as follows:

$$
\begin{aligned}
\eta_w^{\text{across}} &= \frac{E[y_{ij}|T_i^w = 1, \log w_2] - E[y_{ij}|T_i^w = 1, \log w_1] - (E[y_{ij}|T_i^w = 0, \log w_2] - E[y_{ij}|T_i^w = 0, \log w_1])}{\log w_2 - \log w_1} \\
&= \frac{E\left[\varphi_2 + \beta_w \left(\alpha_w \log w_2 + (1-\alpha_w) \log w_{i2}^{\text{prior}}\right) + \epsilon_{i2} | T_i^w = 1, \log w_2\right]}{\log w_2 - \log w_1} \\
&\quad - \frac{E\left[\varphi_1 + \beta_w \left(\alpha_w \log w_1 + (1-\alpha_w) \log w_{i1}^{\text{prior}}\right) + \epsilon_{i1} | T_i^w = 1, \log w_1\right]}{\log w_2 - \log w_1} \\
&\quad - \frac{E\left[\varphi_2 + \beta_w \log w_{i2}^{\text{prior}} + \epsilon_{i2} | T_i^w = 0, \log w_2\right] - E\left[\varphi_1 + \beta_w \log w_{i1}^{\text{prior}} + \epsilon_{i1} | T_i^w = 0, \log w_1\right]}{\log w_2 - \log w_1} \\
&= \frac{\beta_w \left[\alpha_w \log w_2 + (1-\alpha_w) E[\log w_2^{\text{prior}}]\right] - \beta_w \left[\alpha_w \log w_1 + (1-\alpha_w) E[\log w_1^{\text{prior}}]\right]}{\log w_2 - \log w_1} \\
&\quad - \frac{\beta_w \left(E[\log w_2^{\text{prior}}] - E[\log w_1^{\text{prior}}]\right)}{\log w_2 - \log w_1} \\
&= \alpha_w \beta_w \left(1 - \frac{E\left[\log w_2^{\text{prior}}\right] - E\left[\log w_1^{\text{prior}}\right]}{\log w_2 - \log w_1}\right) \\
&= \alpha_w \beta_w (1 - \theta_w). \qquad (21)
\end{aligned}
$$

The first line states the OLS formula for the regression coefficient η_w^{across} with a binary regressor. 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 difference-in-differences expression, leaving only the updating term. The fourth line simplifies the expression, and the last line introduces the definition of $\theta \equiv \left(E\left[\log w_2^{\text{prior}}\right] - E\left[\log w_1^{\text{prior}}\right]\right) / (\log w_2 - \log w_1)$. θ measures how well job seekers were informed about wage differences in the first place. If $\theta = 1$, the difference in priors was exactly equal to the wage difference shown. In contrast, $\theta = 0$ would indicate that job seekers have very little prior knowledge about wage differences across vacancies. Thus, when using the across-vacancy variation to estimate the elasticity of applications with respect to wages shown, the 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 vacancy information treatment does not have any bite.

Decomposing parameters using survey experiment: To decompose the elasticity of applications into preferences, belief updating, and misperception about wages, we complemented the field experiment with the survey experiment described in section 2.2. As detailed above, the survey participants are randomized into a control group that observes the vacancies without wage and benefit information, whereas the treatment group observes the 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.

A job seeker i's posterior belief about the log wage level in vacancy j, $\log w_{ij}^{\text{post}}$, can be written as follows:

$$\log w_{ij}^{\text{post}} = \kappa_1 + \kappa_2 T_i + \rho_1 \left(\log w_j - \overline{\log w}\right) + \rho_2 T_i \times \left(\log w_j - \overline{\log w}\right) + \epsilon_{ij} \tag{22}$$

where κ_1 measures the misperception about the wage level at average-paying firms held by job seekers in the absence of the information treatment. ρ_1 measures misperceptions about how much more high-paying firms pay among the untreated job seekers. κ_2 and ρ_2 in turn estimate how much the treatment group updates their beliefs based on the information shown in the treatment condition.

To understand how the regression coefficients map into the parameters of our conceptual framework, let us first consider the formula for ρ_1 under a binary regressor:

$$\begin{aligned}
\rho_1 &= \frac{E\left[\log w_{ij}^{\text{post}} \mid T_i = 0, \log w_2\right] - E\left[\log w_{ij}^{\text{post}} \mid T_i = 0, \log w_1\right]}{\log w_2 - \log w_1} \\
&= \frac{E\left[\log w_2^{\text{prior}}\right] - E\left[\log w_1^{\text{prior}}\right]}{\log w_2 - \log w_1} \\
&= \theta.
\end{aligned} \tag{23}$$

Thus, ρ_1 exactly identifies θ. ρ_2 from the regression can be expressed as:

$$\begin{aligned}
\rho_2 &= \frac{\left(E\left[\log w_{ij}^{\text{post}} \mid T_i = 1, \log w_2\right] - E\left[\log w_{ij}^{\text{post}} \mid T_i = 1, \log w_1\right]\right)}{\log w_2 - \log w_1} \\
&\quad - \frac{\left(E\left[\log w_{ij}^{\text{post}} \mid T_i = 0, \log w_2\right] - E\left[\log w_{ij}^{\text{post}} \mid T_i = 0, \log w_1\right]\right)}{\log w_2 - \log w_1} \\
&= \frac{\alpha_w \left(\log w_2 - \log w_1 - \left(E\left[\log w_2^{\text{prior}}\right] - E\left[\log w_1^{\text{prior}}\right]\right)\right)}{\log w_2 - \log w_1} \\
&= \alpha_w \left(1 - \theta\right).
\end{aligned} \tag{24}$$

The first line uses the OLS expression for the coefficient under a binary regressor, line two plugs in the updating framework from equation 2, and line three uses the definition of θ. Thus, ρ_2 is a mix between the updating parameter α_w and the level of misperception $1 - \theta$. 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 $\frac{\rho_2}{1-\rho_1}$ from the regression equation (22).

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

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

The two groups of terms correspond to the two independent experiments that contribute to $\widehat{\mathrm{WTP}}^{k}$: the field experiment and the belief-elicitation survey. Within the field experiment, $\hat{\eta}_k^{\text{across}}$ and $\hat{\eta}_w^{\text{across}}$ are jointly estimated from the regression specification in Equation (9), and we take their full joint covariance into account. 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 are run on 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 $\operatorname{Var}(\hat{\sigma}) = (1-\hat{\theta})^2 \operatorname{Var}(\hat{\alpha}) + \hat{\alpha}^2 \operatorname{Var}(\hat{\theta}) - 2(1-\hat{\theta})\,\hat{\alpha}\operatorname{Cov}(\hat{\alpha}, \hat{\theta})$.

F Belief Spillovers: Evidence from the Field Experiment

This appendix develops the second, field-experimental test of belief spillovers summarized in Section 6. It derives how belief spillovers from wage information to beliefs about benefits would be reflected in the posted-wage elasticity of ad views, states the corresponding estimating equations, and reports the results. Throughout, α_{wb} denotes the belief-spillover parameter introduced in Equation (16), and β_w and β_b denote job seekers' valuation of wages and benefits as defined in Section 3.

Our field experiment provides another independent way to infer the extent of belief spillovers α_{wb}. Certain treatment groups only received information on wages, while others received information on both wages and benefits. Intuitively, if job seekers use the wage information to update their beliefs on benefits that were not shown, the elasticity of applications with respect to wages should differ in the two groups. The following formalizes this intuition. As established in Section 3, using the within job-ad variation the elasticity of applications with respect to wages conditional on benefit information—that is, for the treatment group that observes both wage and benefit information—can be written as

$$\eta_w^{\text{across}} = \alpha_w \beta_w (1 - \theta_w). \tag{26}$$

How does this elasticity differ for job seekers who do not observe benefit information but instead use wage information to update their beliefs according to equation (16)? Plugging the belief updating into our conceptual framework, the elasticity of applications with respect to wages using the within vacancy variation is

$$\begin{aligned} \eta_w^{\text{within}} &= \frac{Cov(\varphi_j + \beta_w \left[\alpha_w \log w_{ij} + (1-\alpha_w) \log w_{ij}^{\text{prior}}\right], \log w_{ij})}{Var(\log w_{ij})} \\ &\quad + \frac{Cov(\beta_b \left[\alpha_{wb}(\log w_{ij} - \log w_{ij}^{\text{prior}}) + (1-\alpha_{wb}) b_{ij}^{\text{prior}}\right] + \epsilon_{ij}, \log w_{ij})}{Var(\log w_{ij})} \\ &= \alpha_w \beta_w + \alpha_{wb} \beta_b. \end{aligned} \tag{27}$$

Here, we again exploit the random assignment of the displayed wage, which implies that the displayed wage level ($\log w_{ij}$) is independent of vacancy fixed effects, job seekers' prior beliefs, and the residual term ϵ_{ij}. Consequently, all covariance terms involving these variables vanish, yielding the final expression above. The formula for the between-vacancy variation is similar, only with the additional term $1 - \theta$.

If job seekers use wage information to update their prior beliefs about benefits, then, provided they value those benefits ($\beta_b \neq 0$), we would expect the application elasticities to differ across the different treatment groups, i.e., $\eta_w \neq \eta_{wb}$. We can formally test this with the following regression for the within-vacancy variation:

$$\begin{aligned} y_{s(i)j} &= \eta_w^{within} \log(\widetilde{w_{s(i)j}}) \times T_{s(i)j}^{w} + \eta_{wb}^{within} \log(\widetilde{w_{s(i)j}}) \times T_{s(i)j}^{wb} \\ &\quad + \sum_{k=1}^{K} \eta_k b_{s(i)j}^{k} \times T_{s(i)}^{k} + \varphi_j + \chi_s + \pi_r + \varepsilon_{s(i)j}, \end{aligned} \tag{28}$$

or alternatively the following regression equation for the across-vacancy variation:

$$\begin{aligned} y_{s(i)j} &= \gamma_w T^w_{s(i)j} + \gamma_{wb} T^{wb}_{s(i)j} + \eta^{across}_w \log(\widetilde{w_{s(i)j}}) \times T^w_{s(i)j} + \eta^{across}_{wb} \log(\widetilde{w_{s(i)j}}) \times T^{wb}_{s(i)j} \\ &\quad + \sum_{k=1}^{K} \eta_k b^k_{s(i)j} \times T^k_{s(i)} + \varphi_j + \chi_s + \pi_r + \varepsilon_{s(i)j}. \end{aligned} \tag{29}$$

$T^w_{s(i)j}$ is an indicator of belonging to the wage only treatment group, and $T^{wb}_{s(i)j}$ an indicator of belonging to the treatment group that displayed both wages and benefits. η_w and η_{wb} are the application elasticities in the two groups.

Table A20 reports the estimates from the across-ad and within-ad specifications in equations (28) and (29) using ad-view outcomes. Columns 1–2 include all wage-and-benefit treatment arms, while Columns 3–4 restrict the wage-and-benefit treatment group (T^{wb}) to the eight treatment arms that display at least two of the four benefits with the highest willingness-to-pay estimates in Figure 5: parking, childcare, work-from-home options, and a company car. In all specifications, the point estimates of η_w and η_{wb} are nearly identical. Consistent with the absence of economically meaningful cross-belief updating, we cannot reject the null hypothesis that $\eta_w = \eta_{wb}$ at conventional significance levels in any specification.

Table A20: Wage elasticities by treatment arm: wage only versus wage and benefits

	All wage+benefit arms		**Only arms with ≥2 valued benefits**	
	(1)	(2)	(3)	(4)
	Across	Within	Across	Within
$\widetilde{\log(w_j^{\mathrm{k}})} \times T^{w}_{s(i)j}$	.013***	.032***	.014***	.02**
	(.0026)	(.0091)	(.0026)	(.0099)
$\widetilde{\log(w_j^{\mathrm{k}})} \times T^{wb}_{s(i)j}$	.015***	.033***	.017***	.022**
	(.0016)	(.009)	(.0018)	(.0098)
$T^{w}_{s(i)j}$	.0077***		.0079***	
	(.00084)		(.00084)	
$T^{wb}_{s(i)j}$	.009***		.0089***	
	(.0005)		(.00057)	
Control mean	.038599	.038896	.038593	.038896
Posted wage elasticity, wage only	.346	.81	.354	.488
Posted wage elasticity, wage+benefits	.4	.843	.433	.537
p-value, $H_0 \colon \eta_w = \eta_{wb}$	0.367	0.469	0.223	0.308
Sample	All	All	Valued	Valued
Model	LPM	LPM	LPM	LPM
Ad fixed effects	Yes	Yes	Yes	Yes
Rank fixed effects	Yes	Yes	Yes	Yes
Session fixed effects	Yes	Yes	Yes	Yes
Observations	10,565,784	10,921,672	6,280,581	6,636,118

Notes: This table reports estimates of equations (28) and (29). The outcome is an indicator of ad view. The unit of observation is an ad shown to a job seeker in the list of search results (ad impression). The sample corresponds to the analysis sample defined in Section 2.1.4. Columns 1 and 3 report estimates from the across-ad specification, which exploits variation in posted wages across treatment arms; Columns 2 and 4 report estimates from the within-ad specification, which exploits experimental variation in average versus median wages within the same ad. Columns 1–2 use all treatment arms; Columns 3–4 restrict the wage-and-benefit cell (T^{wb}) to the eight treatment arms that display at least two of the four benefits with the largest estimated effects in Figure 5 (parking, childcare, work from home, company car). The Sample row reports which sample each column uses. $T^{w}_{s(i)j}$ equals one for impressions in which the wage is shown in a treatment arm that displays only wage information; $T^{wb}_{s(i)j}$ equals one for impressions in which the wage is shown in a treatment arm that displays both wage and benefit information. Both indicators equal zero for control impressions and for impressions in which the wage is not displayed. The posted-wage elasticities $\hat{\eta}_w$ and $\hat{\eta}_{wb}$ reported at the bottom of the table are calculated by dividing the corresponding 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. The reported p-value is for the test of equality $H_0 \colon \eta_w = \eta_{wb}$. Standard errors are clustered at the user-by-month level.

G Follow-up Survey Instrument

This appendix reproduces the complete follow-up survey instrument used in the belief–elicitation experiment. The original survey was programmed in Qualtrics and administered in German to job seekers in Switzerland. Below we provide an English translation of all screens and questions seen by respondents, together with a description of the randomization and survey flow. Hidden technical questions (timers, embedded data, browser information) are omitted.

G.1 Overview and Randomization

After consent, respondents first reported their job search profile (occupation and region). Based on these answers, the survey dynamically selected real job ads from jobchannel's platforms that matched the respondent's occupation and region.

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

- **Control group (Group 1):** Job ads shown without any additional information on wages or benefits beyond the standard platform layout.
- **Treatment Group 1 (Group 2):** Job ads shown with additional information from kununu.com. 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 (Group 3):** Job ads shown with wage and benefit information explicitly labeled as coming from kununu.com.

Each respondent then saw a preview of sample job ads and subsequently evaluated a sequence of (typically eight) job ads. For each job ad, they answered questions on wages, benefits, and one of three additional dimensions (chance of getting the job, workplace atmosphere, or wage after five years). At the end of the survey, respondents evaluated the job ads and the usefulness of wage and benefit information, interpreted benefit icons, and provided demographics and employment history. Respondents could optionally leave an email address for a lottery and for receiving a summary of the results or for future research contact.

Where the original survey used dynamically generated previews of real job ads, we indicate this with a textual placeholder.

G.2 Introduction and Consent

Screen C1: Welcome and consent

> **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?

Response options (single choice):

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

Screen C2: Confirmation if initially selecting "No" Displayed only if the respondent chose "No" on the previous screen.

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.

Response option (single choice):

- I would like to participate in the survey.

Only respondents who ultimately agree to participate proceed.

G.3 Job Search Profile (Target Occupation and Region)

Screen C3: Occupation

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

Response format: Drop-down list of occupational fields (ISCO-3 categories).

Screen C4: Region

In which of the following regions would you be most likely to search for a job?

Response format: Drop-down list of Swiss regions (NUTS-2 regions).

Answers to the occupation and region questions are used to select the job ads shown in the following parts of the survey.

G.4 Group-Specific Instructions and Preview of Job Ads

Respondents next received instructions that depended on their randomization group and, for Group 2, on their subgroup.

Screen C5: Control group (Group 1, no additional information) Displayed only to respondents in the control group.

> In the following, we will show you various job ads in the occupational field you selected that were posted on the job platforms of jobchannel.
>
> 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:

Job-ad previews (dynamically generated in the original survey):

> *[Job-ad preview image: Job 1]*
> *[Job-ad preview image: Job 2]*
> *[Job-ad preview image: Job 3]*

Screen C6: Treatment Group 1, subgroup TG1a (wage and benefit information) Displayed only to respondents in Group 2, subgroup TG1a.

> In the following, we will show you various job ads in the occupational field you selected that were posted on the job platforms of jobchannel.
>
> 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:

Job-ad previews with additional information (dynamically generated):

> *[Job-ad preview image: Job 1 with additional information]*
> *[Job-ad preview image: Job 2 with additional information]*
> *[Job-ad preview image: Job 3 with additional information]*

(Within this subgroup, along with the usual job-ad information, respondents see a grey bar with wage and benefit information; the source is not explicitly labeled.)

Screen C7: Treatment Group 1, subgroup TG1b (wage information only) Displayed only to respondents in Group 2, subgroup TG1b.

> In the following, we will show you various job ads in the occupational field you selected that were posted on the job platforms of jobchannel.
>
> 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:

Job-ad previews (dynamically generated; an additional grey bar shows wage information).

Screen C8: Treatment Group 1, subgroup TG1c (benefit information only) Displayed only to respondents in Group 2, subgroup TG1c.

Text identical to TG1a and TG1b above; only the additional grey bar content differs (benefits only, no explicit wage information).

Screen C9: Treatment Group 2 (Group 3, explicit kununu source) Displayed only to respondents in Group 3.

In the following, we will show you various job ads in the occupational field you selected that were posted on the job platforms of jobchannel.

Additionally, in a grey bar you will see information on wages and additional benefits that employees reported when rating their employer on `https://www.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:

Job-ad previews with explicit source information (dynamically generated):

[Job-ad preview image with kununu wage and benefits bar]
[Job-ad preview: Job 2]
[Job-ad preview: Job 3]

G.5 Main Belief-Elicitation Block (Repeated for Each Job Ad)

For each job ad, respondents saw the job details (and, depending on group, additional information in the grey bar) and then answered the following questions.

Screen C10: Job ad display

[Full job ad as shown on jobchannel, matching the respondent's occupation and region. Depending on treatment group, a grey bar displays additional information on wages and/or benefits; for Group 3, this bar is explicitly labeled as based on employee information from kununu.com.]

Question C1: Wage belief

1. What do you personally consider realistic?

How high do you believe your monthly gross wage in this job would be, scaled up to a full-time position (including one-twelfth of any special payments and one-twelfth of any 13th-month salary)?

Response format: open numerical entry (in CHF).

Question C2: Beliefs about specific benefits

2. What do you think about the following additional benefits: would the employer...

For each of the following items, respondents choose "Yes" or "No":

- **Flexible working hours:**
  "...allow you to determine your working hours yourself?"
- **Work from home:**
  "...allow you to work from home?"
- **Childcare contribution:**
  "...provide a financial contribution to the costs of external childcare?"
- **Company car:**
  "...provide you with a company car?"

Question C3: Additional job-related belief (randomized) For each job, respondents are randomly assigned to answer one of the following three questions (the choice of which question is asked is randomized at the respondent–job level):

1. **Chance of getting the job**

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

   Response format: slider or numerical input (0–100%).

2. **Workplace atmosphere**

   3. How would you rate the workplace atmosphere in this company?
   Please try to answer even if you do not know the company.

   Response format: 8-point scale:

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

3. **Wage after five years**

3. What do you think your monthly gross wage in this job would be five years after starting, scaled up to a full-time position (including one-twelfth of any special payments and one-twelfth of any 13th-month salary)?

Response format: open numerical entry (in CHF).

The block above (job display + questions C1–C3) is repeated for each job ad in the respondent's sequence (typically eight jobs).

G.6 Source Attribution for Wage and Benefit Information

After the job-evaluation blocks, respondents in treatments that saw additional information are asked about the perceived source of that information.

Screen C11: Source question – Groups 1 and 2 (where applicable) Displayed to respondents in Group 1 or Group 2 if they saw wage and/or benefit information in the grey bar.

What do you think, where do the details on wages and additional benefits come from?

Response format: multiple choice with an open "Other" field. (Exact response categories are not consequential; the key is that respondents could indicate their perceived source and write in an alternative.)

Screen C12: Source question – Group 2, wage only (TG1b) Displayed to respondents in subgroup TG1b, who saw wage information only.

What do you think, where do the details on the wages come from?

Response format: multiple choice with open "Other".

Screen C13: Source question – Group 2, benefits only (TG1c) Displayed to respondents in subgroup TG1c, who saw benefit information only.

What do you think, where do the details on the additional benefits come from?

Response format: multiple choice with open "Other".

Screen C14: Source question – Group 3 (kununu explicitly mentioned) Displayed to respondents in Group 3, after having seen wage and benefit information explicitly labeled as coming from kununu.com.

What do you think, where do the details on the wages and additional benefits come from?

Response format: multiple choice with open "Other".

G.7 Follow-Up Evaluation of Ads and Information

Question C4: Relevance of shown jobs

> How relevant were the jobs you just saw, compared to the jobs you typically search for?

Response format: 7-point Likert scale:

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

Question C5: Usefulness of showing benefit information (Groups 1 and 2) Displayed to respondents in Group 1 or Group 2.

> Would you find it useful if job platforms displayed information about additional benefits in job ads? (for example flexible working hours, childcare, company car)

Response format: 7-point Likert scale:

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

Question C6: Usefulness of showing benefit information from kununu (Group 3) Displayed only to respondents in Group 3.

> Would you find it useful if job platforms displayed information about additional benefits from kununu.com in job ads? (for example flexible working hours, childcare, company car)

Response format: same 7-point scale as above.

Question C7: Usefulness of showing wage information (Groups 1 and 2) Displayed to respondents in Group 1 or Group 2.

> Would you find it useful if job platforms displayed information about wages in job ads?

Response format: 7-point Likert scale (1 = not at all useful, 7 = very useful).

Question C8: Usefulness of showing wage information from kununu (Group 3) Displayed only to respondents in Group 3.

> Would you find it useful if job platforms displayed information about wages from kununu.com in job ads?

Response format: same 7-point scale as above.

G.8 Demographics and Job Search

Question C9: Attention check

> 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" as your answer below.

Response format: multiple choice with color options, where the correct combination is to select "red" and "yellow".

Question C10: Year of birth

Could you tell us your year of birth?

Response format: open numerical entry.

Question C11: Gender

What is your gender?

Response format: multiple choice (e.g., female, male, other / prefer not to say).

Question C12: Country of residence

In which country do you live?

Response format: drop-down list of countries.

Question C13: Postal code (Switzerland only) Displayed only if the respondent indicated Switzerland as the country of residence.

What is your postal code?

Response format: open entry.

Question C14: Children aged 0–6

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

Response format: numerical entry (0, 1, 2, …).

Question C15: Children aged 7–18

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

Response format: numerical entry.

Question C16: Highest completed education

What is your highest completed education?

Response format: multiple choice (e.g. compulsory schooling, vocational training, upper secondary, tertiary education, etc.; exact categories as in the original German survey).

Question C17: Duration of job search

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

Response format: numerical entry.

G.9 Interpretation of Benefit Icons

Respondents are next shown icons representing different job benefits and are asked how they would interpret these icons if they saw them on a job ad.

Question C18: Flexible working hours icon

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

[Icon for flexible working hours]

Which statement fits better?

Response format: multiple choice between two interpretations, plus an open "Other" field. (In the original survey, the two statements distinguished between weaker and stronger forms of flexibility; for example, some flexibility vs. largely self-determined working hours.)

Question C19: Work-from-home icon

What would you expect if you saw the following symbol on an open position?

[Icon for work from home]

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

Response format: numerical entry or categorical choice for expected number of work-from-home days per week, plus an "Other" option.

Question C20: Childcare support icon

What would you expect if you saw the following symbol on an open position?

[Icon for childcare support]

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

Response format: ordered categories for the percentage of childcare costs covered (e.g., 0%, 1–24%, 25–49%, 50–74%, 75–100%), plus an "Other" option.

G.10 Employment and Wage History

Question C21: Years of work experience in Switzerland

For how many years have you been working in Switzerland?

Response format: numerical entry.

Question C22: Current employment status

Are you currently employed (as an employee or self-employed)?

Response options:

- Yes
- No

Question C23: Current wage and employment degree Displayed only if the respondent is currently employed.

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)?

Response format: numerical entry (CHF).

What is your current employment degree (in %)?

Response format: numerical entry (0–100).

Question C24: Previous wage and employment degree Displayed only if the respondent is not currently employed.

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)?

Response format: numerical entry (CHF).

What was your employment degree in your last job (in %)?

Response format: numerical entry (0–100).

G.11 Conclusion, Lottery, and Debriefing

Screen C25: Lottery and email collection

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). 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 (GDPR) 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 as soon as 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 `jobsucheumfrage@kof.ethz.ch` to request access to your email address, to correct it, or to have it deleted.

Would you like to register for the lottery?

Response format:

- Yes, with this email address: *[open text field]*
- No, I do not want to participate in the lottery.

Screen C26: Debriefing (purpose of the study)

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

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

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

Screen C27: Request for summary of results

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 as soon as the summary has been sent.

Response options:

- Yes, please send it to this email address: *[open text field]*
- No, I do not wish to receive a summary.

Screen C28: Permission for future contact

May we contact you for possible new scientific surveys?

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

Response options:

- Yes, under this email address: *[open text field]*
- No, please do not contact me.

Screen C29: End of survey

We thank you for your participation in this survey. Your response has been recorded.