STAFF REPORTS

NO. 1124
SEPTEMBER 2024

The Pay and Non-Pay Content of Job Ads

Richard Audoly | Manudeep Bhuller | Tore Adam Reiremo

FEDERAL RESERVE BANK of NEW YORK

The Pay and Non-Pay Content of Job Ads
Richard Audoly, Manudeep Bhuller, and Tore Adam Reiremo
Federal Reserve Bank of New York Staff Reports, no. 1124
September 2024
https://doi.org/10.59576/sr.1124

Abstract

How informative are job ads about the actual pay and amenities offered by employers? Using a comprehensive database of job ads posted by Norwegian employers, we develop a methodology to systematically classify the information on both pay and non-pay job attributes advertised in vacancy texts. We link this information to measures of employer attractiveness, which we derive from a job search model estimated on observed wages and worker mobility flows. About 55 percent of job ads provide information related to pay and nearly all ads feature information on non-pay attributes. We show that publicly advertised job attributes are meaningful predictors of employer attractiveness, and non-pay attributes are about as predictive as pay-related attributes. High-pay employers mention pay-related attributes more often, while high-amenity employers are more likely to advertise flexible working hours and contract duration.

JEL classification: J23, J32, J33, J62, J63
Key words: amenities, non-pecuniary job attributes, compensating differentials, worker mobility, information frictions, pay transparency, job ads, text analysis, vacancy duration

---

Audoly: Federal Reserve Bank of New York (email: richard.audoly@ny.frb.org). Buller: University of Oslo, Statistics Norway, CEPR, IZA, CESifo (email: manudeep.bhuller@econ.uio.no). Reiremo: Norwegian School of Economics (email: tore.reiremo@nhh.no). The authors thank Statistics Norway and the Norwegian Public Employment Agency (NAV) for providing us access to their data. For helpful feedback and suggestions, they thank Christoph Heidricht, Lisa Kahn, Gizem Kosar, Thomas Le Barbanchon, Steve Machin, Alan Manning, Ryan Michaels, Chris Moser, Suresh Naidu, Barbara Petrongolo, Giorgio Topa, as well as conference and seminar participants at the FAIR Midway Conference 2022 (Tromso), the 2023 Nordic Meeting on Register Data and Economic Modelling (Oslo), the 2023 EEA-ESEM Meetings (Barcelona), the 3rd CRC Labor Workshop (Mannheim), the Federal Reserve System Applied Micro Conference (Chicago), the 2024 CEPR Annual Symposium in Labour Economics (London), and the 2024 NBER SI Labor Studies/Personnel Economics Meeting (Cambridge). This project has received funding from the European Union's Horizon Europe Research and Innovation Programme under Grant Agreement No. 101043127 and the Norwegian Research Council under Grant Agreement No. 275123.

1 Introduction

There is a long tradition in labor economics of seeing jobs as more than pay [e.g., Brown, 1980, Duncan and Holmlund, 1983, Rosen, 1986]. Phrases commonly used in the literature such as "compensating differentials," "amenities," or "workplace differentiation" all conceptualize the notion that jobs can be represented as a multidimensional bundle of attributes that workers trade off when considering opportunities with alternative employers. Such trade-offs can be substantial. Recent experimental evidence shows that workers at a US call center are willing to give up 8 percent of their salary for the option to work from home [Mas and Pallais, 2017], while undergraduates at a top US university would accept a 5.1 percent salary cut for a job offering the possibility to work part-time [Wiswall and Zafar, 2018].

What is unclear from this evidence, however, is how workers gather information on the pay and amenity attributes of a job while exploring their options in the labor market. The content of job postings represents one important search channel through which workers can gather information about these attributes.[1] But as employers ultimately decide on the content of publicly advertised job postings, it remains an open question whether such content in fact reflects the actual pay and amenities associated with different employers.

In this paper, we study in detail the information content of job ads. Our first objective is to systematically quantify what workers can learn about the pay and non-pay attributes of jobs from the texts of job ads. Our second objective is to provide a detailed description of how different employers–characterized both in terms of their pay premiums and amenity values–advertise different types of pay and non-pay attributes in job ads. Our third objective is to provide statistical measures of how well publicly advertised job attributes predict employer pay premiums and amenity values. Our fourth and final objective is to introduce quantitative measures of ad informativeness in an economic framework that allows us to

[1] Recently, Carrillo-Tudela et al. [2023] provide survey evidence from Germany showing that 55.3 percent of firms used job postings in their latest hire, while 88.1 percent of workers used job postings in their search process. Notably, however, job postings are not exclusive of other search methods, such as social networks. Carrillo-Tudela et al. [2023] find that firms and workers use, respectively, 1.9 and 2.3 channels on average.

study counterfactual job mobility flows under alternative informational environments.

To characterize the distribution of job attributes advertised by employers in job ads, we use the near-universe of publicly posted vacancies in Norway across several years. A key contribution of our paper is to systematically extract job attributes from these texts. Using tools from natural language processing, we identify nearly fifty commonly advertised pay and non-pay attributes, such as "competitive pay," "flexible work hours," and "nice work environment." We build a large collection of unique natural language expressions associated with each attribute, which we make publicly available online for other researchers to use.[2] This approach allows us to identify a large number of job attributes and to transparently map each attribute to many alternative expressions.

We find that about 55 percent of job ads provide some pay-related information, with 27 percent revealing explicit information on salaries, such as by mentioning an actual salary number, a salary bracket, or indicating that pay is set by a collective agreement. Moreover, nearly all ads feature information on some amenity attributes, such as contract duration, irregular hours, shift work, flexible work hours, workplace attributes, task-related attributes, or other minor perks. These results clearly suggest that employers use the texts of publicly posted job ads to advertise characteristics of jobs that workers potentially value.

A key strength of our data set is that each ad can be directly traced to the establishment posting the job. To shed light on how different employers advertise different types of information in publicly posted job ads and to assess the information content of these publicly advertised job attributes, we estimate workers' valuation of employers using job flows and hourly wages derived from matched employer-employee data. Building on Sorkin [2018] and Morchio and Moser [2024], we use a structural revealed preference approach and obtain measures of workers' valuation of different employers. Within our framework, this valuation can be further decomposed into a pay component, a non-pay component, and a job security component. Measuring these components separately matters because many of the attributes

[2] The full list of expressions is available in an online repository, accessible through this **link**.

advertised in job listings, such as workplace quality, are not necessarily reflected in pay.

By linking information on publicly advertised attributes in job ads to the employer-level estimates from our structural framework, we can measure the type of attributes posted by different employers. We find that high-pay employers are more likely to advertise compensation-related information, while high-amenity employers more often mention job amenities, such as "possibility to work flexible hours" and "regular daytime work schedule". This systematic relationship between workers' estimated valuation of employers, which is derived from actual labor market outcomes, and the propensity of these employers to mention specific attributes is evidence that employers supply credible pay and non-pay content in their job listings.

We show that the content of job ads has predictive power for workers' valuation of employers. Notably, we find that the attributes we detect in job ads can explain about 60 percent of the variation in overall employer values and employer pay premiums, and above 50 percent of amenity values and 40 percent of job security differences across employers. Job ads have predictive power over and above commonly observed job characteristics. Once we remove variation in employer values that can be attributed to flexible industry, occupation, and location controls, job attributes listed in ads can still explain 15-20 percent of the remaining variation across employers. In this regard, we find that attributes related to having flexible working hours and shift work are especially important predictors. Further, controlling for industry, occupation and location, we show that employers with higher pay premiums and amenity values associated with the information posted in job ads have shorter duration of posted vacancies, consistent with such employers being more attractive.

Finally, we use our model estimates to gauge the impact on worker mobility of having access to the content of job ads. We predict the counterfactual value of employers under alternative assumptions on workers' information sets and derive the implied mobility flows in each case. In our baseline, we find that aggregate job-to-job transitions to more highly-valued employers increase by 1.1 percent in an economy where workers have access to information on job attributes in job ads, as compared to a benchmark economy where only "one-liner" vacancies listing occupation, industry, and location are posted. This increase corresponds to 25 percent of the additional mobility towards better employers in the full-information model. Our results further suggest that the information in job ads can contribute to men and women self-selecting into distinct employers based on their preferences.

We see our approach as the first attempt at quantifying the breadth and quality of information present in the texts of job ads. We show that job openings contain a wide range of attributes relevant to workers' valuation of employers and that a subset of these attributes aligns with these valuations: the content of job ads is informative about the actual attractiveness of employers. These findings have implications for a wide range of models in the job search literature, such as models with a multi-dimensional job ladder [Hwang et al., 1998, Jarosch, 2023]. We leave it to future research to shed light on employers' endogenous choices of providing more or less content in their job ads, and the type of content they provide, in relation to their recruitment objectives.

Our results also have implications for the design of policies related to the information contained in job ads. There is a current policy drive to improve pay transparency within organisations in many countries, and several jurisdictions are also implementing policies to promote pay transparency in job openings [Cullen, 2024].[3] Taken together, our results on the prevalence of specific job attributes in vacancy texts and on the type of attributes advertised by more attractive employers can be helpful to guide further policy efforts in that direction.

Related literature Several recent studies have used the job description of vacancies as a source of information on the type of jobs advertised by employers. A large portion of this literature centers on skill requirements. Marinescu and Wolthoff [2020] focus on the role of job titles in accounting for the number of applicants received by each vacancy. Deming and Kahn [2018], Atalay et al. [2020], and Deming and Noray [2020] study variations in the demand for specific skills across local labor markets and over time using information

[3] As an example, in November 2022 New York City mandated employers to systematically include a salary bracket in their job ads.

extracted from the job description of vacancies. Many recent contributions specifically focus on the demand for computer science and artificial intelligence skills [see, e.g., Alekseeva et al., 2021, Acemoglu et al., 2022, Braxton and Taska, 2023, Contractor and Taska, 2022].

We contribute to this literature by instead considering a broad set of pay and non-pay attributes that are advertised by employers in job listings. In a similar vein, several recent contributions focus on a single specific attribute, such as the type of flexible work arrangements advertised in job ads [Adams-Prassl et al., 2023a], the possibility to work remotely [Hansen et al., 2023], and whether the employer offers training on the job [Adams-Prassl et al., 2023b]. We add to these important contributions by retrieving a comprehensive set of job attributes advertised in job ads, which encompasses flexible work arrangements, remote work and on-the-job training, but also what employers communicate about compensation or the quality of the workplace. In a similar spirit, Sockin [2022] uses text analysis to retrieve information on a large set of amenities from employer reviews posted by workers, rather than using vacancy texts. We see our results using vacancy texts as complementary since they instead reflect the information advertised by employers.[4]

A large body of theoretical and empirical work has emphasized that non-pay attributes are important determinants of how much workers value alternative employers. One strand of this literature elicits workers' preferences over specific attributes in survey experiments or quasi-experimental settings. Examples of non-pay attributes analyzed using this approach include workplace flexibility [Mas and Pallais, 2017, Wiswall and Zafar, 2018, Drake et al., 2023], job security [Datta, 2019], commuting time [Le Barbanchon et al., 2021], working from home [Nagler et al., 2022, Lewandowski et al., 2022], and shift work [Desiere and Walter, 2023], among others. Our paper instead provides a composite measure of employer amenity value that is linked to a large number of job attributes from job ads.

We follow the second strand of this literature, which recovers structural estimates of the amenity value of employers by modelling where workers choose to work in a labor market

[4]In other recent and related contemporary work, Arold et al. [2024] and Lagos [2024] use text analysis methods to retrieve information on amenities contained in collective bargaining agreement texts.

with heterogeneous employers. In this approach, estimates of the amenity value of employers are obtained by reconciling workers' choice of employer with their level of pay [Rosen, 1986]. The choice of workers are captured using either the size of employers in a static setting [e.g., Card et al., 2018, Lamadon et al., 2022] or mobility flows between employers [e.g., Sorkin, 2018, Taber and Vejlin, 2020, Morchio and Moser, 2024]. We rely on the latter approach since it can be recast in a job search framework, allowing us to further disentangle the amenity value of employers from differences in their recruitment effort. Relative to these important studies, we link the estimates of workers' valuation of employers to the attributes advertised in job openings and quantify the value of this information.[5]

We also contribute to the expanding body of work that aims at unpacking search frictions in the labor market. Jäger et al. [2024] elicit the beliefs of workers on their reemployment wage if they were to lose their job and compare these wages to actual labor market outcomes. Horton et al. [2024] and Belot et al. [2022] design experiments within a controlled environment to study how job seekers respond to variation in the wage featured in job vacancies. Our approach adds to these studies in that we focus on the texts of actual job postings as a source of information on alternative employers. Our evidence also relates to recent studies on how posted wages relate to the duration of vacancies [e.g., Mueller et al., 2024, Bassier et al., 2023, Faberman and Menzio, 2018].[6] We add to this literature by showing how measures of posted pay premiums and amenities relate to the duration of vacancies.

Outline The rest of the paper is organized as follows. Section 2 describes the data. Section 3 explains how we retrieve information on the attributes advertised by employers in vacancy texts and provide evidence on their prevalence in our data. Section 4 introduces our structural model and describes the estimated values of employers. Section 5 provides our evidence based on the linking of information in job ads to our structural estimates and

[5]Morchio and Moser [2024] also relate their estimates of the amenity values of employers to some of the workplace characteristics observed in their data, such as "workplace hazards" and "working hours flexibility".

[6]Several recent studies have also examined the elasticity of applications [e.g., Azar et al., 2022, Banfi and Villena-Roldan, 2019] and hires [e.g., Bassier et al., 2022, Hirsch et al., 2022] with respect to posted wages.

reports our results on the information content of the texts of job ads. Section 6 concludes.

2 Data and Institutional Context

2.1 Vacancy Data

We have access to almost 3 million job ads covering the near universe of publicly posted vacancies in Norway between 2002 and 2019. These data are maintained by the Norwegian Public Employment Agency (NAV). The employment agency collects information about vacancies from several sources, including online job boards and newspapers, as well as direct job opening reports from employers.[7] The share of online postings has increased gradually over the two decades, and about 80 percent of job postings recorded in the agency's database since 2015 were retrieved from various online job boards.[8] The remaining job ads were either (i) scanned or transcribed from newspapers by caseworkers in the employment agency, or (ii) enclosed in the notifications of job openings sent by employers directly to the agency. The ability to observe virtually all publicly posted vacancy ads in the economy with full-text corpora is an important advantage of our setting, which differentiates it from the existing literature using vacancy texts that often relies on information from selected online job portals, typically covering job ads posted over a shorter time span.

In addition to the job title and actual text of each job opening, our data contain the following structured information about each ad: unique establishment identifiers, the dates when the ad was registered and filled or removed (i.e., vacancy duration), the number of job openings per vacancy, and some additional information about job characteristics submitted by employers. The establishment identifiers are central to our analysis. They allow us to link each job opening to matched employer-employee administrative data and to compare the

[7] In accordance with the Labor Market Act §7, Norwegian employers are required by law to report publicly posted job openings to the agency, which maintains a comprehensive database of publicly posted vacancy ads with the stated goal of providing job-seekers with current information on suitable job opportunities.

[8] Bhuller et al. [2023] study the consequences on labor market matching of increased online job search and recruitment triggered by a roll-out of broadband internet across Norway during the early 2000s.

information we extract from the text to actual outcomes at the establishment level.[9] Each job ad in our data set also has an occupational code based on the 4-digit ISCO classification. This occupational code is assigned by caseworkers based on the job title and textual information on job descriptions and skill requirements stated by the employers in the vacancy, and is a novel feature of our data. Besides the structured information, the text of a typical job ad contains about 200 words. As we describe in Section 3, we use tools from the natural language processing literature to retrieve information on job attributes from these texts.

2.2 Matched Employer-Employee Data

We also have access to Norwegian administrative matched employer-employee data for all years between 2000 and 2019. The data come as a series of employment spells in each year, where each spell has information on the individual's employer, as well as their pre-tax earnings, hours of work, and employment start and end date. Individuals can have several recorded spells in any given month if they receive earnings from more than one employer. For each worker, we aggregate all employment spells to the annual level and retain observations corresponding to the main employer, defined as the establishment with the largest annual earnings. Using information on annual earnings and annual contracted hours of work, we further calculate the average hourly wage for each worker in their main job in each year. This data set has additional information on several background characteristics of employees (gender, education) and employers (location, industry). In the most recent years, the data set also has some information on a small subset of job attributes (e.g., shift work, contract type) for each employment spell that employers must report to Statistics Norway.

[9] Notably, around 7 percent of job ads in our data were posted through recruitment or temporary employment agencies, and we drop these from most of our main analysis, as we are unable to link such ads to the actual establishment where the job is placed. Further, in around 13 percent of job ads, the posting employer has for various reasons decided not to disclose the establishment name in the publicly posted information, but we do have the corresponding establishment identifier in our data as the agency maintains this information and could share this with us for research purposes (according to the Labor Market Act §7-4).

2.3 Sample Selection and Analysis Periods

While our vacancy and employer-employee data sets cover almost two decades, both in our descriptive and structural analyses we split the data into four five-year periods: 2000-2004, 2005-2009, 2010-2014, and 2015-2019. We do this in order to capture potential changes in workers' valuation of employers in terms of their pay and non-pay attributes, as well as employers' posting behavior. In most of our analysis, we focus on the 2015-2019 period as our baseline and provide supplementary results for the remaining three five-year periods.

We impose several sample restrictions at various stages of our analysis. Table 1 describes these restrictions and shows how they reduce the number of observations used in our baseline analysis for the 2015-2019 period.[10] As shown in the first row, prior to the selection restrictions, our initial matched employer-employee data contains about 300,000 establishments covering about 3.1 million unique workers and 13.3 million worker-year observations. During this period, we observe about 900,000 job ads in the employment agency's database.[11]

In Panel A of Table 1, we focus on the sample used for our text analysis in Section 3, where we limit attention to job ads with complete (non-missing) vacancy texts with at least one section written in Norwegian. As shown in Column (1), this restriction drops about 6 percent of job ads. Note, however, that less than half of the establishments in our initial sample publicly posted any job ads in the 2015-2019 period. Once we consider job ads with non-missing Norwegian texts, we are left with about 41 percent of establishments, as shown in Column (2). Meanwhile, the posting establishments tend to be substantially larger, and thus we still retain establishments covering about 90 percent of workers and worker-year observations in our text analysis, as shown in Columns (3)-(4). In Appendix Table A.2, we provide a comparison of how the various sample restrictions impact the sample composition.

Next, in Panel B, we consider the sample used for the structural model estimation in Section 4. For the structural estimation, we focus on prime-aged workers, i.e., aged 20–60,

[10] Appendix Table A.1 shows the sample restrictions for each five-year period, and jointly for 2000-2019.
[11] An ad is on average associated with 1.6 job openings, so these ads cover 1.4 million job openings.

Table 1: Overview of Sample Selection: Baseline Analysis Period.

	Job Ads (1)		Establishments (2)		Workers (3)		Worker-Years (4)	
All Observations	915,789	(100%)	292,327	(100%)	3,123,052	(100%)	13,278,337	(100%)
Panel A: Text Analysis								
Observe Vacancy Text	860,467	(94.0%)	118,741	(40.6%)	2,882,624	(92.3%)	11,893,847	(89.6%)
Panel B: Model Estimation								
Prime-age Workers	880,892	(96.2%)	268,565	(91.9%)	2,671,777	(85.6%)	11,253,721	(84.8%)
Strongly Connected Set	744,133	(81.3%)	78,136	(26.7%)	2,353,355	(75.4%)	9,436,842	(71.1%)
Panel C: Main Analysis								
Observe Vacancy Text	860,467	(94.0%)	118,741	(40.6%)	2,882,624	(92.3%)	11,893,847	(89.6%)
Prime-age Workers	829,137	(90.5%)	118,107	(40.4%)	2,486,782	(79.6%)	10,182,255	(76.7%)
Strongly Connected Set	701,117	(76.6%)	78,123	(26.7%)	2,353,237	(75.4%)	9,436,263	(71.1%)
Excluding Staffing Agencies	623,341	(68.1%)	77,164	(26.4%)	2,324,318	(74.4%)	9,239,281	(69.6%)

Note: This table documents the sample sizes for different sample restrictions imposed in the estimation for the 2015-2019 baseline period. The first row shows the total number of observations in our data on job ads (Column 1) and matched employer-employee data (Columns 2-4), prior to the various sample selection steps. "Observe Vacancy Text" refers to the set of job ads with non-missing vacancy text and at least one section written in Norwegian. "Prime-age Workers" are the set of workers between 20 and 60 (inclusive)."Strongly Connected Set" refers to the strongly connected set of employers used in the structural estimation. "Excluding Staffing Agencies" is the set of employers that are not recruitment or temporary employment agencies. Panels A, B and C show the sample restrictions imposed on the text analysis, the structural model estimation, and the main analysis, respectively. Each sample restriction is imposed sequentially within each panel. Appendix Table A.1 provides an overview of sample selection for each of the four five-year periods (2000-2004, 2005-2009, 2010-2014, and 2015-2019) and the combined 2000-2019 period, respectively.

following standard practice in the literature. This restriction alone drops about 15 percent of workers and worker-year observations and 8 percent of establishments. Further, we restrict the estimation to establishments that are in the set of strongly connected establishments, i.e., the set of establishments across which we observe two-way job mobility flows during the 2015-2019 period. With both of these restrictions, we exclude more than 70 percent of the establishments in our initial employer-employee data, but can still retain establishments covering about 75 (70) percent of workers (worker-year observations). Comparing the average characteristics of establishments used in the text analysis and the model estimation in Appendix Table A.2, we find that the two samples are broadly similar in terms of industry and worker composition, while establishments used in the model estimation are larger.

Lastly, in Panel C, we focus on the sample used in our main analysis in Section 5 where we relate the employer-level model estimates to the information extracted from job ads. For this analysis, we impose both sets of restrictions used in the earlier stages. In addition, we also drop job ads posted by staffing agencies, as for such ads we observe the establishment identifier for the staffing agency and not the actual employer.[12] As shown in the final row of Table 1, for the main analysis, we retain about 68 percent of job ads, 26 percent of establishments, and 74 (70) percent of workers (worker-year observations) in our initial data.

2.4 Job Postings and Recruitment

An underlying assumption in our analysis is that job postings are an important source through which employers recruit workers and where workers learn about available jobs. In order to assess the importance of job postings for recruitment, we link each ad in our vacancy data to a potential hire observed in our employer-employee data. Building this linkage requires us to assign a potential hire to each job ad by finding a job spell that starts (i) in the same establishment as the job posting, (ii) soon after the ad is posted, and (iii) in the same or a close occupation as the one in the ad. We are able to link 87.7 percent of all publicly posted non-staffing agency job ads to a hire within the following six months; this applies to 94.1 percent for ads posted by public sector state employers, 90.3 percent of ads by non-state public sector employers (e.g., local municipalities, public schools, public hospitals), and 87.5 percent of ads by private sector employers. Conversely, 42.6 percent of all hires can be related to a posted ad within the previous six months; this applies to 88.4 percent of hires by public sector state employers, 55.3 percent of hires by non-state public employers, and 35.7 percent of private sector hires.[13] These numbers suggest that job postings are indeed an important channel for job search and recruitment in the Norwegian labor market.

[12] Around 92.5 percent of job ads classified as "staffing agency" ads in our data were posted by temporary employment agencies (e.g., Manpower), while recruitment agencies account for the remaining share.

[13] While all employers have an obligation to report the vacancies they post to the Norwegian employment agency, it is up to them to decide whether or not to post a vacancy for the purposes of recruitment. For instance, employers may also use other recruitment channels, such as social networks, or directly solicit applications. However, recruitment by public sector state employers is regulated by state legislation, which requires as a general rule that all job openings are publicly advertised (the State Employees Act §4).

3 Pay and Non-Pay Attributes in Vacancy Texts

This section describes how we retrieve the job attributes publicly advertised by employers in job ads. These attributes are extracted from vacancy texts using tools developed in natural language processing. To the best of our knowledge, we are among the first to systematically extract a comprehensive set of pay and amenity attributes from the texts of job ads.

3.1 Extracting Advertised Job Attributes

Consider the sample vacancy text for an IT-consultant position in Figure 1. This text contains a variety of information on the type of tasks associated with the position, the required skills and experience necessary to perform these tasks, and the pay and non-pay attributes of the job. As this example makes clear, these attributes cover many different aspects of the position, ranging from the duration of the contract and regular hours ("full time, permanent position") to characteristics of the workplace ("good working environment") and information about the level of pay ("competitive conditions"). Our goal is to systematically extract this information from the open text of all vacancies in the database.

The key difficulty with extracting job attributes from the raw texts of job ads is that a single attribute can be expressed in many distinct ways in natural language. Besides, while prior work on the pay and non-pay attributes of jobs has centered on specific amenities, we do not have a definitive list of the job attributes advertised by employers in vacancy texts. We therefore proceed in three broad steps to extract these attributes from all job postings: (i) we pin down a list of attributes to extract, (ii) we ascribe a set of expressions to each of these attributes, and (iii) we label the entire corpus of vacancy texts based on these expressions.

Step 1: List of job attributes to include Our choice of which pay and non-pay job attributes to include in the analysis is based partly on information that is commonly stated in ads, partly on our knowledge of the particular context of the Norwegian labor market,

Figure 1: Sample Vacancy Text: IT-consultant.

IT-consultant

About the job: Full-time, Permanent position

We are looking for a new team player for our IT department! As an IT consultant, you will be responsible for our IT environment. It also includes getting to know large parts of our organization well through having internal user support as well as employee training. You simply get a varied working day, with different tasks to brush up on. In the future, you will also assist in larger IT development projects [...]

Tasks:

- IT security.
- SharePoint.
- Data warehouse.
- SOX Compliance.

Required qualifications:

- Likes IT technical challenges.
- Sociable and likes to communicate with people.
- Has relevant education.
- Has experience from similar positions.

We can offer:

- A workplace with many varied tasks and a high focus on quality.
- Training in our key products/services.
- Good working environment with solid professional expertise.
- Flexible working hours.
- Competitive conditions.

[...]

Note: Translation from Norwegian by the authors. In Appendix Figures A.1-A.2, we provide additional examples of representative vacancy texts covering a teaching substitute and a civil engineer, respectively.

and partly on attributes that are considered important in the existing literature.

We identify commonly mentioned attributes in publicly posted vacancies using two distinct strategies: human recognition on a subset of job ads in our data set and identifying commonly used phrases in the "we offer" sections of vacancy texts. We first directed several research assistants read through the text of 1,200 randomly chosen job ads in the vacancy corpus. The research assistants were asked to make a list of the commonly advertised attributes that are valuable to workers, to take note of the corresponding expressions associated with these attributes, and to indicate the presence of these attributes in each job posting.

As a second source of information on commonly advertised attributes, we isolate and extract common phrases from the "we offer" sections of job ads. We do this in two steps: First, we identify and extract all structured lists from the vacancy corpus.[14] As an example, in the sample vacancy text in Figure 1, we retrieve three distinct lists: "Tasks," "Required qualifications," and "We can offer." We find that 50 percent of the vacancy texts in our sample feature one or more such lists. Second, we isolate the "we offer" lists by applying unsupervised topic modeling [Blei et al., 2003] to the collection of lists retrieved in the previous step. Further methodological details are provided in Appendix B. The most common expressions in the "we offer" lists constitute the second source of information on commonly advertised job attributes that we are able to identify.

Based on the two approaches described above, we are able to identify 47 distinct job attributes, as well as an initial set of expressions associated with each attribute. For the purpose of illustration, we summarize the 47 attributes in ten broad categories within classes of pay and non-pay job attributes, as shown in Table 2. Taken together, these attributes cover a variety of different job characteristics advertised in job ads, such as information on career opportunities, convenient/inconvenient hours, and the quality of the workplace. The

[14] These lists are found by searching for consecutive sentences starting with a hyphen or bullet point and by searching for the HTML tags used to generate lists in online vacancies. By design, this approach identifies attributes that can be organised in structured lists. In the later steps, we search for the common phrases identified in structured lists across the full corpus of job ads, including ads that were not posted online or lacked structured lists, as well as search for alternative similar phrases identified using text analysis tools.

Table 2: Main Categories of Job Attributes.

Main Category	Examples
A. Pay Attributes	
– Compensation Scheme	Compensation level; Collective agreement pay; Incentive pay
– Financial Attributes	Insurance scheme; Pension scheme; Mortgage possibility
– Career Opportunities	On-the-job training; Career opportunities
B. Non-Pay Attributes	
– Hours of Work	Full-time; Part-time; Full-time/part-time choice
– Convenient Hours	Regular daytime work; Possibility to work flexible hours
– Inconvenient Hours	Shift-work; Weekend/evening/nights; On-call employment
– Contract Duration	Permanent job; Temporary job
– Workplace Attributes	Social environment; Good colleagues; Remote work
– Task-Related Attributes	Interesting tasks; Challenging tasks
– Other Minor Perks	Central location; Company vehicle; Company gym

Note: This table documents broad categories of attributes detected in vacancies, with notable examples of individual attributes. The full list of job attributes with descriptions is provided in Appendix Tables B.3-B.5.

full list of job attributes with descriptions is provided in Appendix Tables B.3-B.5.

Step 2: Expressions corresponding to each job attribute Each job attribute retrieved in Step 1 is associated with an initial set of expressions. For instance, the attribute "flexible working hours," is associated with the target expressions "flexible working hours" and "flexible work time arrangements."[15] In Step 2, we seek to enrich this initial set by including other common ways of referring to the same attributes.

To expand on the set of possible expressions in our analysis, we rely on the Continuous Bag-of-Words (CBOW) Word2Vec algorithm [Mikolov et al., 2013].[16] This approach measures the degree of similarity between phrases by exploiting that words with similar meanings tend to occur in similar contexts. As an example, information about flexible working hours is commonly given with information about whether the spell is full-time or part-time in our corpus. We train the CBOW model using our collection of job ads, thus associating words

[15] The corresponding expressions in Norwegian are "fleksibel arbeidstid" and "fleksible arbeidstidsordninger," which denote working hour schemes that offer workers the possibility to choose their own schedule.

[16] See Atalay et al. [2020] for an application of this algorithm to vacancy job titles.

that are similar in the context of job ads but not necessarily in other contexts. We then loop through the full collection of job ads and use the trained model to store phrases of one, two, or three consecutive words most similar to the phrases in our original lists from Step 1. Finally, we remove expressions that are wrongly classified as similar by the model. An example of a removed phrase is "flexible job adaptation," which is classified as similar to "flexible working hours" by the model, but does not imply flexibility in one's schedule.[17] This discussion highlights the main advantage of our approach. We are able to identify a large number of attributes that can be described by many different expressions, but we still retain control on which specific expressions are included.

In total, our dictionary consists of 1,772 unique expressions, and each of the 47 attributes is associated with around 39 expressions on average, although some attributes have many fewer associated expressions than others. We have made the full list of expressions is available in an online repository, which is accessible through this **link**. Differences in the number of associated phrases reflect both the number of ways an attribute is presented in vacancy texts and the specificity of an attribute. For instance, the attribute "inclusive work-life scheme" is usually discussed using specific terms in Norwegian, which is why the model identifies only a small number of expressions for this attribute. Conversely, the attribute "shift work" is associated with many alternative expressions since there are many ways jobs can involve shift work.

Step 3: Apply to all vacancies We simply use our 1,772 unique expressions to generate job attribute indicators for all ads in the data. A job ad is defined as advertising a given attribute if any of the expressions associated with the attribute is found in the ad text.

[17] Notably, as there are two variants of written Norwegian (Bokmål and Nynorsk), we translate all phrases from Bokmål to Nynorsk, and append phrases that are not already captured by our dictionary.

3.2 Validation

A potential drawback of our approach is that it might fail to detect attributes implied by more complex phrases. Consider the sentence "we offer flexibility in starting date as well as in working hours." As there is some text between "flexibility" and "working hours," our dictionary approach would not detect flexibility in working hours in this case. More generally, we can expect that a job posting does offer a specific attribute when our procedure identifies it (high precision rate), but we can be less certain that a job posting does not contain that same attribute when it is not identified by our procedure (low sensitivity rate).

To check how important this concern is in our setting, we performed an additional round of manual recognition. We selected another random sample of 400 job ads posted during the last five-year period, i.e., 2015-2019, and directed another group of research assistants to manually classify job attributes in this sample. We deliberately recruited a distinct group of assistants from the ones who contributed to the initial manual classification (Step 1 of Section 3.1). We gave each research assistant our full list of 47 attributes and a few sample expressions for each attribute, and tasked them to search for each of these attributes in the full text of each job ad. We can then systematically compare the attributes retrieved with our automated procedure to those detected in the manual human recognition.

Table 3 shows the results of this comparison for ten broad categories of job attributes. Column (1) reports the prevalence rates measured based on our text analysis for each of these categories in the sample of job ads from 2015–2019. Column (2) reports the corresponding rates for the random sample of 400 job ads, also based on our text analysis procedure, while Column (3) reports the prevalence rates from manual recognition for this random sample. Finally, Columns (4)-(6) provide rates of success, precision, and sensitivity by comparing prevalence rates for each category across text analysis and manual recognition.

Across the ten broad categories of job attributes, we find that our procedure performs well in terms of success, precision, and sensitivity. Success and precision rates are above or around 80 percent for most categories. The notable exceptions are "task-related attributes", where

Table 3: Validation of Detected Job Attributes.

	All Ads	Job Ads in the Validation Sample				
	Text Analysis	Text Analysis	Manual Recognition	Success Rate	Precision Rate	Sensitivity Rate
	(1)	(2)	(3)	(4)	(5)	(6)
A. Pay Attributes	**70.8**	**79.2**	**73.2**	**89.0**	**89.3**	**96.6**
— Compensation Scheme	55.9	64.8	61.8	92.5	91.9	96.4
— Financial Attributes	32.4	37.8	34.2	96.5	90.7	100.0
— Career Opportunities	43.6	48.8	41.8	79.0	71.3	83.2
B. Non-Pay Attributes	**96.5**	**98.0**	**97.0**	**97.5**	**98.2**	**99.2**
— Hours of Work	77.0	76.0	71.8	92.2	92.1	97.6
— Convenient Hours	20.5	15.8	15.8	94.0	81.0	81.0
— Inconvenient Hours	27.0	28.0	22.8	92.8	77.7	95.6
— Contract Duration	65.0	67.2	66.8	88.0	90.7	91.4
— Workplace Attributes	50.7	55.0	51.0	82.0	80.0	86.3
— Task-Related Attributes	55.9	65.0	46.5	72.0	64.2	89.8
— Minor Perks	39.8	46.0	18.5	66.5	33.7	83.8
Any Observed Attribute	97.0	98.2	98.0	98.2	99.0	99.2
Average Number of Attributes	6.11	6.64	5.56	–	–	–
Number of Job Ads	860,467	400	400	400	–	–

Note: This table compares prevalence rates for ten broad categories of job attributes measured using our text analysis approach and manual recognition. A broad category of job attribute is considered present if at least one of the underlying distinct job attributes is detected. Column (1) is for the full sample of job ads with non-missing vacancy text posted between 2015 and 2019, while Columns (2)-(6) consider a random sample of 400 job ads used in the validation. Columns (1)-(3) report prevalence rates, while Columns (4)-(6) provide summary statistics that compare prevalence rates from text analysis and manual recognition. Success rate is defined as the share of job ads where the text analysis yields the same result as manual recognition. Precision is the rate of agreement between the two methods given that an attribute was detected in our text analysis, while, sensitivity is the rate of agreement given that an attribute was detected in the manual recognition. Precision is a measure of how many detected attributes are false positives, and sensitivity is a measure of how good our method is at recovering true attributes. See Appendix Table A.3 for results from validation exercises for each of the 47 underlying job attributes that contribute to the ten broad categories shown here.

we find a success rate of 72 percent and precision of 64.2 percent, and "minor perks", where we find a success rate of 66.5 percent and precision of 33.7 percent. The former is driven by differences in detection rates for whether the job "involves leadership responsibilities" and "work involves travelling," while the latter is driven by "central location," which is more frequently detected in our text analysis.[18] Overall, high levels of precision indicate that

[18]See Appendix Table A.3, which shows the results of the same validation exercise for each of the 47

the attributes recovered by text analysis reflect the actual content of the attributes well. Sensitivity is always equal to or larger than precision, above 80 percent for all attributes.

3.3 The Prevalence of Advertised Job Attributes

We now use the detailed advertised job attributes to highlight several salient descriptive statistics. Figure 2 reports the share of job ads posted between 2015 and 2019 advertising each of the 47 attributes. About 55 percent of job ads provide some compensation-related information. However, explicit information about the actual level of pay is scarce. Less that 10 percent of ads mention a salary number or bracket ("compensation level") and slightly more than 25 percent of ads indicate that pay is set by a collective wage bargaining agreement ("collective agreement pay").[19] As these two attributes can overlap, taken together, we find that 27 percent of job ads feature explicit information about the actual pay level.

Beyond information on the more tangible attributes of the job, such as pay, hours of work and contract duration, Figure 2 also shows that employers frequently choose to advertise more subjective aspects of the working environment, such as characteristics of the workplace ("workplace attributes") or some appreciation of the type of tasks involved ("task-related attributes"). Around 30 percent of job ads describe these tasks as "challenging" or "interesting." Similarly, 45 percent of vacancies feature some language about the quality of the work environment and more than 20 percent advertise "good colleagues."

Next, we analyze to what extent differences in advertised attributes are explained by additional observable characteristics of the job ad, such as the industry of the posting establishment or the occupation associated with the job title stated in the job ad. To describe the variation in advertised job attributes that can be explained by these characteristics, we estimate a series of logistic regressions for the probability that a given attribute is advertised in the text of a vacancy with several sets of fixed effects: industries, locations, occupations, and establishments. Figure 3 reports the estimated pseudo-R^2, a summary measure of goodness

underlying the ten broad categories in Table 3.

[19]By comparison, Batra et al. [2023] find that 13 percent of job ads have salary information in US data.

Figure 2: The Prevalence of Job Attributes Advertised in Vacancy Texts.

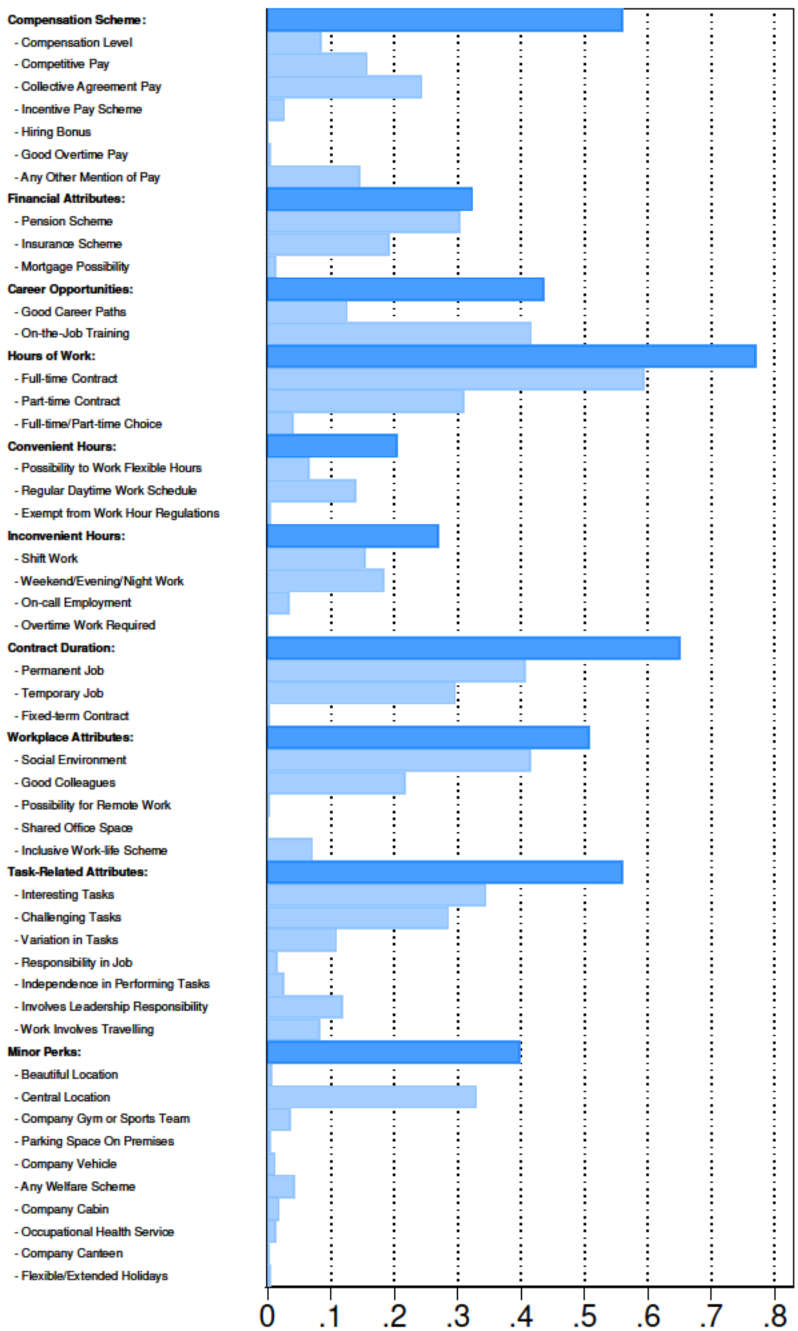

Note: This figure documents the prevalence of job attributes detected in job ads posted in Norway between 2015 and 2019 [N=860,467]. The light blue bars show the share of ads detected having each of the distinct job attributes. The dark blue bars show the share of ads detected with at least one attribute within ten broad categories. Appendix Figure A.3 shows the corresponding results for the 2002-2019 period as a whole.

Figure 3: Explained Variation in Publicly Advertised Job Attributes.

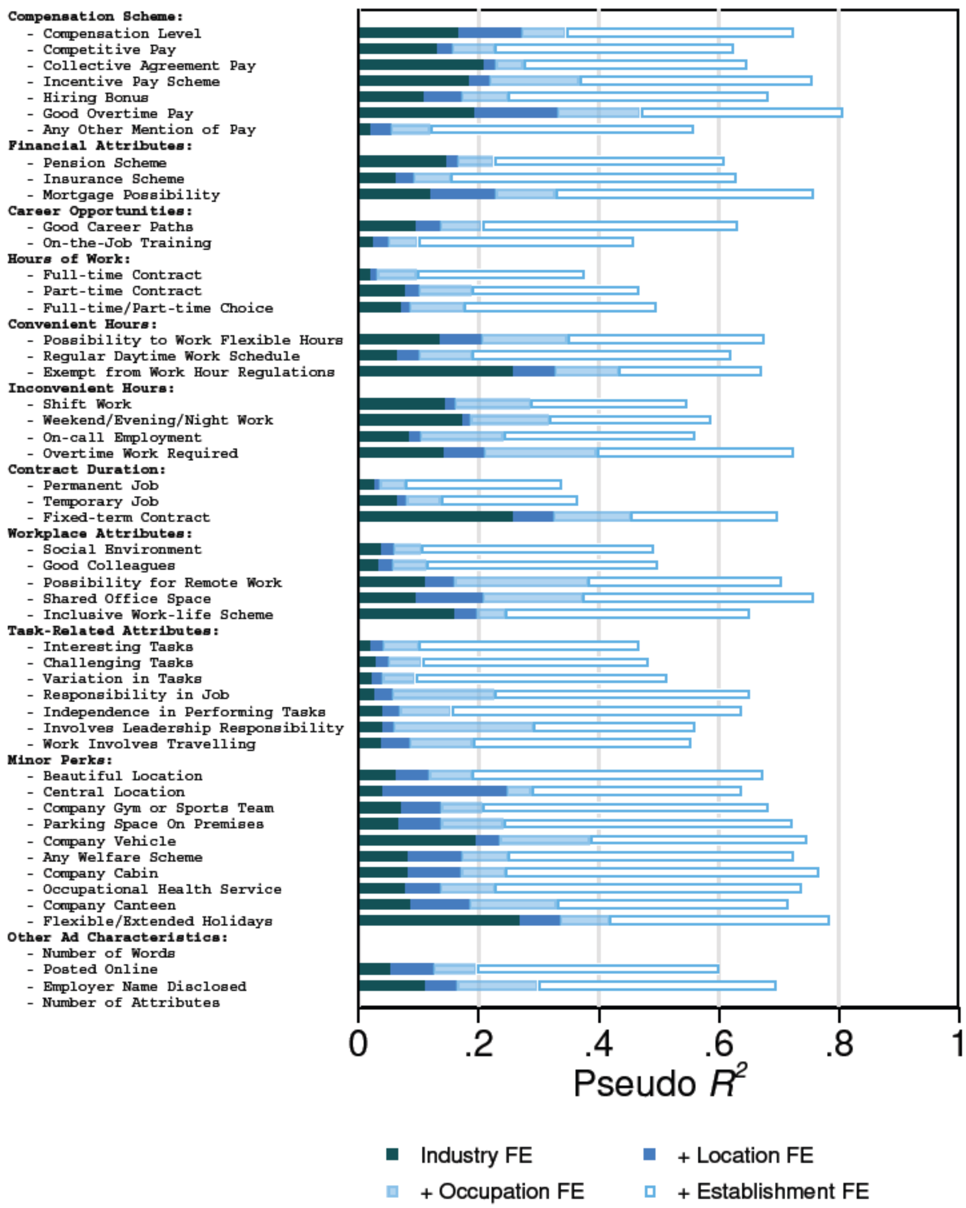

Note: This figure shows pseudo R^2 [McFadden, 1974] from separate logistic regressions of binary job attributes detected in job ads posted in Norway between 2015 and 2019 [N=719,538] on fixed effects denoting unique combinations of 2-digit industries (88 groups), location indicators (10 groups), 2-digit occupations (42 groups), and 59,259 establishments. We start by including industry fixed effects, continue by including industry×location fixed effects, and so on. The last set of regressions controls for 114,380 unique combinations. The regression excludes 140,929 ads with a unique combination of establishment, occupation, location, and industry indicators, since these observations are perfectly predicted by the last set of regressions by construction. Location indicators group 422 municipalities into ten groups based on the number of workers such that municipalities assigned to the same group have a similar number of workers, with a specific indicator for job postings in Oslo. Appendix Figure A.4 shows the corresponding results for the 2002-2019 period as a whole.

of fit for binary outcomes [McFadden, 1974], from logistic regressions estimated separately for each attribute.[20] This measure therefore captures the variation in the content of job ads within industries, locations, occupations, and establishments.

A possibility is that the attributes advertised for a given position are largely determined by its industry or occupation, so the pay and non-pay attributes mentioned in vacancy texts directly reflect the type of job advertised. For example, it can be expected that working as a hospital nurse involves shift-work, in which case there is no value-added to mentioning this information in the corresponding job description. The data instead suggest that there is variation in advertised job attributes for very similar positions. Figure 3 shows that industries, occupations, and location fixed-effects explain at most 30 to 40 percent of the variation in the attributes advertised by employers across the 47 categories we retrieve. By contrast, when we add establishment fixed-effects, we find that the explained variation jumps up to 50 to 70 percent across attributes, which suggests that advertised attributes are in large part correlated across vacancies for establishments observed with several job ads. This last result is consistent with the existence of an "establishment fixed-effect" in advertised attributes, and it justifies the establishment-level analysis we develop in the next sections.

4 Estimating the Value of Employers

In this section, we estimate how workers value alternative employers using data on realized worker flows and wages. Our goal is to obtain measures of workers' valuation of employers to which we can compare the information advertised in job ads. To this end, we estimate a version of the Burdett and Mortensen [1998] model in which employers post a value bundle using the revealed preference approach put forward by Sorkin [2018]. In this setting, workers' preferences over heterogeneous employers can be recovered from their mobility decisions. We

[20] As an alternative, we also estimated linear probability models using OLS with the same set of fixed effects separately for each attribute (results available upon request). However, as all of our pay and non-job attributes are binary, one could be concerned about interpreting standard goodness of fit from such regressions, which restrict the range of the coefficient of determination R^2 [Cox and Wermuth, 1992].

expand on this methodology by adding information on the wages paid by employers, from which we can break down workers' valuation of employers into a pay and non-pay component.

4.1 Model

Consider an economy in discrete time populated by a fixed measure of infinitely-lived workers and a fixed discrete number J of employers. Workers have discount factor $\beta \in (0,1)$ and can be either employed at one of the J employers or non-employed. Let $j \in \mathcal{J} = \{1, \ldots, J\}$ denote employers and V_j the value to a worker of being currently employed at employer j. Let N denote non-employment and V_N the value to a worker of being currently non-employed.

Both employed and non-employed workers search for better employers. Search is random. When searching, workers make contact with one of the $j \in \mathcal{J}$ employers with probability f_j, where f captures the relative weight of employers in workers' search. We maintain the wage-posting (here "value-posting") assumption from Burdett and Mortensen [1998]. The value of working at j is summarized by V_j and is not renegotiated even in the event of a change in the worker's outside option.

Let M_{jk} denote the measure of workers making a job-to-job move from employer j to employer k. M_{jk} is assumed to be the sum of involuntary relocation flows M_{jk}^R and voluntary flows M_{jk}^V. Relocation job-to-job flows are given by

$$M_{jk}^R = L_j \cdot \rho_j \cdot f_k, \tag{1}$$

where L_j is employment at j, ρ_j is the probability that workers employed at j are forced to reallocate to another employer j, and f_k is the probability that they find a job at employer k.[21] In Equation (1), workers do not make a choice and are forced to relocate to employer

[21] There are no such relocation shocks in Burdett and Mortensen [1998], but they are a common feature in random search models with on-the-job search (see, e.g., Jolivet et al. [2006] and Bagger and Lentz [2019]).

k. This is in contrast with voluntary job-to-job flows, which are given by

$$M_{jk}^{V} = L_j \cdot (1 - \rho_j - \delta_j) \cdot \lambda_1 \cdot f_k \cdot \Pr(k \succ j). \tag{2}$$

Voluntary worker flows from j to k are given by the measure of workers employed at j (L_j), who are neither forced to reallocate (with probability ρ_j) nor forced to move to non-employment (with probability δ_j), search for an alternative employer (with probability λ_1), and get an offer from employer k with probability f_k.[22] In Equation (2), workers choose whether to remain with employer j or move to k, which occurs with probability $\Pr(k \succ j)$. This is the sense in which these flows are voluntary and reveal workers' preferences over the set of employers $\mathcal{J}$. We follow Sorkin [2018] and assume there are i.i.d. taste shocks $\{\varepsilon_i\}_{i \in \{\mathcal{J}, N\}}$ that make the value of working at employer j and of non-employment worker-specific in each period. We further assume that these shocks are drawn from a Gumbel distribution with location parameter normalized to zero and scale parameter σ^{-1}.[23] With this assumption, the acceptance probability for a worker with a job offer from employer k currently working at j is given by

$$\Pr(k \succ j) = \Pr(V_k + \varepsilon_k \geq V_j + \varepsilon_j) = \frac{\exp(\sigma V_j)}{\exp(\sigma V_j) + \exp(\sigma V_k)}.$$

The sequence of events described by the relocation and voluntary worker flows in Equations (1)-(2) implies the following expression for the value V_j of working at employer j. Let

[22] The probabilities ρ_j and δ_j are assumed to be mutually exclusive.

[23] Normalizing the location parameter to zero is without loss of generality because it shifts the value of working at each employer j and of being in non-employment by the same amount.

u_j denote the flow utility of working at j. V_j can be expressed as

$$
\begin{aligned}
V_j = u_j &+ \beta\delta_j \mathbf{E}\big[V_N + \varepsilon_N\big] \\
&+ \beta\rho_j \sum_{k\in\mathcal{J}} f_k \mathbf{E}\big[V_k + \varepsilon_k\big] \\
&+ \beta(1-\delta_j-\rho_j)\lambda_1 \sum_{k\in\mathcal{J}} f_k \mathbf{E}\big[\max\big\{V_j+\varepsilon_j, V_k+\varepsilon_k\big\}\big] \\
&+ \beta(1-\delta_j-\rho_j)(1-\lambda_1)\mathbf{E}\big[\max\big\{V_N+\varepsilon_N, V_j+\varepsilon_j\big\}\big].
\end{aligned}
\tag{3}
$$

In Equation (3), the continuation value is made of the four following terms. With probability δ_j, the worker transitions to non-employment, an exogenous shock. With probability ρ_j, the worker is hit by a relocation shock and is forced to move to an alternative employer by drawing from the offer distribution f, an exogenous shock. With probability $(1-\delta_j-\rho_j)\lambda_1$, the worker gets an offer from a potential alternative employer by drawing from the offer distribution f and decides whether to stay or move, an endogenous choice. With probability $(1-\delta_j-\rho_j)(1-\lambda_1)$, the worker decides whether to stay with their current employer or to move to non-employment, also an endogenous choice.

We depart from Sorkin [2018] by making an explicit functional form assumption on the utility flow u_j of working at employer j.[24] We assume that u_j depends linearly on a pay component $\ln W_j$ and a non-pay (or amenity) component a_j

$$
u_j = \ln W_j + a_j. \tag{4}
$$

The additive formulation $\ln W_j + a_j$ is standard in the literature that estimates static models of worker preferences over employers [Card et al., 2018, Lamadon et al., 2022]. With an additional assumption to pin down the scale of the i.i.d taste shocks $\{\varepsilon_i\}_{i\in\{\mathcal{J},N\}}$, we can recover the non-pay component a_j associated with working at employer j by inverting the value function in (3), conditional on the wage W_j paid by employer j. We discuss this

[24] Sorkin [2018] makes a similar assumption on the value V_j of employers; Lagos [2024] also follows a similar approach. By contrast, V_j is not additively separable into pay and non-pay components in our formulation.

assumption below.

Given the form of the utility flow u_j derived from working at employer j, workers' valuation of j can be summarized by the vector $\left(\ln W_j, a_j, \delta_j, \rho_j\right)$. In what follows, we refer to $\ln W_j$ as the pay component, a_j as the non-pay component, and $s_j = 1 - \delta_j - \rho_j$ as the job security component of workers' preferences over employers.

We complete the description of the model by noting that there are similar expressions for workers' flows to and from non-employment (counterparts to Equation (1) and Equation (2)) and for the value of worker in non-employment (a counterpart to Equation (3)), where workers in non-employment get a draw from the job offer distribution $\{f_j\}$ at rate λ_0. These expressions follow naturally and are relegated to Appendix C.

4.2 Identification

The model has several employer-level parameters, such as workers' valuations of alternative employers $\{V_j\}_{j\in\mathcal{J}}$, and aggregate-level parameters, such as the offer arrival rate for employed workers λ_1. These parameters are identified using the matched employer-employee data described in Section 2, which contain the relevant information on wages and employer-to-employer flows. We now lay out the identification argument in a series of heuristic steps. At each step, we highlight the key assumptions required for identification.

Step 1: Parameters identified from worker flows The employer-level parameters $\{\sigma V_j, L_j, f_j, \rho_j, \delta_j\}_{j\in\mathcal{J}}$ and the transition parameters λ_0 and λ_1 are identified from (i) the realized mobility flows between employers M_{jk}, and (ii) which employers are shrinking or growing [see Sorkin, 2018, Section V.B]. The notation σV_j emphasizes that worker flows pin down the value of working at employer j up to the scale of the idiosyncratic taste shocks, which is retrieved separately in the next steps.

We briefly outline the main source of identification for each parameter intuitively and refer to Sorkin [2018] for details. The main source of identification for σV_j is the aggregation of the realized mobility flows M_{jk}. An estimate for L_j is obtained directly from the data as the share of worker-year observations at employer j. f_j is identified from the share of non-employment flows to employer j. The separation shocks δ_j and ρ_j are identified from the separation flows at shrinking firms, which give information on involuntary worker flows. Finally, the transition parameters λ_0 and λ_1 are pinned down by accounting for the aggregate transition flows, respectively from non-employment to employment and between employers.

Two sample restrictions are required for identification at this step. First, the aggregation of mobility flows only identifies V_j within the strongly connected sets of employers.[25] Second, each employer must hire at least one worker from non-employment for its sampling weight f_j to be well-defined.

Step 2: Flow utility parameters Given $\{\sigma V_j, L_j, f_j, \rho_j, \delta_j\}_{j\in\mathcal{J}}$ and the aggregate transition rate λ_1, σu_j can be recovered directly from Equation (3), again up to the scale of the i.i.d. taste shocks. This calculation is straightforward since, given that $\{\sigma\varepsilon_i\}_{i\in\{\mathcal{J},N\}}$ are i.i.d. draws from a Type 1 Extreme Value distribution with scale one, all expectations in Equation (3) admit the following closed-form solutions

$$\begin{aligned} \sigma\mathbf{E}\big[V_j+\varepsilon_j\big] &= \sigma V_j+\gamma, \\ \sigma\mathbf{E}\big[\max\big\{V_j+\varepsilon_j, V_k+\varepsilon_k\big\}\big] &= \ln\big(\exp(\sigma V_j)+\exp(\sigma V_k)\big)+\gamma, \end{aligned} \tag{5}$$

where γ denotes Euler's constant.[26] This step further requires to assume a value for the discount factor β, which we set in line with a 5 percent annual discount rate.

Step 3: Pay parameters We make an additional set of assumptions on wage determination in the model such that the pay parameters $\ln(W_j)$ are identified by the employer

[25] Within a strongly connected set of employers, each employer has at least one worker moving in and one worker moving out. This condition is required for the fixed-point associated with the appropriately scaled matrix of worker flows to exist. See, e.g., Jackson [2010] for details on social network definitions.

[26] Gyetvai et al. [2022] make a related point in the context of a continuous time random search model.

fixed-effect $\psi_{j(i,t)}$ in the standard two-sided unobserved heterogeneity regression

$$\ln W_{it} = X_{it}'\beta + \alpha_i + \psi_{j(i,t)} + e_{it}, \tag{6}$$

where X_{it} are exogenous worker characteristics (such as age), α_i is the unobserved worker effect, $j(i,t)$ is worker i's employer in period t, and e_{it} is an idiosyncratic shock. $\psi_{j(i,t)}$ is only identified within the connected sets of employers, which is already required to identify V_j from worker flows (as a strongly connected set is connected by definition). Specifically, we assume that employers post piece-rate wage contracts as part of the utility bundle V_j. Piece-rate contracts are a commonly used assumption in random search models with wage-posting.[27] This assumption implies that the wage premium paid by employers is log-additive in workers' experience (X_{it}), ability (α_i), and a time-varying and non-persistent productivity shock (e_{it}). Equation (6) is then correctly specified as long as individual-specific productivity does not interact with mobility decisions. We therefore also require that (i) individual-specific productivity is log-linear in the flow value of non-employment, and (ii) time-varying productivity shocks e_{it} are mean-independent conditional on the realization of i.i.d. taste shocks. These assumptions ensure that individual-specific productivity does not affect the choice to move between employers and between employment and non-employment, in line with the model of mobility described so far.

Step 4: Non-pay parameters Conditional on u_j and a measure for $\ln W_j$, the non-pay parameters can be inferred from Equation (4). To recover u_j, however, we also need to pin down the parameter σ that controls the scale of the i.i.d. taste shocks. We do this by assuming the moment condition $\text{Var}(u_j) = \text{Var}\big(\ln W_j\big)$, so the variance of flow utility is the same as the variance in the flow utility of pay $\ln W_j$. An intuitive justification for this moment condition is that the flow utility of pay can be expected to be an important component of the flow utility of a job, so their variance should be of similar magnitude. However,

[27] See, for instance, Barlevy [2008] and Engbom and Moser [2022].

this moment condition is not without loss of generality because it restricts the correlation between $\ln W_j$ and a_j to be non-positive. It is immediate to check that $\text{Var}(u_j) = \text{Var}\big(\ln W_j\big)$ implies $\text{Corr}(\ln W_j, a_j) \leq 0$. In this specific sense, the non-pay parameters we recover can be interpreted as having a compensating differentials component by definition.

4.3 Estimation

The model is estimated on the Norwegian matched employer-employee data described in Section 2. We use the establishment identifiers in the data as the direct counterpart to employer heterogeneity j in the model. We follow the employer-employee fixed-effect literature and assign each individual one main employer in each year by selecting the establishment with their largest annual earnings. Our data feature the exact start date and end date of each employment spell with an employer in each year, so we can make a precise distinction between employer-to-employer moves (a data counterpart to M_{jk} for moves between j and k) and moves with a non-employment spell in-between (data counterparts to M_{jN} and M_{Nk} for moves between j and k). Moves between two main employers that include an intervening spell with another employer also qualify as employer-to-employer moves.[28]

Our sample restrictions follow from the model structure and the conditions required for identification. In our baseline estimation, we restrict the sample to the period 2015-2019 and retain workers aged 20 to 60 (inclusive). In additional analyses, we estimate the model separately for each of the remaining three five-year periods, using the same sample restrictions as in our baseline. We also restrict attention to relatively larger and more stable employers. Specifically, we require each establishment to be observed in at least two separate years and to have at least five non-singleton workers on average in each year, where a non-singleton worker is a worker observed at some later point within the sample period. In addition, the identification conditions outlined in Section 4.2 impose extra restrictions on the set of establishments included in the estimation sample. We restrict the sample to the largest

[28]We define employer-to-employer flows as any two consecutive employment spells with a gap of at most 31 days.

strongly connected set of establishments and impose that each of these establishments hire at least one individual from non-employment.[29] These last two restrictions are interdependent, and they are imposed recursively until the set of establishments converges.

In the model, we make the assumption that workers have identical preferences over employers and that they face similar search frictions. The piece-rate wage contract assumption allows wages to depend on a worker's ability and experience, but the search parameters (such as the rate of arrival for employed workers λ_1 and offer distribution $\{f_j\}$) and preferences over employers are the same for all workers. As a result, there is no sorting on unobserved ability. Any two workers presented with the same choice of employer j and k make the same decision, up to their draw of taste shocks $(\varepsilon_j, \varepsilon_k)$. We allow for some degree of heterogeneity in preferences over employers and in search frictions by splitting the sample along observable pre-determined characteristics. We consider two demographic partitions of the data: women with and without a post-secondary education degree, and men with and without a post-secondary education degree. When splitting the sample by gender and by education, we impose the identification restrictions separately in each partition of the data.

There is a trade-off between allowing for a greater degree of observable worker heterogeneity and the mobility restrictions required for identification. Constant parameters and identical worker preferences call for considering shorter time spans and narrower worker groups (such as workers with a specific degree). Conversely, the requirement to focus on stable employers with sufficient mobility calls for longer samples and broader worker groups. In what follows, we consider the pooled sample as our baseline, but we also report results by gender and by education.

The model has employer-specific parameters $\{V_j, L_j, f_j, \rho_j, \delta_j, \psi_j, a_j\}_{j\in\mathcal{J}}$. We reduce the number of these parameters by grouping employers using a clustering algorithm. The goal is to increase the number of movers underlying each of these parameters to estimate them

[29] Focusing on the largest connected set is common in the literature. See, among others, Card et al. [2013] and Bonhomme et al. [2020] in the context of the two-sided unobserved heterogeneity regression model (6) and Sorkin [2018] and Morchio and Moser [2024] in the context of the fixed-point revealed preferences model.

more accurately. We use a k-means algorithm to create G employer groups from the J employers in each selected sample in a first pre-estimation step [Bonhomme et al., 2019]. The k-means algorithm requires two inputs: a vector of establishment-level variables on which the classification operates and a number of groups $G \leq J$. The vector of clustering variables includes wages (we use the empirical cumulative distribution function of wages in each establishment, computed at the deciles of the distribution of wages in the whole sample), job flows (we use the job creation and destruction flow rates, separately for direct transitions between employers and to and from non-employment), and some additional establishment characteristics (we use a set of indicators for the industry and occupation composition of the establishment, as well as its location).[30] We choose the number of clusters G with the objective to allow for a large degree of employer heterogeneity while still significantly reducing the number of parameters to be estimated and set $G = \text{round}(J/50)$ in our baseline.[31]

We obtain estimates of the establishment wage premium from the two-sided unobserved heterogeneity regression (6). We estimate this equation separately in each sample. Our wage measure is log average earnings per hour, where earnings are total pre-tax earnings and hours are total contract hours over the duration of each yearly employment spell. The vector of worker-level characteristics X'_{it} includes year-by-education fixed effects and a separate cubic age polynomial for each education group, restricted to be flat at age 40.[32] In line with our clustering strategy, we estimate wage premiums at the cluster level using the group identifiers obtained from the k-means algorithm.

As shown in Table 1, a minority of employers satisfy the stability and mobility requirements. Across samples, these conditions leave us with around 25 percent of the initial number of establishments. However, due to the long tail of the employer-size distribution, our estimation sample still retains around 75 percent of the initial number of workers.

[30]For location, we cut the distribution of municipality employment into deciles and construct indicators for the establishment's municipality. For example, Oslo accounts for more than 10 percent of Norwegian employment, so there is a specific indicator for establishments in Oslo.

[31]Our baseline model thus features approximately $G = \text{round}(78,133/50) \approx 1,560$ unique employer clusters. We provide robustness of our findings using $G = \text{round}(J/25)$ and $G = \text{round}(J/100)$, respectively.

[32]These are the same set of worker characteristics as in Card et al. [2013].

4.4 Results

Figure 4 shows two-way scatter plots of the parameter estimates in the baseline model. Starting with the joint distribution of the estimated overall employer values V_j and employer pay premiums ψ_j in Panel A, the key object of interest in Sorkin [2018], we find that the correlation between these model components is 0.36.[33] Next, Panel B shows the estimated overall employer values V_j along the horizontal axis and the corresponding employer non-pay values a_j along the vertical axis. We also find a positive association between V_j and employer non-pay values a_j, especially at the lower end of the value distribution. By comparison, differences in pay premiums are the only source of employer heterogeneity in the canonical wage-posting model Burdett and Mortensen [1998], which implies $a_j = 0$ at all employers, representing a useful benchmark to compare our results. Taken together, the preceding results indicate a strong positive correlation between V_j and the flow utility u_j, which captures both pay and non-pay values, as shown in Equation (4). In Panel C, we show the joint distribution of V_j and job security value $s_j = 1 - \delta_j - \rho_j$, finding a strong positive correlation of 0.55. By contrast, in Burdett and Mortensen [1998], this correlation is also zero by construction since job security is assumed to be the same across all employers.

Overall, panels A-C show that there is a positive gradient between the value of employers V_j and its pay, non-pay, and job security components. The plots also suggest that these relationships are to some degree non-linear. Differences in non-pay and job security are mostly a characteristic of employers at the lower end of the V_j distribution, while differences in pay are mostly a characteristic of employers at the upper end of the V_j distribution.

Next, we focus on the joint distributions between pay premiums ψ_j and the other sources of employer heterogeneity that determine overall employer values in our model. Importantly, in Panel D, we find a strong negative correlation of -0.58 between the estimated employer

[33]Sorkin [2018] finds a correlation of 0.53 between overall values V_j and pay premiums ψ_j based on US matched employer-employee data (see Table II in his paper). Besides a different institutional setting and time period, US data on job-to-job transitions is available only at quarterly frequency and lacks information on hours worked, while we use the exact start and end date of each spell and hourly wages.

Figure 4: Model Estimates: Employer Pay, Non-Pay and Job Security Values.

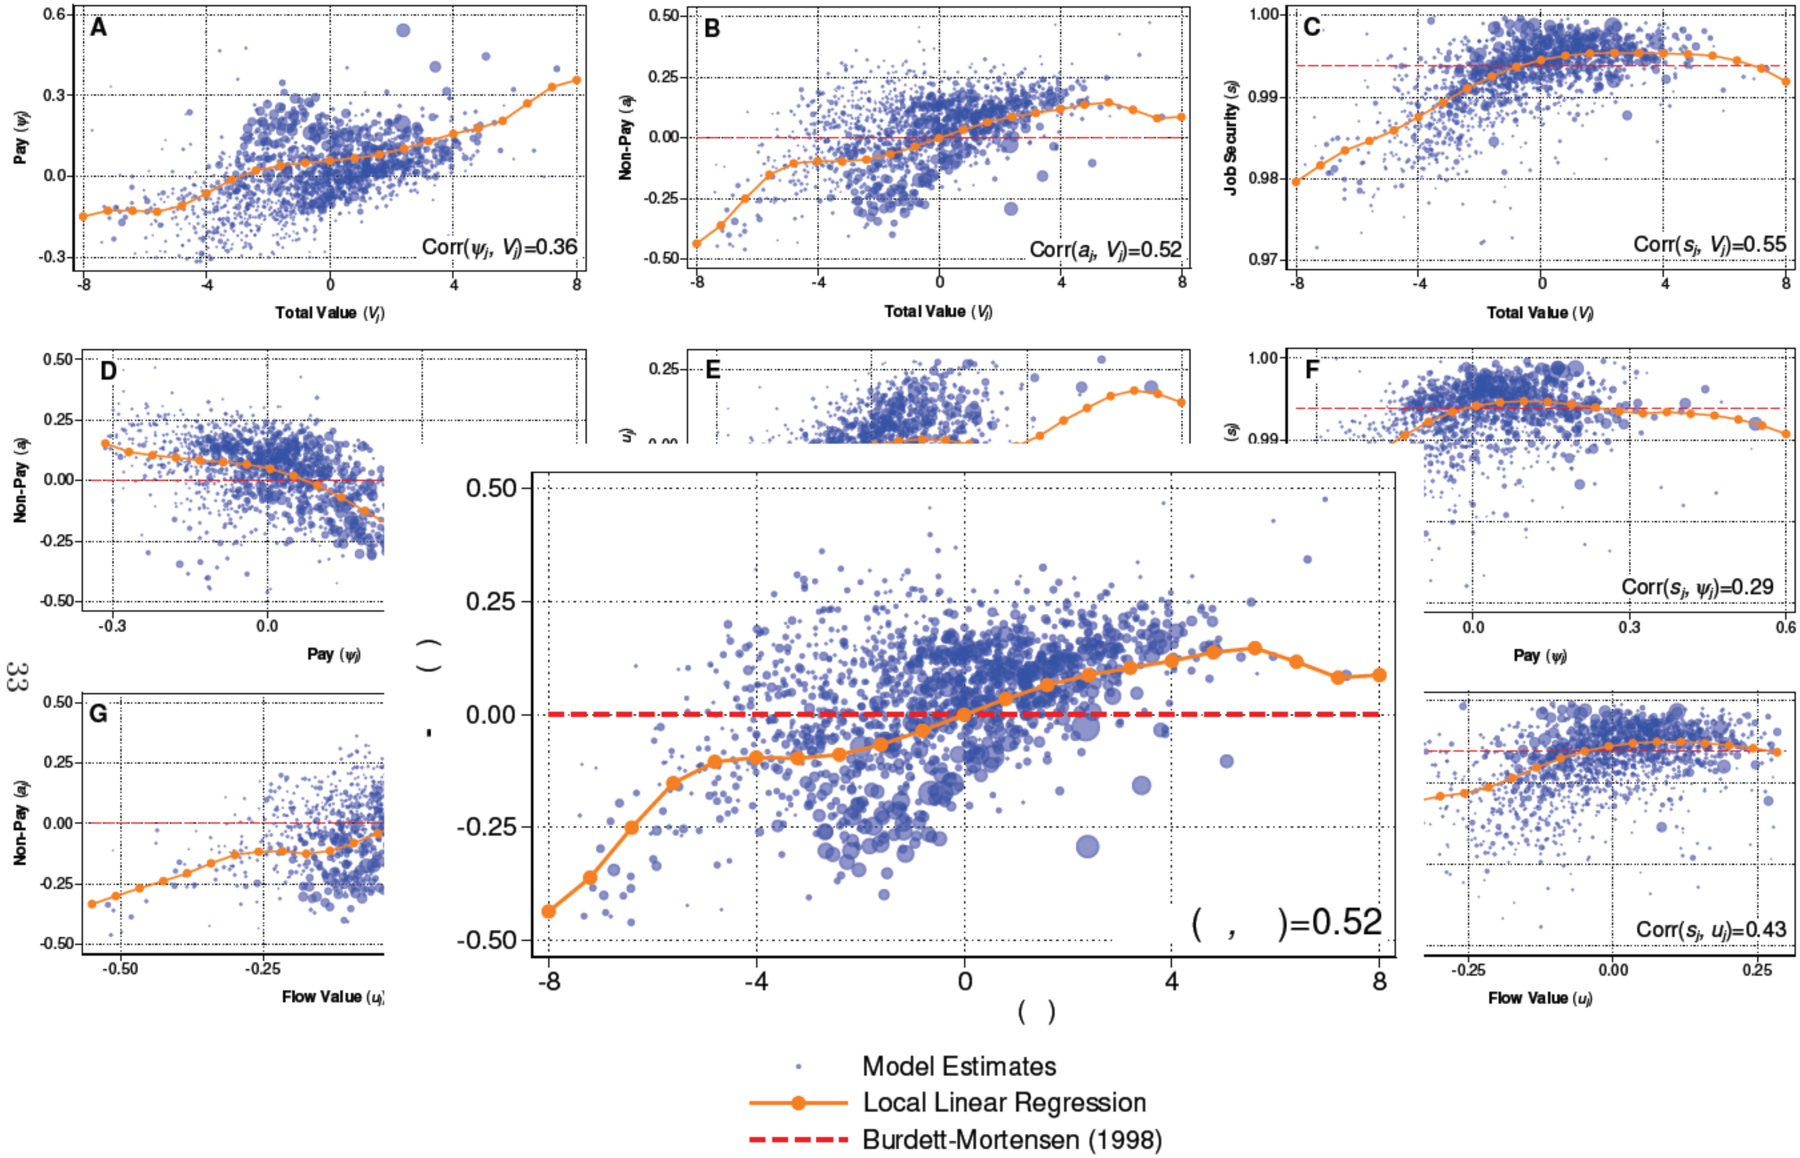

Note: The figure provides two-way scatter plots depicting our baseline model estimates for employers' overall value (V_j), flow utility value (u_j), pay value (ψ_j), non-pay value (a_j), and job security value (s_j). The size of the scatter plot markers are proportional to the number of worker-years in each cluster. Local linear regressions and correlations are weighted by worker-years. Dashed horizontal lines refers to the standard Burdett and Mortensen [1998] model, i.e., identical job security across employers and no amenities. We calculated the former as the weighted average of s_j across employers.

pay ψ_j and non-pay a_j values, reflecting sizable compensating differentials in our framework. By comparison, in the Burdett and Mortensen [1998] model, there are no systematic non-pay differences across employers ($\text{Corr}(\psi_j, a_j) = 0$). Since the correlation is clearly less than zero, there is still substantial variation in non-pay a_j at every estimated level of employer pay ψ_j.

Further, in Panel F, we show that the estimated job security component s_j and pay component ψ_j are positively related, with a correlation of 0.29.[34] In the canonical model, this correlation is also zero by assumption since job security is the same at all employers. Finally, we find that employer flow utilities u_j are positively correlated with pay and non-pay values in Panels E and G, respectively, consistent with the findings in Panels A-C, while there does not appear to be a clear association between employer job security s_j and employer non-pay as shown in Panel H.

In describing the model estimates by gender and by education, we focus on three of the main bi-variate relationships described in Figure 4: $\text{Corr}(\psi_j,V_j)$, $\text{Corr}(\psi_j,s_j)$ and $\text{Corr}(\psi_j,a_j)$. These correlations represent a good summary of the key quantitative properties of the model and are shown in Table 4. The first row in Panel A gives these correlations for the pooled sample, which correspond to the parameter estimates shown in Figure 4. Next, we report these measures by gender groups and education groups, based on separate estimations of the model for each group. The main differences are found across gender; we find a weaker correlation between employer pay and overall value for women, and consequently stronger compensating differentials for women as reflected by a more negative correlation between pay and non-pay values. Further, in Table 4, Panels B-C, we report the same correlations for alternative model specifications, where we change the number of clusters used to group employers. The overall patterns are remarkably similar when we set the number of clusters to round($J/25$) or round($J/100$), as opposed to round($J/50$) as in our baseline estimation.

To sum up, through the lens of a model with rich heterogeneity on the employer side, workers' preferences over alternative employers are only partially captured by pay premiums.

[34]This finding is in line with the estimates in Jarosch [2023], who estimates a positive correlation between match productivity and job security in a related random search framework with job-to-job transitions.

Table 4: Summary of Model Estimates.

	Correlation Between Employer Pay Value ψ_j and:			
	Overall Value V_j (1)	Job Security Value s_j (2)	Non-Pay Value a_j (3)	Number of Clusters G (4)
Panel A: Baseline Model: $G = \text{round}(J/50)$				
Pooled Sample	0.357	0.294	-0.582	1,563
Women	0.200	0.309	-0.653	850
Men	0.474	0.262	-0.518	858
College	0.314	0.307	-0.591	762
Non-College	0.339	0.291	-0.590	937
Panel B: Smaller Clusters: $G = \text{round}(J/25)$				
Pooled Sample	0.339	0.258	-0.589	3,125
Women	0.185	0.266	-0.652	1,699
Men	0.443	0.247	-0.538	1,717
College	0.294	0.255	-0.597	1,525
Non-College	0.314	0.230	-0.595	1,873
Panel C: Larger Clusters: $G = \text{round}(J/100)$				
Pooled Sample	0.366	0.326	-0.577	781
Women	0.192	0.343	-0.659	425
Men	0.497	0.275	-0.503	429
College	0.333	0.311	-0.580	381
Non-College	0.343	0.317	-0.589	468

Notes: The table shows key summary statistics of the estimated model parameters, namely the correlation between employer pay value and overall value (Corr(ψ_j,V_j)), job security value (Corr(ψ_j,s_j)), and non-pay value (Corr(ψ_j,a_j)), respectively, across alternative specifications and sample restrictions. The model estimates are for the period 2015-2019. We weight by the number of worker-year observations in each cluster.

This conclusion is robust to several alternative specifications, as shown in Table 4. As a result, the content of job ads can be relevant for other dimensions of employer heterogeneity besides differences in pay within our framework, as we study next.

5 The Information Content of Job Ads

In this section, we bring together the job attributes advertised by employers extracted from vacancy texts in Section 3 and the value of employers estimated from worker mobility and wages in Section 4. We start this analysis by providing evidence on the job attributes that are publicly advertised by employers depending on their pay and non-pay values. We then analyze the predictive power of publicly advertised job attributes for employers' actual values. Lastly, we provide a model-based measure of the information content of job ads.

5.1 What do Attractive Employers Offer in Job Ads?

Are attractive employers more likely to mention certain job attributes in publicly posted vacancies? We first investigate whether employers estimated as high-pay or low-pay advertise different types of attributes in their job postings. We estimate the regression model

$$\text{Attribute}^k_{g(j)} = \alpha^k_0 + \alpha^k_\psi \cdot \psi_{g(j)} + \alpha'_X X_{g(j)} + \varepsilon^k_{g(j)}, \tag{7}$$

at the establishment-cluster level g, where $\text{Attribute}^k_{g(j)}$ is the share of job ads from establishments in cluster g that advertise attribute k, $\psi_{g(j)}$ is their estimated pay component, $X_{g(j)}$ is a vector of occupation, industry, and location controls, and $\varepsilon^k_{g(j)}$ is an error term. The coefficient of interest is α^k_ψ, which measures to what extent employers with one standard deviation higher pay premium tend to advertise job attribute k conditional on the controls in $X_{g(j)}$.[35] We weight by the number of worker-years in each establishment-cluster.

We proceed similarly to study which attributes employers estimated as having high amenity value are more likely to advertise, substituting $a_{g(j)}$ for $\psi_{g(j)}$ in Equation (7), from which we obtain a second set of coefficients of interest α^k_a for each advertised job attribute k. While the pay and amenity estimates are negatively correlated, as shown in Panel D of

[35]For each establishment-cluster g, the vector $X_{g(j)}$ controls flexibly for the composition of workers and establishments in our matched employer-employee data, with respect to 2-digit industries, 2-digit occupations, and 10 location groups. The patterns we describe below are robust to dropping controls $X_{g(j)}$.

Figure 4, there is substantial variation in the employer non-pay values also conditional on the employer pay premiums, and therefore, α_a^k is not mechanically equal to $-\alpha_\psi^k$.

Figure 5 reports the estimated α_ψ^k and α_a^k for all attributes we derive from the vacancy data. Panel (a) shows that high-pay employers are more likely to mention characteristics associated with a more generous compensation package ("pension/insurance scheme," "competitive pay," "collective agreement pay") and career opportunities ("career paths," "on-the-job training"). Perhaps more surprisingly, we also find that several attributes related to the workplace, such as "good colleagues" and "social environment," are associated with high-pay employers. At the other end of the spectrum, we see that most low-pay employers are more likely to specify the type of hours and the type of contract in their job postings. Panel (b) further shows that high amenity employers are more likely to specify contract duration and hours-related attributes in the text of their job vacancies, especially in relation to a predictable schedule ("regular daytime schedule") or flexible schedule ("possibility to work flexible hours"). Conversely, several career-related ("on-the-job training") and workplace-related ("social environment," "good colleagues") attributes are negatively associated with amenity values. Non-standard hours ("shift work") is also more frequently advertised by employer with low amenity value. Finally, Figure 5 shows that many estimated α_ψ^k and α_a^k are not statistically different from zero. We therefore fail to detect any correlation between many advertised attributes and the estimated value of the posting employers.[36]

Overall, the correlations that we document in Figure 5 are consistent with the existence of compensating differentials, in line with the negative correlation between the pay and amenity model estimates (Panel D in Figure 4). Notably, jobs offered by employers with higher non-pay value appear to more frequently have convenient hours, echoing the findings obtained in experimental settings [Mas and Pallais, 2017]. However, a more accurate interpretation of the coefficients α_ψ^k and α_a^k is as measures of what high-pay and high non-pay employers are more likely to advertise in their ads. The advertised attributes need not translate directly

[36]Appendix Figure A.5 shows results corresponding to Figure 5 for the overall employer value V_j, while Appendix Figures A.6–A.7 show results, respectively, using smaller and larger employer-clusters.

Figure 5: Publicly Advertised Job Attributes By Employer Pay and Non-Pay Values.

(a) Pay (ψ_j) (b) Non-Pay (a_j)

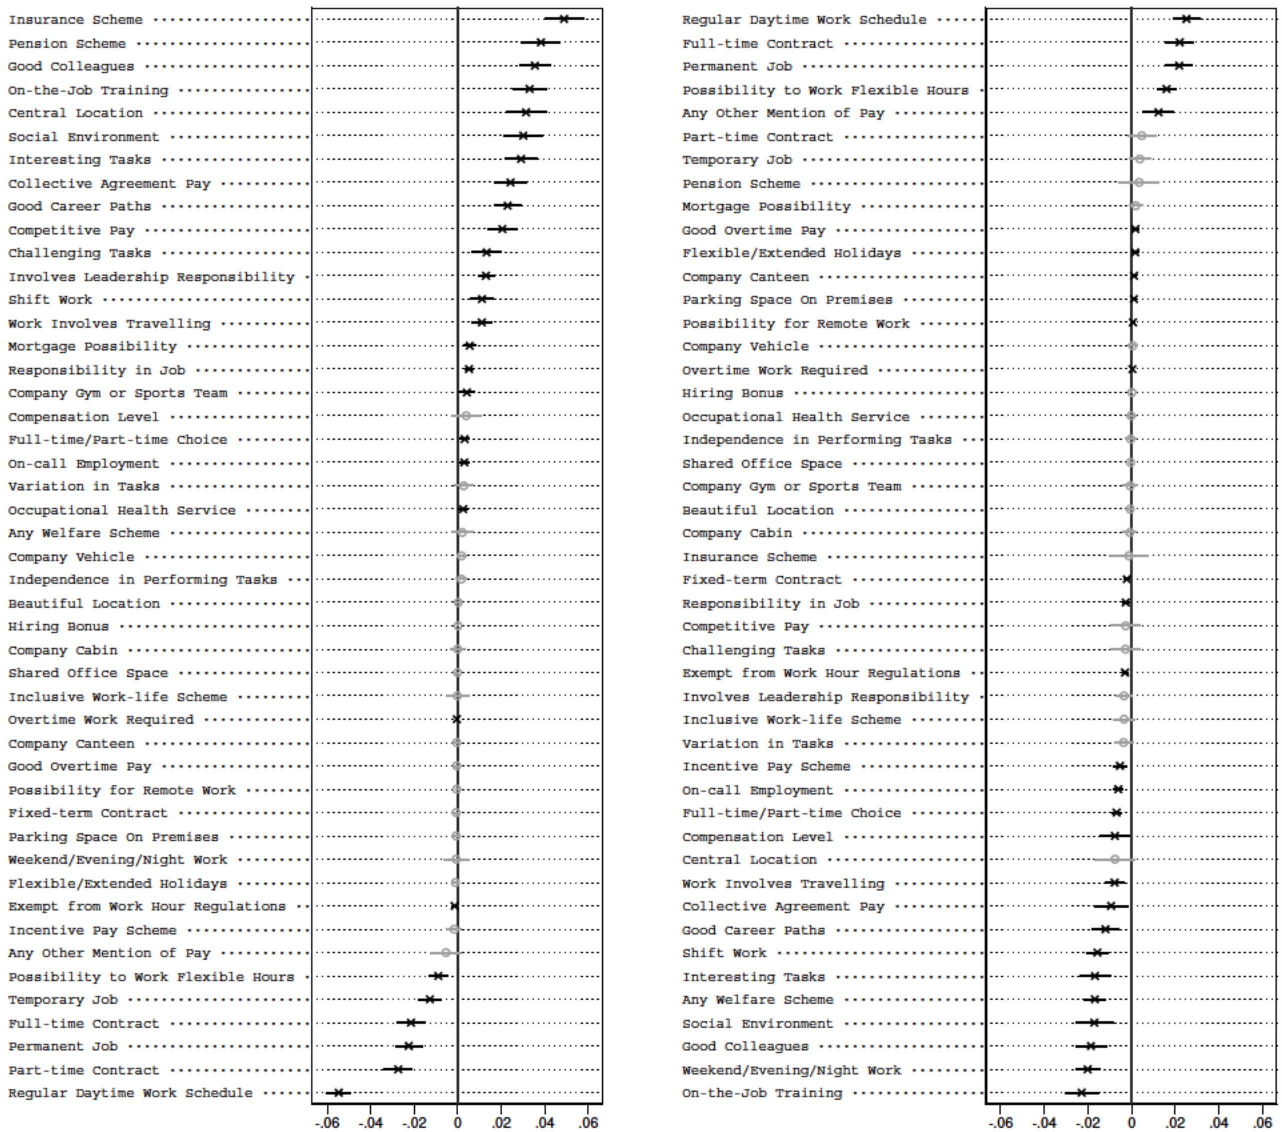

Note: This figure shows parameter estimates and 90% confidence intervals from separate regressions of the share of each ad attribute on employer pay (ψ_j) and non-pay (a_j) values, capturing the change in the fraction of ads with each job attribute associated with a standard deviation increase in employer value. Estimations are done at the employer-cluster level for the five-year period from 2015 to 2019, and are weighted by the number of worker-years. We control for occupation and industry using two-digit occupation and two-digit industry fixed effects and location based on deciles of the number of workers in the municipality of the firm (Oslo as an own group). Appendix Figure A.5 shows corresponding results by overall employer values (V_j) for the baseline model specification, while Appendix Figures A.6-A.7 shows results for model estimates using smaller and larger employer clusters, respectively.

into attributes of the actual jobs for the following reasons. First, it is unclear to what extent employers can deliver on many of these attributes, especially given that some of them are vague ("social environment," "good colleagues"). Second, employers can omit specific characteristics from the text of the ad even though they are actual attributes of the position. As an example, they may actually offer a competitive salary without mentioning it explicitly in the corresponding vacancy text. As such, we are cautious in interpreting the pay and non-pay attributes mentioned in job ads as reflecting actual job attributes in general.

5.2 Predicting the Value of Employers from Job Ads

How predictive are publicly advertised job attributes of the employers' actual pay and non-pay values? To assess the predictive power of the pay and non-pay content of job ads, we now estimate the following regression models, with the overall value of employers $V_{g(j)}$ or one of its components (i.e., pay $\psi_{g(j)}$ or non-pay $a_{g(j)}$) as the outcome variable:

$$V_{g(j)} = \beta_0 + \sum_k \beta_k \cdot \text{Attribute}^k_{g(j)} + \beta'_X X_{g(j)} + \varepsilon^V_{g(j)}. \tag{8}$$

In Equation (8), $\text{Attribute}^k_{g(j)}$ is the share of ads posted by employers in establishment-cluster g that features job attribute k, $X_{g(j)}$ is a set of industry, location and occupation controls that is meant to capture other commonly observed characteristics listed in job ads, and $\varepsilon^V_{g(j)}$ is an error term. We weight by the number of worker-years in each cluster.

Summarizing the overall predictive power of job ads, Figure 6 shows the adjusted-R^2 of the regression model (8) with and without controlling for the variables in X_g.[37] This plot shows that the information contained in job ads meaningfully predicts overall employer

[37] The formula used to calculate predictive power is as follows:

$$\text{Adjusted } R^2 = 1 - \frac{SS_{res}/df_{res}}{SS_{total}/df_{total}}$$

where SS_{res} is residual sum of squares, df_{res} is the number of establishment clusters (approximately 1,560 in our baseline) minus the number of explanatory variables, and SS_{total} and df_{total} are from regressions on a constant term only. The partial R^2 is equal to the share of the variation that is unexplained in a regression values and their components. We find that the non-pay attributes (blue bars) are about as predictive of the overall employer values as the pay-related attributes (red bars). In total, solely based on the job attributes we are able to detect in ads (i.e., without including the controls in $X_{g(j)}$), we are able to explain about 60 percent of the variation in employer pay and their overall values, and about 50 percent of non-pay and 40 percent of the variation in job security. Once we purge the information contained in X_g, we find that detected job attributes still explain around 10 to 20 percent of the remaining variation in employer values, as shown by the partial-R^2 in Figure 6.[38]

We further delve into the adjusted-R^2 measures reported in Figure 6 with the following two exercises. First, we decompose the adjusted-R^2 measures reported above into contributions from different categories of job attributes in Appendix Figure A.8. Panel (a) documents that financial attributes explain a large share of the variation in employer pay values, while hours of work and convenient hours contain some of the residual variation. Similarly, Panel (b) shows that a large fraction of the variation in non-pay employer values can be explained by detected job attributes related to convenient and inconvenient hours, while pay-related attributes contribute less to explanatory power. In a similar spirit, in Appendix D, we further document that the patterns discussed here are robust to focusing on specific job attributes and the order in which we add various attributes. Specifically, we show evidence from a series of regressions where we sequentially add the individual job attributes that provide the largest increase in predictive power. Based on this approach, we confirm that a substantial share of the predictive power with respect to employer pay values is contained in the information related to financial attributes, while the information on convenient or inconvenient hours remains key in the predictions of employer non-pay values.

---

on baseline controls only but explained when data extracted from job ads is added to the right hand side:

$$\text{Partial } R^2 = \frac{R^2_{full} - R^2_{baseline}}{1 - R^2_{baseline}}$$

where R^2_{full} ($R^2_{baseline}$) is adjusted-R^2 from a regression on baseline controls and (without) job attributes.

[38] Appendix Figure A.10 shows that alternatively using smaller or larger employer-clusters yields the same overall pattern in terms of the predictive power of ad attributes for each component of the value of employers.

Figure 6: Predictive Power of Publicly Advertised Job Attributes.

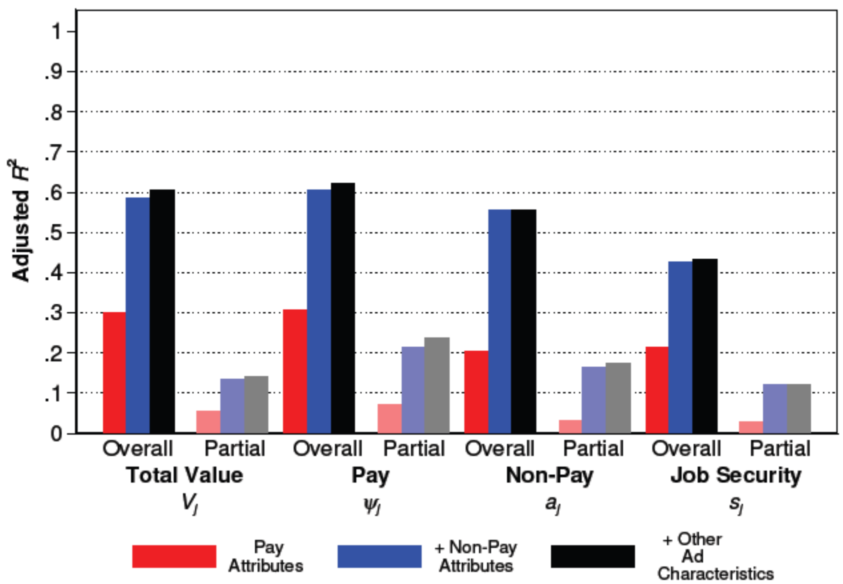

Note: This figure shows adjusted R^2s from regressions of estimated model parameters on pay and non-pay job attributes and other ad characteristics retrieved in the text analysis, including the number of words in the ad, indicators for ad posted online, disclosure of company name, and the overall number of attributes in the ad. Estimations are done at the employer-cluster level for the five-year period from 2015 to 2019, and are weighted by the number of worker-years in each cluster. Overall R^2s are from regressions that include ad attributes and other ad characteristics, while partial R^2s are from regressions that also control for the composition of industry, occupation and location in each cluster. Industry and occupation controls are defined at the two-digit level. Location indicators split local municipalities into deciles, with municipalities in the same group having similar number of workers (Oslo as an own group). The sample corresponds to the last row in Table 1. Appendix Figure A.10 reports the corresponding results using employer values estimated on larger and smaller clusters, while Appendix Figure A.11 provides results separately for each of the four five-year periods between 2000 and 2019.

Second, we consider the predictive power of job ads across different industries in Appendix Figure A.9. We estimate the regression models in Equation (8) separately by industry.[39] While job ads hold strong predictive power for the overall values of employers across all industry groups, we do find sizeable differences across industries. Notably, the adjusted-R^2 ranges from 0.5 in the Construction industry to 0.8 in the Retail and Services industry.

[39] We assign an industry to each cluster g based on modal industry among the establishments in the cluster.

Posted Employer Value and Vacancy Duration Next, we consider how the employers' values posted in job ads–as can be predicted using the regression model in Equation (8)–relate to the average duration of vacancies that these employers open. Recent evidence provided by Mueller et al. [2024] and Bassier et al. [2023], respectively, shows that the employer pay premiums and the posted pay are negatively related to the duration of posted vacancies, consistent with the notion that high-pay employers are more attractive in the labor market. In Appendix Table A.4, we show that the average vacancy duration is 13.9 percent shorter for job ads with one standard deviation higher overall posted employer value, while controlling flexibly for industries, occupations and locations listed in the job ads. Decomposing the overall posted employer value in its various components, we find a vacancy duration elasticity of -0.7 with respect to the posted pay premium and an elasticity of about -1 with respect to the posted amenity value. Further, we get a negative relationship between posted job security and vacancy duration, finding 5.9 percent shorter average duration for job ads with a standard deviation higher job security value. These results show that both pay premiums, amenity values and job security shape the overall attractiveness of employers.

Heterogeneity by Gender and by Education Returning to the regression model in Equation (8), Figure 7 shows the adjusted-R^2 by gender and by education, with and without controlling for $X_{g(j)}$. The outcome variables in these regressions–the employer values–are estimated separately for each sub-sample, and all regressions are weighted by the corresponding number of worker-years in each sub-sample. These plots show that the information contained in vacancy texts meaningfully predicts the value of employers in all sub-samples. Across sub-samples, the magnitude of the overall R^2s is around 60 percent for the estimated overall employer value, broadly similar to those reported for the pooled sample in Figure 6.

Overall, Figure 7 suggests that the most salient difference across sub-samples in terms of the explanatory power of the content of job ads is found between college and non-college workers. In particular, the explanatory power of ads is at least 10 percentage points lower

Figure 7: Heterogeneity in the Predictive Power of Publicly Advertised Job Attributes.

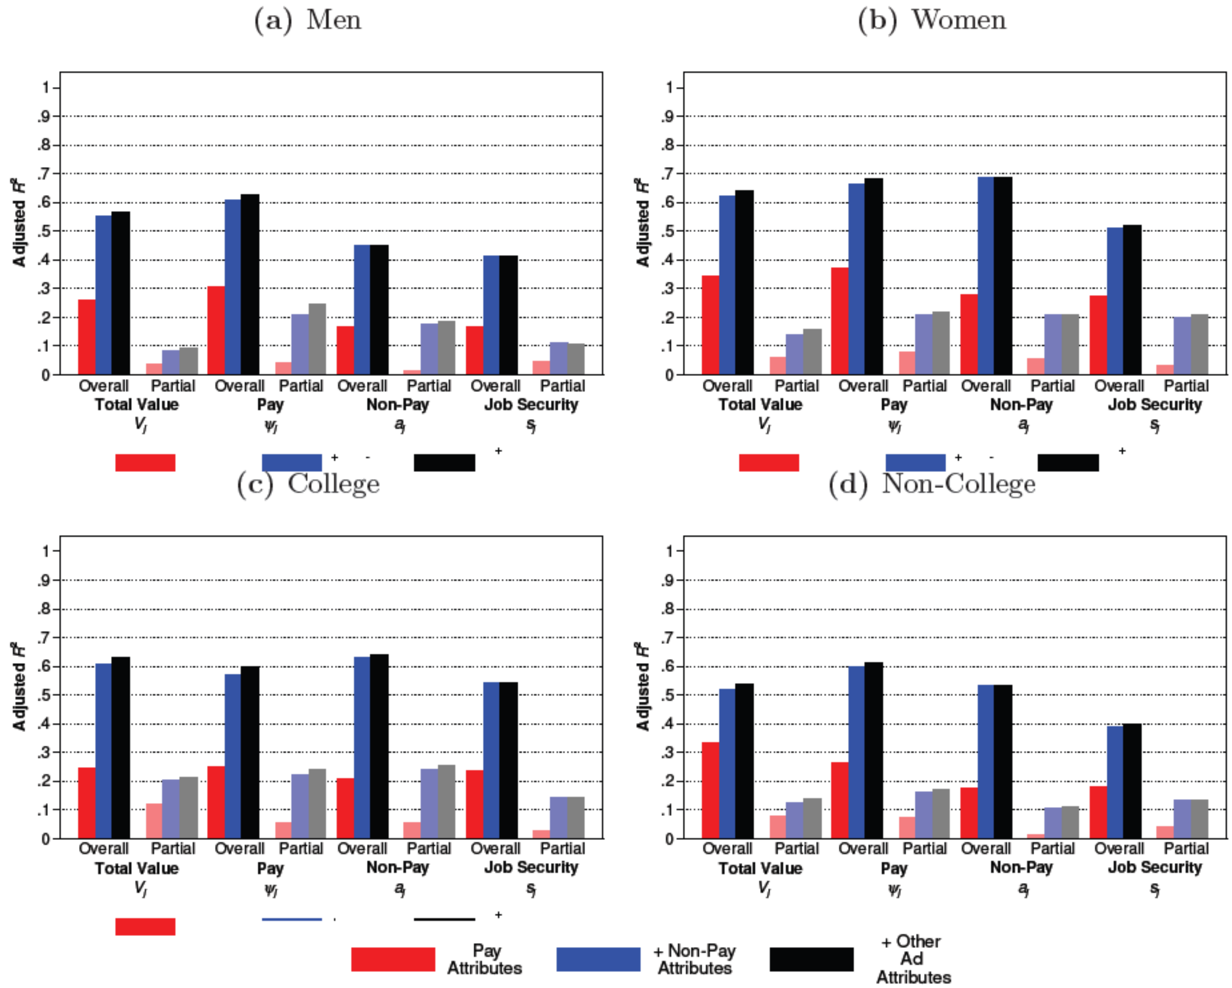

Note: This figure shows adjusted R^2s from regressions of estimated population group-specific model parameters on pay and non-pay job attributes and other ad characteristics retrieved in the text analysis, including the number of words in the ad, indicators for ad posted online, disclosure of company name, and the overall number of attributes in the ad. Estimations are done at the employer-cluster level for the five-year period from 2015 to 2019, separately by gender (Panels (a)-(b)) and education (Panels (c)-(d)), and are weighted by the number of worker-years in each cluster and population group. Overall R^2s are from regressions that include ad attributes and other ad characteristics, while partial R^2s are from regressions that also control for the composition of industry, occupation, and location in each cluster. Industry and occupation controls are defined at the two digit level. Location indicators split local municipalities into deciles, with municipalities in the same group having similar number of workers (Oslo as an own group). Appendix Figure A.12 shows a corresponding plot with weights fixed across sample splits using the total number of worker-years.

for the estimated overall and non-pay employer values for non-college workers relative to college workers, both for the overall and partial R^2. By comparison, the overall R^2 suggests that the content of job ads has a higher predictive power in terms of employers' non-pay value for women than for men. However, this result disappears once we control for $X_{g(j)}$.[40]

5.3 A Model-Based Measure of Ad Informativeness

What is the information content of job ads in terms of pay and non-pay attributes offered by employers? We now use the link between the estimated value of employers and the information extracted from vacancy texts to provide a measure of ad informativeness based on the model introduced in Section 4. Within the model, the R^2 calculations can be translated into upward and downward mobility flows. Specifically, we introduce the notation

$$\widehat{M}_{jk} = L_j \cdot \Big[\rho_j \cdot f_k + (1 - \rho_j - \delta_j) \cdot \lambda_1 \cdot f_k \cdot \Pr(\widehat{V}_k \geq V_j)\Big], \tag{9}$$

for the job-to-job flows from employer j to employer k, where $\widehat{V}_k$ is the value of employers predicted from one of the regression models in Equation (8).[41] In words, we interpret $\widehat{M}_{jk}$ as the counterfactual flows from employer j to employer k if workers were to use the information advertised on employer k's vacancies to decide whether to move. We further make a distinction between the job-to-job flows where the actual value of the destination employer V_k is greater or lower than the value of employer j, which we denote, respectively, $\widehat{M}_{jk}(\uparrow) = \widehat{M}_{jk} \cdot \mathbf{1}\{V_k \geq V_j\}$ and $\widehat{M}_{jk}(\downarrow) = \widehat{M}_{jk} \cdot \mathbf{1}\{V_k < V_j\}$.

We compute economy-wide summary measures of overall counterfactual job-to-job mobility by aggregating across all bilateral employer flows as $\widehat{M} = \sum_{j',k'} \widehat{M}_{j'k'}$ for each alternative prediction model. We proceed similarly to aggregate job-to-job flows to employers whose actual value is above and below the worker's current employer, with, respectively, $\widehat{M}(\uparrow) = \sum_{j,k} \widehat{M}_{jk}(\uparrow)$ providing a summary measure of upward job mobility and $\widehat{M}(\downarrow) = \sum_{j,k} \widehat{M}_{jk}(\downarrow)$ providing a summary measure of downward job mobility.

In Table 5, we investigate how job-to-job mobility changes under alternative information

[40] Appendix Figure A.11 provides similar heterogeneity results over time, separately for each of the four five-year periods between 2000 and 2019. These results do not suggest a clear time profile.

[41] The regression models (8) are weighted by the job offer distribution f_j. This ensures that the average job offer $\sum_{j'} f_{j'} \widehat{V}_{j'}$ remains the same across specification by construction.

Table 5: Counterfactual Analysis: Informativeness of Job Ads for Job-to-Job Mobility.

	Relative Change in Job-to-Job Mobility (%)								
	Pay			Pay & Non-Pay			Full Model		
	Δ Mobility			Δ Mobility			Δ Mobility		
	Overall	↑	↓	Overall	↑	↓	Overall	↑	↓
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
% change relative to baseline scenario:									
Pooled sample	0.1	0.5	-0.3	0.3	1.1	-0.8	1.1	4.3	-3.2
Gender:									
Women	-0.0	0.5	-0.5	0.2	1.2	-1.0	0.7	3.7	-3.0
Men	0.1	0.2	-0.1	0.4	0.8	-0.4	1.3	3.6	-2.3
Education:									
College	0.1	0.5	-0.4	0.4	1.4	-1.0	0.9	4.0	-3.1
Non-college	0.3	0.6	-0.3	0.6	1.1	-0.5	1.6	4.1	-2.5

Note: The table reports relative changes in job-to-job mobility rates (in percent) under alternative information sets used to predict the overall employer values, as compared to a benchmark scenario where workers can only use a linear projection based on industry, occupation, and location. In Columns (1)-(3), we report mobility rates where workers use pay information extracted from job ads, besides industry, occupation and location, while in Columns (4)-(6), workers use information on both pay and non-pay attributes as well as other ad characteristics. Finally, in Columns (7)-(9), we use the actual employer values obtained from the model. The columns denoted by ↑ and ↓ show the rates of changes in upward job mobility ($V_k \geq V_j$) and downward job mobility ($V_k < V_j$), respectively.

sets used to predict overall employer values. All results are shown relative to the benchmark linear prediction model in Equation (8) with occupation, industry, and location controls, but no information from vacancy texts. The numbers reported in the top row suggest that having detailed information on job attributes from job ads would increase (reduce) inflows of workers to better (worse) jobs by 1.1 (0.8) percent, as compared to the benchmark scenario where only "one-liner" vacancies listing occupation, industry, and location are posted. This suggests that publicly advertised information on both pay and non-pay attributes can have important implications for the strength and direction of job-to-job flows. Using our full model, we find that upward mobility goes up by 4.3 percent relative to the benchmark. The corresponding figure for downward mobility is 3.2 percent. This result suggests that publicly advertised information on job attributes explains around 25 percent of the upward mobility in job-to-job flows that we observe in our model. Using this measure of ad informativeness, we also find evidence of heterogeneity by gender and by education. The heterogeneity results in Table 5 imply that the information on job attributes in ads explains slightly more of the upward mobility in job-to-job flows for women than for men (30 percent vs 20 percent) and for college graduates than for non-college workers (35 percent vs 25 percent).

In Table 6, we extend the analysis to upward and downward mobility in employer pay and amenity values. Using the same linear models to define counterfactual mobility flows, we now quantify to what extent detailed ad information translates into additional mobility toward better-paying employers $\sum_{j,k} \widehat{M}_{jk} \cdot \mathbf{1}\{\psi_k \geq \psi_j\}$ or employers offering better amenities $\sum_{j,k} \widehat{M}_{jk} \cdot \mathbf{1}\{a_k \geq a_j\}$. We similarly define downward mobility in pay and non-pay.

We draw two main conclusions from this exercise. First, we find that the magnitude of the changes in job-to-job flows to employers offering better pay and better non-pay is smaller across prediction models, as compared to Table 5. For instance, in the top row of Table 5, Column (5), the increase in job-to-job flows to employers with higher value is 1.1 percent using the predictive model with both pay and non-pay attributes in the pooled sample. As shown in Table 6, Column (3), job-to-job flows to employers offering better pay (non-pay) increase by only 0.2 percent (0.6 percent). These results reflect the relatively limited correlation between the overall value of employers V_j and its components: the correlation between ψ_j and V_j is 0.36 in the pooled sample; see Table 4. Second, this exercise suggests that there is heterogeneity in the direction of job-to-job flows predicted by the pay and non-pay attributes contained in vacancy texts. Focusing on gender differences, we find that men are more likely to move to employers offering better pay as their information set is expanded. Relative to the baseline model, their job-to-job mobility to better paying employers increases by 0.3 percent using the full set of attributes advertised in vacancies (Panel A, Column (3) in Table 6). By contrast, women's job-to-job mobility to employers offering higher pay decreases with the full set of advertised attributes. The flip side of this decrease in job-

Table 6: Counterfactual Analysis: Informativeness of Job Ads for Job-to-Job Mobility.

	Relative Change in Job-to-Job Mobility (%)					
	Pay		Pay & Non-Pay		Full Model	
	Δ Mobility		Δ Mobility		Δ Mobility	
	↑	↓	↑	↓	↑	↓
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Pay mobility						
% change relative to baseline scenario:						
Pooled sample	0.3	-0.1	0.2	0.1	0.5	0.7
Gender:						
Women	-0.1	0.0	-0.4	0.6	-0.0	0.7
Men	0.2	-0.1	0.3	0.1	0.8	0.5
Education:						
College	-0.0	0.1	-0.1	0.5	0.3	0.6
Non-college	0.8	-0.5	0.9	-0.3	1.4	0.2
Panel B: Non-pay mobility						
% change relative to baseline scenario:						
Pooled sample	0.1	0.0	0.6	-0.3	2.6	-1.5
Gender:						
Women	0.3	-0.3	1.1	-0.8	2.2	-1.5
Men	0.0	0.1	0.4	-0.1	2.3	-1.0
Education:						
College	0.5	-0.4	1.3	-0.9	2.7	-1.8
Non-college	-0.1	0.4	0.3	0.3	2.0	-0.4

Note: The table reports relative changes in job-to-job mobility rates (in %) under alternative information sets used to predict the overall employer values, as compared to a benchmark scenario where workers can only use a linear projection based on industry, occupation, and location. Panel A reports upward and downward mobility in Pay (ψ_j) for each sub-sample. Panel B reports upward and downward mobility in Non-pay (a_j) for each sub-sample. The columns denoted by ↑ shows the rate of change in upward mobility in Pay ($\psi_k \geq \psi_j$) and Non-pay ($a_k \geq a_j$), respectively. The columns denoted by ↓ shows the rate of change in downward mobility in Pay ($\psi_k < \psi_j$) and Non-pay ($a_k < a_j$), respectively.

to-job flows to higher-paying employers is a one percent increase in the flows to employers offering better amenities. The intuition behind the gender results is straightforward. In the estimated model, pay is a more important component of the value of employers for men than for women, as shown in Table 4. We also find that advertised attributes are slightly more informative of the value of employers for women than for men (Table 5). As a result, with more information, women self-select more into high-amenity jobs. This last finding is consistent with gender differences in preferences over employer attributes carrying over to the job search stage [e.g., Le Barbanchon et al., 2021, Fluchtmann et al., 2024].

6 Conclusion

In this paper, we study the pay and amenity content of job ads. Our first set of results is derived from analyzing the texts of a comprehensive database of job ads. We find that the texts of these ads specify many job attributes relevant to workers' preferences over alternative employers. Almost every job ad mentions at least one of these attributes with an average of six attributes per ad. By considering a wide range of advertised job attributes, our analysis adds to prior work focusing on specific content, such as pay information [Batra et al., 2023], remote work [Hansen et al., 2023], or flexible work arrangements [Adams-Prassl et al., 2023a].

We then recover workers' valuation for alternative employers using a structural model estimated on matched employer-employee data. Our second set of results follows from linking these estimates to the pay and non-pay content of job ads. High-pay employers are more likely to advertise aspects of the compensation package, while high-amenity employers more often mention working arrangements, such as contract duration and hours worked, in their postings. We confirm using several alternative metrics that the content of job ads has predictive power for workers' valuation of employers above other observable characteristics, such as industries and occupations. Taken together, these results suggest that the pay and non-pay content of job vacancies represents one reliable source of information for workers assessing their options in the labor market.

Our approach based on extracting a wide range of job attributes from vacancy texts points to several directions for future work. First, the pay and non-pay content of job postings can be expected to differ across labor markets. As one example, health insurance is largely absent from job ads in Norway, where healthcare is universal, but it can be expected to be an important job attribute wherever employers are commonly involved, such as in France or the US. Second, there might be a cyclical element to the content of job ads, with recruiters potentially advertising more and different attributes in tight labor markets. Third, while prior experimental work has focused on the salary component of vacancies [Belot et al., 2022], our results show that there is scope to manipulate a much broader range of job attributes to further unpack the search behavior of job seekers.

Finally, our results have implications for the design of policies related to the information contained in job ads. Several jurisdictions have passed legislation mandating employers to disclose a salary range in their job openings. Our analysis suggests that there is scope to regulate the provision of information on additional job characteristics besides pay, such as working arrangements. Similarly, our results can be useful for the design of online job platforms in that they provide guidance on the attributes employers should be encouraged to provide in their postings.

References

D. Acemoglu, D. Autor, J. Hazell, and P. Restrepo. Artificial intelligence and jobs: evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340, 2022.

A. Adams-Prassl, M. Balgova, M. Qian, and T. Walters. Firm concentration & job design: The case of schedule flexible work arrangements. 2023a.

A. Adams-Prassl, T. Le Barbanchon, and A. Marcato. On-the-job training and labor market concentration. *Available at SSRN 4477642*, 2023b.

L. Alekseeva, J. Azar, M. Gine, S. Samila, and B. Taska. The demand for ai skills in the labor market. *Labour economics*, 71:102002, 2021.

B. W. Arold, E. Ash, W. B. MacLeod, and S. Naidu. Do Words Matter? The Value of Collective Bargaining Agreements. Technical report, 2024.

E. Atalay, P. Phongthiengtham, S. Sotelo, and D. Tannenbaum. The evolution of work in the united states. *American Economic Journal: Applied Economics*, 12(2):1–34, 2020.

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

J. Bagger and R. Lentz. An empirical model of wage dispersion with sorting. *Review of Economic Studies*, 86(1):153–190, 2019.

S. Banfi and B. Villena-Roldan. Do high-wage jobs attract more applicants? Directed search evidence from the online labor market. *Journal of Labor Economics*, 37:715–746, 2019.

G. Barlevy. Identification of search models using record statistics. *Review of Economic Studies*, 75(1):29–64, 2008.

I. Bassier, A. Dube, and S. Naidu. Monopsony in Movers The Elasticity of Labor Supply to Firm Wage Policies. *Journal of Human Resources*, 57:S50–S86, 2022.

I. Bassier, A. Manning, and B. Petrongolo. Vacancy Duration and Wages. Technical report, 2023.

H. Batra, A. Michaud, and S. Mongey. Online job posts contain very little wage information. Technical report, National Bureau of Economic Research, 2023.

M. Belot, P. Kircher, and P. Muller. How wage announcements affect job search—a field experiment. *American Economic Journal: Macroeconomics*, 14(4):1–67, 2022.

M. Bhuller, D. Ferraro, A. R. Kostøl, and T. C. Vigtel. The internet, search frictions and aggregate unemployment. Technical report, 2023.

D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.

S. Bonhomme, T. Lamadon, and E. Manresa. A distributional framework for matched employer employee data. *Econometrica*, 87(3):699–739, 2019.

S. Bonhomme, K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler. How much should we trust estimates of firm effects and worker sorting? 2020.

J. C. Braxton and B. Taska. Technological change and the consequences of job loss. *American Economic Review*, 113(2):279–316, 2023.

C. Brown. Equalizing differences in the labor market. *Quarterly Journal of Economics*, 94 (1):113–134, 1980.

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

D. Card, J. Heining, and P. Kline. Workplace heterogeneity and the rise of west german wage inequality. *Quarterly journal of economics*, 128(3):967–1015, 2013.

D. Card, A. R. Cardoso, J. Heining, and P. Kline. Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1):S13–S70, 2018.

C. Carrillo-Tudela, L. Kaas, and B. Lochner. Matching through search channels. 2023.

Z. Contractor and B. Taska. The effect of software adoption on skill demand. 2022.

D. R. Cox and N. Wermuth. A comment on the coefficient of determination for binary responses. *The American Statistician*, 46(1):1–4, 1992.

Z. Cullen. Is pay transparency good? *Journal of Economic Perspectives*, 38(1):153–180, 2024.

N. Datta. Willing to pay for security: a discrete choice experiment to analyse labour supply preferences. Technical report, 2019.

D. Deming and L. B. Kahn. Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, 36(S1):S337–S369, 2018.

D. J. Deming and K. Noray. Earnings dynamics, changing job skills, and stem careers. *Quarterly Journal of Economics*, 135(4):1965–2005, 2020.

S. Desiere and C. Walter. The Shift Premium: Evidence from a Discrete Choice Experiment. Technical report, 2023.

M. Drake, N. Thakral, and L. T. To. Wage Differentials and the Price of Workplace Flexibility. Technical report, 2023.

G. J. Duncan and B. Holmlund. Was adam smith right after all? another test of the theory of compensating wage differentials. *Journal of Labor Economics*, 1(4):366–379, 1983.

N. Engbom and C. Moser. Earnings inequality and the minimum wage: Evidence from brazil. *American Economic Review*, 112(12):3803–47, 2022.

J. Faberman and G. Menzio. Evidence on the Relationship between Recruiting and the Starting Wage. *Labour Economics*, 50:67–79, 2018.

J. Fluchtmann, A. M. Glenny, N. A. Harmon, and J. Maibom. The gender application gap: Do men and women apply for the same jobs? *American Economic Journal: Economic Policy*, 16(2):182–219, 2024.

A. Gyetvai, P. Arcidiacono, A. Maurel, and E. Jardim. Identification and estimation of continuous-time job search models with preference shocks. Technical report, 2022.

S. Hansen, P. J. Lambert, N. Bloom, S. J. Davis, and B. Taska. Remote work across jobs, companies, and space. Technical report, 2023.

B. Hirsch, E. J. Jahn, A. Manning, and M. Oberfichtner. The Wage Elasticity of Recruitment. Technical report, 2022.

J. J. Horton, R. Johari, and P. Kircher. Sorting through cheap talk: Theory and evidence from a labor market. Technical report, CESifo Working Papers No. 11163, 2024.

H.-s. Hwang, D. T. Mortensen, and W. R. Reed. Hedonic wages and labor market search. *Journal of Labor Economics*, 16(4):815–847, 1998.

M. O. Jackson. *Social and Economic Networks*. Princeton University Press, paperback edition, 11 2010. ISBN 978-0691148205.

S. Jäger, C. Roth, N. Roussille, and B. Schoefer. Worker beliefs about outside options. *Quarterly Journal of Economics*, 139(23):1505–1556, 2024.

G. Jarosch. Searching for job security and the consequences of job loss. *Econometrica*, 91 (3):903–942, 2023.

G. Jolivet, F. Postel-Vinay, and J.-M. Robin. The empirical content of the job search model: Labor mobility and wage distributions in europe and the us. *Contributions to Economic Analysis*, 275:269–308, 2006.

L. Lagos. Union Bargaining Power and the Amenity-Wage Tradeoff. Technical report, 2024.

T. Lamadon, M. Mogstad, and B. Setzler. Imperfect competition, compensating differentials, and rent sharing in the us labor market. *American Economic Review*, 112(1):169–212, 2022.

T. Le Barbanchon, R. Rathelot, and A. Roulet. Gender differences in job search: Trading off commute against wage. *Quarterly Journal of Economics*, 136(1):381–426, 2021.

P. Lewandowski, K. Lipowska, and M. Smoter. Working from Home during a Pandemic - A Discrete Choice Experiment in Poland. Technical report, 2022.

I. Marinescu and R. Wolthoff. Opening the black box of the matching function: The power of words. *Journal of Labor Economics*, 38(2):535–568, 2020.

A. Mas and A. Pallais. Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–3759, 2017.

D. McFadden. Conditional logit analysis of qualitative choice behavior. In P. Zarembka, editor, *Frontiers in Econometrics*, pages 105–139. New York [u.a.] : Academic Press, 1974.

T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. 2013. doi: 10.48550/ARXIV.1301.3781.

I. Morchio and C. Moser. The gender pay gap: Micro sources and macro consequences. Technical report, National Bureau of Economic Research, 2024.

A. I. Mueller, D. Osterwalder, J. Zweimuller, and A. Kettemann. Vacancy durations and entry wages: Evidence from linked vacancy–employer–employee data. *Review of Economic Studies*, 91(3):1808–1841, 2024.

M. Nagler, J. Rincke, and E. Winkler. How Much Do Workers Actually Value Working from Home? Technical report, 2022.

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

J. Sockin. Show me the amenity: Are higher-paying firms better all around? Technical report, 2022.

I. Sorkin. Ranking firms using revealed preference. *Quarterly Journal of Economics*, 133(3): 1331–1393, 2018.

C. Taber and R. Vejlin. Estimation of a roy/search/compensating differential model of the labor market. *Econometrica*, 88(3):1031–1069, 2020.

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

Appendix

A Additional Tables and Figures

Table A.1: Overview of Sample Selection: All Five-Year Periods.

	Job Ads (1)		Establishments (2)		Workers (3)		Worker-Years (4)	
2000-2004								
All Observations	339,791	(100%)	203,084	(100%)	2,468,397	(100%)	9,820,695	(100%)
Text Analysis	326,582	(96.1%)	79,744	(39.3%)	2,226,611	(90.2%)	8,482,931	(86.4%)
Model Estimation	221,966	(65.3%)	54,766	(27.0%)	1,927,755	(78.1%)	7,271,922	(74.0%)
Main Analysis	207,747	(61.1%)	54,488	(26.8%)	1,915,369	(77.6%)	7,197,974	(73.3%)
2005-2009								
All Observations	948,589	(100%)	215,133	(100%)	2,663,702	(100%)	10,706,267	(100%)
Text Analysis	910,636	(96.0%)	105,908	(49.2%)	2,483,551	(93.2%)	9,721,545	(90.8%)
Model Estimation	708,884	(74.7%)	61,476	(28.6%)	2,060,892	(77.4%)	7,864,068	(73.5%)
Main Analysis	628,035	(66.2%)	60,868	(28.3%)	2,039,348	(76.6%)	7,744,231	(72.3%)
2010-2014								
All Observations	734,996	(100%)	228,191	(100%)	2,854,890	(100%)	11,797,651	(100%)
Text Analysis	698,923	(95.1%)	100,721	(44.1%)	2,622,071	(91.8%)	10,564,100	(89.5%)
Model Estimation	572,180	(77.8%)	62,925	(27.6%)	2,157,573	(75.6%)	8,442,020	(71.6%)
Main Analysis	519,459	(70.7%)	62,122	(27.2%)	2,125,632	(74.5%)	8,264,304	(70.1%)
2015-2019								
All Observations	915,789	(100%)	292,327	(100%)	3,123,052	(100%)	13,278,337	(100%)
Text Analysis	860,467	(94.0%)	118,741	(40.6%)	2,882,624	(92.3%)	11,893,847	(89.6%)
Model Estimation	744,133	(81.3%)	78,136	(26.7%)	2,353,355	(75.4%)	9,436,842	(71.1%)
Main Analysis	623,341	(68.1%)	77,164	(26.4%)	2,324,318	(74.4%)	9,239,281	(69.6%)
2000-2019								
All Observations	2,939,165	(100%)	449,570	(100%)	3,867,230	(100%)	45,602,952	(100%)
Text Analysis	2,796,608	(95.1%)	195,788	(43.6%)	3,721,622	(96.2%)	40,662,424	(89.2%)
Model Estimation	2,247,163	(76.5%)	126,931	(28.2%)	3,335,908	(86.3%)	33,014,852	(72.4%)
Main Analysis	1,978,582	(67.3%)	125,504	(27.9%)	3,312,618	(85.7%)	32,445,790	(71.1%)

Note: This table documents the sample sizes for different sample restrictions imposed in the estimation for each of the four five-year periods and the 2000-2019 period. The first row ("All Observations") in each panel shows the total number of observations in our data on job ads (Column 1) and matched employer-employee data (Columns 2-4) for each five-year period, prior to the various sample selection steps. The remaining three rows in each panel summarize the sample restrictions imposed in the three separate estimation stages. "Text Analysis" refers to the sample of job ads with non-missing text content and at least one section written in Norwegian. "Model Estimation" restricts the sample to the strongly connected set of employers with workers between 20 and 60 years old (inclusive). "Main Analysis" imposes both of the previous sets of restrictions, and additionally removes recruitment and employment agencies. Job ads are only observed between 2002 and 2019 due to data availability, meaning that the number of observations on job ads in the top and bottom panel refers to the number of job ads posted in 2002-2004 and 2002-2019, respectively.

Table A.2: Overview of Sample Composition.

		Text Analysis	Model Estimation		Main Analysis
	All Observations	Observe Vacancy Text	Prime-age Workers	Strongly Connected Set	All Sample Restrictions
	(1)	(2)	(3)	(4)	(5)
Worker Characteristics:					
Women	43.7%	49.9%	43.7%	51.2%	51.2%
College	34.9%	37.9%	34.9%	40.2%	40.2%
Hourly Wage	$25.34	$25.17	$25.34	$25.45	$25.45
Number of Workers	3,123,052	2,882,624	2,671,777	2,353,355	2,324,318
Number of Worker-Years	13,278,337	11,893,847	11,253,721	9,436,842	9,239,281
Establishment Characteristics:					
Size	11.6	19.9	11.6	26.7	26.7
Public Sector	17.3%	25.8%	17.3%	30.7%	30.7%
Industry Composition:					
Agriculture, Forestry and Fishing	2.8%	1.3%	2.8%	1.0%	1.0%
Mining and Quarrying	0.4%	0.4%	0.4%	0.5%	0.5%
Manufacturing	5.8%	6.5%	5.8%	6.6%	6.6%
Energy Supply	0.4%	0.5%	0.4%	0.4%	0.4%
Water Supply	0.5%	0.6%	0.5%	0.6%	0.6%
Construction	10.3%	8.2%	10.3%	8.0%	8.0%
Wholesale and Retail Trade	23.3%	24.6%	23.3%	22.9%	22.9%
Transportation and Storage	5.5%	4.5%	5.5%	4.6%	4.6%
Accommodation and Food Service	4.2%	4.9%	4.2%	4.7%	4.7%
Information and Communication	3.1%	2.8%	3.1%	2.9%	2.9%
Financial and Insurance	1.4%	1.5%	1.4%	1.5%	1.5%
Real Estate	3.4%	1.5%	3.4%	1.1%	1.1%
Professional, Scientific and Technical	8.2%	5.8%	8.2%	5.1%	5.1%
Administrative and Support Service	4.0%	4.2%	4.0%	4.4%	4.4%
Public Administration and Defence	2.9%	4.2%	2.9%	4.3%	4.3%
Education	3.8%	5.7%	3.8%	7.1%	7.1%
Health and Social Work	12.9%	17.4%	12.9%	20.0%	20.0%
Arts, Entertainment and Recreation	2.0%	1.8%	2.0%	1.7%	1.7%
Other Service Activities	4.6%	3.6%	4.6%	2.5%	2.5%
Number of Establishments	292,327	118,741	268,565	78,136	77,164

Note: This table documents how the sample restrictions imposed in the estimation affects the sample composition between 2000 and 2019. The table shows averages over establishment-years, meaning that employers that are present in all years are implicitly weighted more than employer that are present in half of the years. "Worker Characteristics" refers to average worker characteristics across establishment-years, weighting workers in small employers more than large employers. For instance, the first row is constructed by first calculating the average share of women within employer-years, and then averaging over employer-years. The first column refers to the initial sample prior to the various sample selection steps, while the remaining columns refer to the three estimation stages, as in Table 1. Industry composition is based on the Classification of Economic Activities in the European Community (NACE).

Table A.3: Validation of Detected Pay and Non-Pay Job Attributes.

	All Ads	Job Ads in the Validation Sample				
	Text Analysis	Text Analysis	Manual Recognition	Success Rate	Precision Rate	Sensitivity Rate
	(1)	(2)	(3)	(4)	(5)	(6)
Compensation Scheme						
— Competitive Pay	15.8	17.5	19.2	94.2	88.6	80.5
— Collective Agreement Pay	24.3	30.8	33.8	87.0	83.7	76.3
— Incentive Pay Scheme	2.7	2.0	1.5	98.5	50.0	66.7
— Hiring Bonus	0.2	0.0	0.0	100.0	–	–
— Good Overtime Pay	0.4	0.2	0.8	99.5	100.0	33.3
— Any Other Mention of Pay	14.5	14.5	8.8	86.2	32.8	54.3
Financial Attributes						
— Pension Scheme	30.4	34.8	32.8	98.0	94.2	100.0
— Insurance Scheme	19.1	23.8	21.5	96.2	87.4	96.5
— Mortgage Possibility	1.4	2.0	3.8	98.2	100.0	53.3
Career Opportunities						
— Good Career Paths	12.6	16.8	6.2	85.0	23.9	64.0
— On-the-Job Training	41.6	46.8	39.2	78.0	68.4	81.5
Hours of Work						
— Full-time Contract	59.3	55.8	53.2	95.0	93.3	97.7
— Part-time Contract	30.9	34.0	28.2	91.2	78.7	94.7
— Full-time/Part-time Choice	4.1	3.0	3.0	94.0	0.0	0.0
Convenient Hours						
— Possibility to Work Flexible Hours	6.5	6.8	7.5	98.2	92.6	83.3
— Regular Daytime Work Schedule	14.0	8.8	8.8	97.0	82.9	82.9
— Exempt from Work Hour Regulations	0.5	1.2	0.2	98.5	0.0	0.0
Inconvenient Hours						
— Shift Work	15.4	15.0	11.8	95.8	75.0	95.7
— Weekend/Evening/Night Work	18.4	21.8	17.0	94.8	77.0	98.5
— On-call Employment	3.4	3.8	2.2	98.0	53.3	88.9
— Overtime Work Required	0.2	0.0	0.2	99.8	–	0.0
Contract Duration						
— Permanent Job	40.6	38.8	40.8	93.0	93.5	89.0
— Temporary Job	29.6	32.2	29.2	91.0	81.4	89.7
— Fixed-term Contract	0.4	0.8	0.8	100.0	100.0	100.0
Workplace Attributes						
— Social Environment	41.4	44.0	40.8	80.8	74.4	80.4
— Good Colleagues	21.7	26.0	28.2	77.2	60.6	55.8
— Possibility for Remote Work	0.4	0.5	0.5	100.0	100.0	100.0
— Shared Office Space	0.2	0.0	0.0	100.0	–	–
— Inclusive Work-life Scheme	7.1	7.2	6.0	98.2	79.3	95.8
Task-Related Attributes						
— Interesting Tasks	34.5	38.2	22.2	82.5	56.2	96.6
— Challenging Tasks	28.6	32.0	19.0	85.0	56.2	94.7
— Variation in Tasks	10.9	10.2	15.8	87.5	65.9	42.9
— Responsibility in Job	1.6	1.8	10.2	90.5	71.4	12.2
— Independence in Performing Tasks	2.5	2.8	5.0	95.8	63.6	35.0
— Involves Leadership Responsibility	11.8	13.5	9.0	87.0	35.2	52.8
— Work Involves Travelling	8.2	10.2	3.5	93.2	34.1	100.0
Minor Perks						
— Beautiful Location	0.8	2.2	4.8	96.5	77.8	36.8
— Central Location	33.0	37.5	5.2	65.8	11.3	81.0
— Company Gym or Sports Team	3.7	5.0	6.5	95.5	70.0	53.8
— Parking Space On Premises	0.5	0.5	0.2	99.8	50.0	100.0
— Company Vehicle	1.2	1.5	1.0	99.5	66.7	100.0
— Any Welfare Scheme	4.2	4.8	2.5	97.2	47.4	90.0
— Company Cabin	1.9	2.0	3.2	98.8	100.0	61.5
— Occupational Health Service	1.4	1.2	1.2	100.0	100.0	100.0
— Company Canteen	0.3	0.8	0.8	99.0	33.3	33.3
— Flexible/Extended Holidays	0.5	0.0	0.0	100.0	–	–
Number of Job Ads	860,467	400	400	400	–	–

Note: See details in the notes to Table 3.

Table A.4: Association Between Posted Employer Values and Vacancy Duration.

	Outcome: Log Vacancy Duration
Panel A: Overall Employer Value	
Posted Overall Value ($\tilde{V}_j$)	-0.139***
	(0.023)
Panel B: Pay, Non-Pay, and Job Security Values	
Posted Pay Premium ($\tilde{\psi}_j$)	-0.689***
	(0.205)
Posted Non-Pay Value ($\tilde{a}_j$)	-1.032***
	(0.170)
Posted Job Security Value ($\tilde{s}_j$)	-0.059***
	(0.015)
Controls:	
Occupation Fixed Effects	✓
Industry Fixed Effects	✓
Location Fixed Effects	✓
Observations:	
Employer Clusters	1,559
Establishments	77,164
Ads	623,341
Worker-Years	9,239,060

Robust standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: This table documents associations between posted employer values and vacancy duration. Vacancy duration is the number of days between the dates when the ad is posted and unlisted, censored between 1 and 90 days to avoid outliers, averaged at the cluster level before taking logs. We compute the "posted" employer value ($\tilde{V}_j, \tilde{\psi}_j, \tilde{a}_j,\, \tilde{s}_j$) associated with each ad as the prediction from regressions of employer values ($V_j, \psi_j, a_j,\, s_j$) on indicators for posted ad attributes. The "posted" employer pay premium ($\tilde{\psi}_j$) and non-pay (amenity) value ($\tilde{a}_j$) are measured in logs, so the estimated coefficient for each of these can be interpreted as an elasticity, while the overall employer value ($\tilde{V}_j$) and job security ($\tilde{s}_j$) are scaled by their standard deviation, so the estimated coefficient for each of these represents a standard deviation increase in employer value. The specification in Panel A focuses only on the overall employer value ($\tilde{V}_j$), while Panel B includes its components ($\tilde{\psi}_j, \tilde{a}_j,\, \tilde{s}_j$) jointly in the same estimation. Estimations are done at the employer-cluster level for the five-period from 2015 to 2019, and are weighted by the number of worker-years in each cluster. We control for occupation and industry using two-digit occupation and two-digit industry fixed effects and location based on deciles of the number of workers in the municipality of the firm (Oslo as an own group).

Figure A.1: Sample Vacancy Text: Teacher Substitute.

Substitute in a professional position:

One year of temporary work in a 100% position from 01.08.2016 - 31.07.2017.

The position entails travel activities in the region and the applicant must have their own car.

We are looking for an employee with:

- Special needs teacher with competence at the Master's or Bachelor's level.
- Professional experience from kindergarten/school and relevant skills.
- Good collaboration skills, flexibility, and personal suitability.

The following applies to the position:

- Challenging tasks.
- Good working environment in an interdisciplinary team.
- Appointment and working conditions according to current agreements and regulations.
- Wages in accordance with the agreement.
- Compulsory membership in KLP (public sector occupational pension scheme).

[...]

Note: Translation from Norwegian by the authors.

Figure A.2: Sample Vacancy Text: Civil Engineer.

[...]

We need to increase capacity within water- and environmental engineering, and are therefore looking for a civil engineer/engineer with experience within water supply, sewage and urban drainage systems.

The position involves exciting and challenging tasks in an interdisciplinary environment.

The work will include planning, engineering, and follow-up of private and municipal urban drainage-, water-, and sewage systems.

Desired competence:

- Engineer / civil engineer
- Relevant experience
- Good ability to work with colleagues, partners, and clients
- Good knowledge of written and spoken Norwegian
- New graduates with good results may also be relevant

Preferred abilities:

- Ability to work independently as well as in a team
- Solution-oriented and accurate
- Resourceful
- Positive and persistent

We offer:

- An independent position in a solid company in growth and development
- Professionally inspiring with talented colleagues
- A pleasant and orderly workplace with bright and open premises in central Oslo
- Competitive conditions

[...]

Note: Translation from Norwegian by the authors.

Figure A.3: The Prevalence of Job Attributes Advertised in Vacancy Texts: 2002-2019.

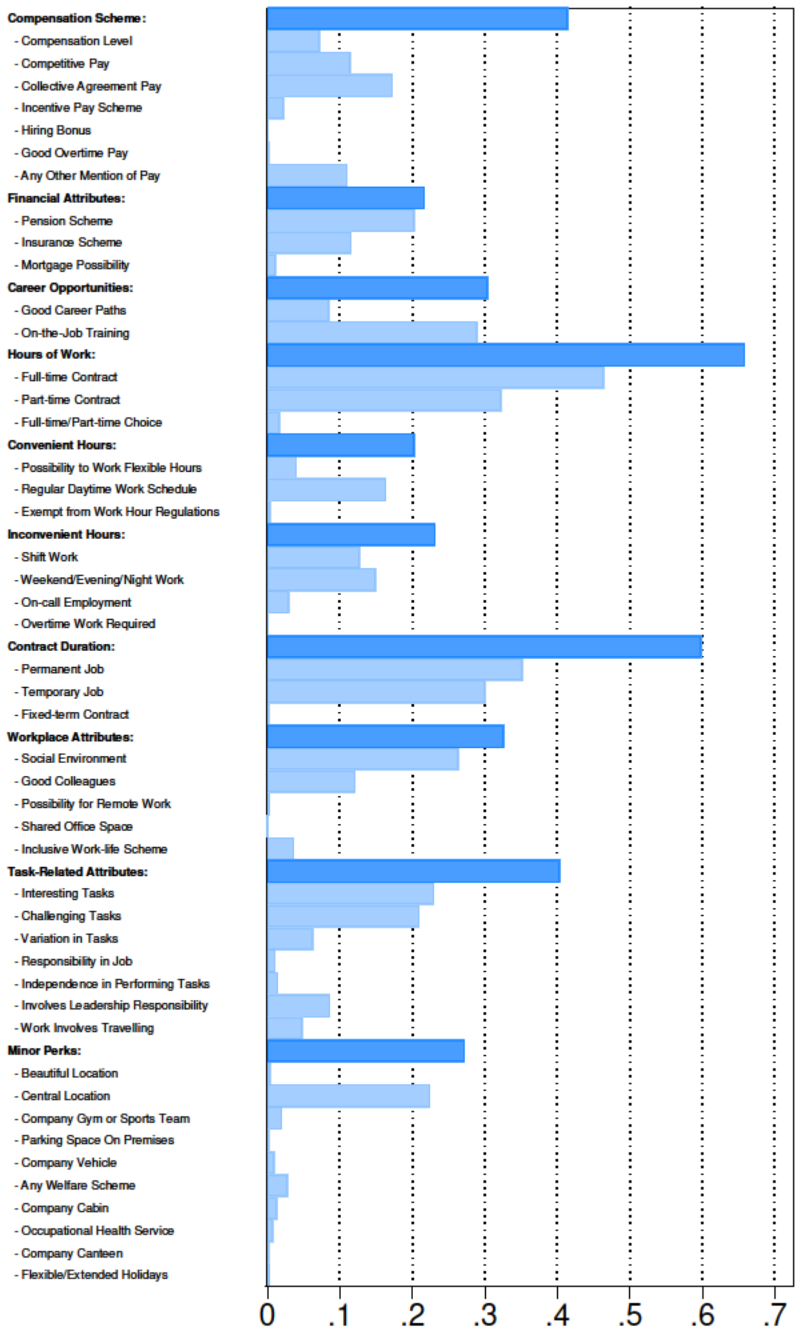

Note: This figure documents the prevalence of job attributes detected in job ads posted in Norway between 2002 and 2019 [N=2,813,645]. The light blue bars show the share of ads detected having each of the distinct job attributes. The dark blue bars show the share of ads detected with at least one attribute within ten broad categories.

Figure A.4: Explained Variation in Publicly Advertised Job Attributes: 2002-2019.

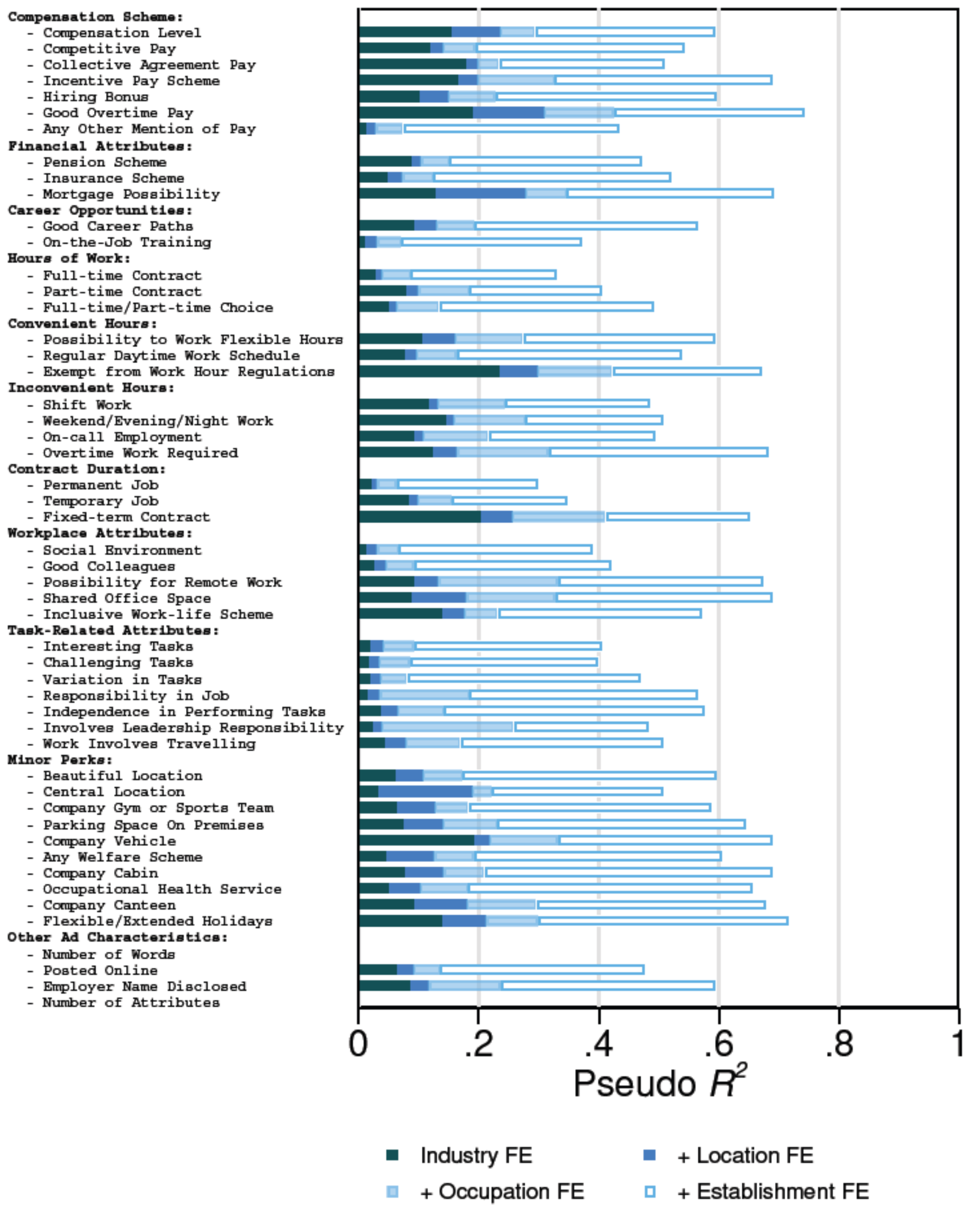

Note: This figure shows pseudo R^2 [McFadden, 1974] from separate logistic regressions of binary job attributes detected in job ads posted in Norway between 2002 and 2019 [N=2,533,559], on fixed effects denoting unique combinations of 2-digit industries (91 groups), location indicators (10 groups), 2-digit occupations (43 groups), and 126,890 establishments. We start by including industry fixed effects, continue by including industry$\times$location fixed effects, and so on. The regression excludes 263,049 ads with a unique combination of establishment, occupation, location, and industry indicators, as these observations are precisely predicted by the last set of regressions by construction. The last set of regressions controls for 297,142 unique combinations. Location indicators group 422 municipalities into ten groups based on the number of workers such that municipalities assigned to the same group have a similar number of workers, with a specific indicator for job postings in Oslo.

Figure A.5: Publicly Advertised Job Attributes By Employer Overall Values.

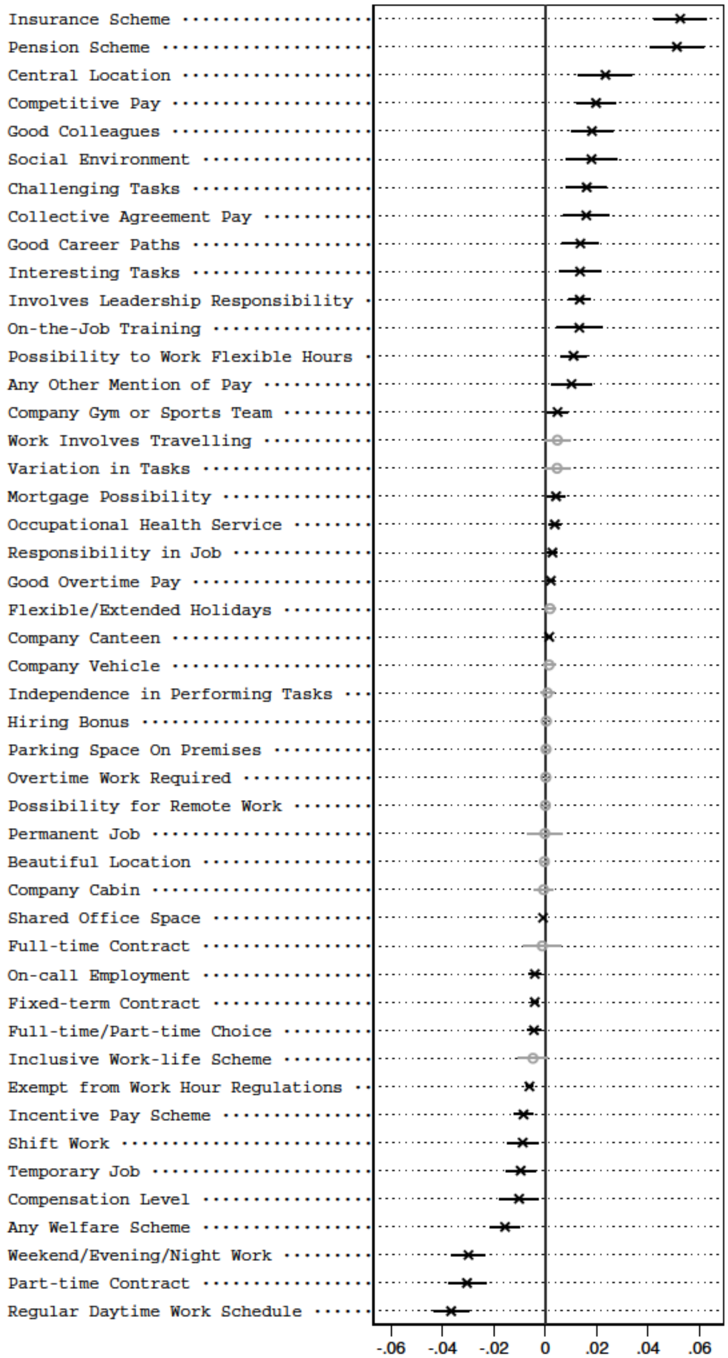

Note: This figure shows parameter estimates and 90% confidence intervals from separate regressions of the share of each ad attribute on overall employer values (V_j), capturing the change in the fraction of ads with each job attribute associated with a one standard deviation increase in employer value. Estimation are done at the employer-cluster level for the five-year period from 2015 to 2019, and are weighted by the number of worker -years in each cluster. We control for occupation and industry using two-digit occupation and two-digit industry fixed effects and location based on deciles of the number of workers in the municipality of the firm.

Figure A.6: Publicly Advertised Job Attributes By Employer Values: Smaller Clusters.

(a) Pay (ψ_j) (b) Non-Pay (a_j)

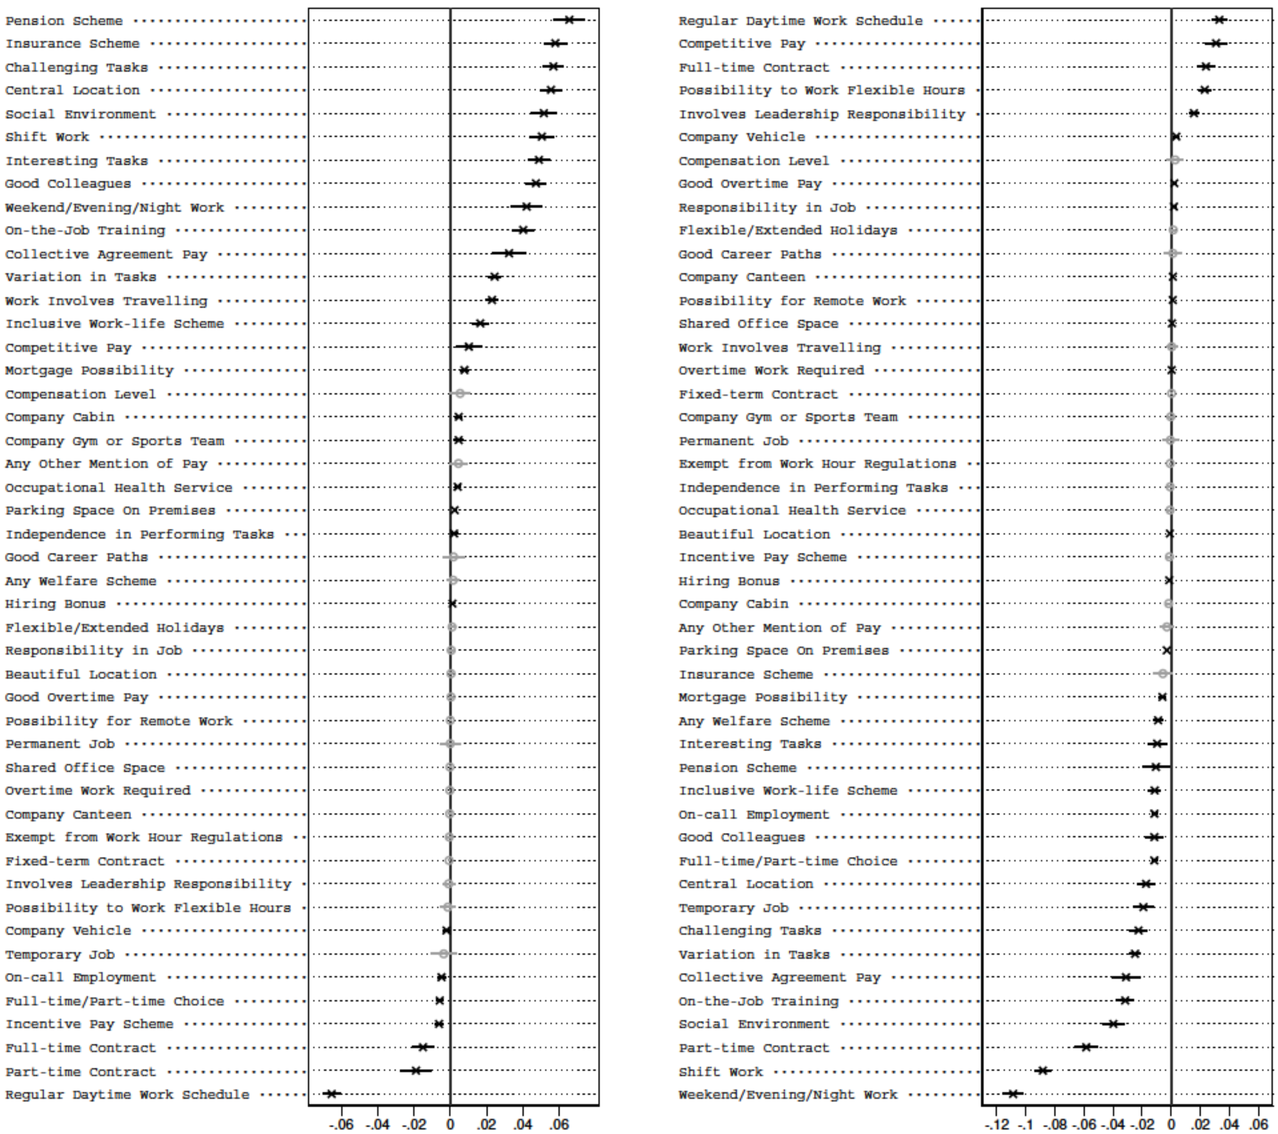

Note: This figure shows parameter estimates and 90% confidence intervals from separate regressions of the share of each ad attribute on employer pay (ψ_j) and non-pay (a_j) values, capturing the change in the fraction of ads with each job attribute associated with a standard deviation increase in employer value. Estimations are done at the employer-cluster level for the five-year period from 2015 to 2019, and are weighted by the number of worker-years. We control for occupation and industry using two-digit occupation and two-digit industry fixed effects and location based on deciles of the number of workers in the municipality of the firm (Oslo as an own group). On average, each cluster has 25 establishments, i.e., $G = \text{round}(J/25)$, while Figure 5 shows the corresponding results in our baseline with $G = \text{round}(J/50)$.

Figure A.7: Publicly Advertised Job Attributes By Employer Values: Larger Clusters.

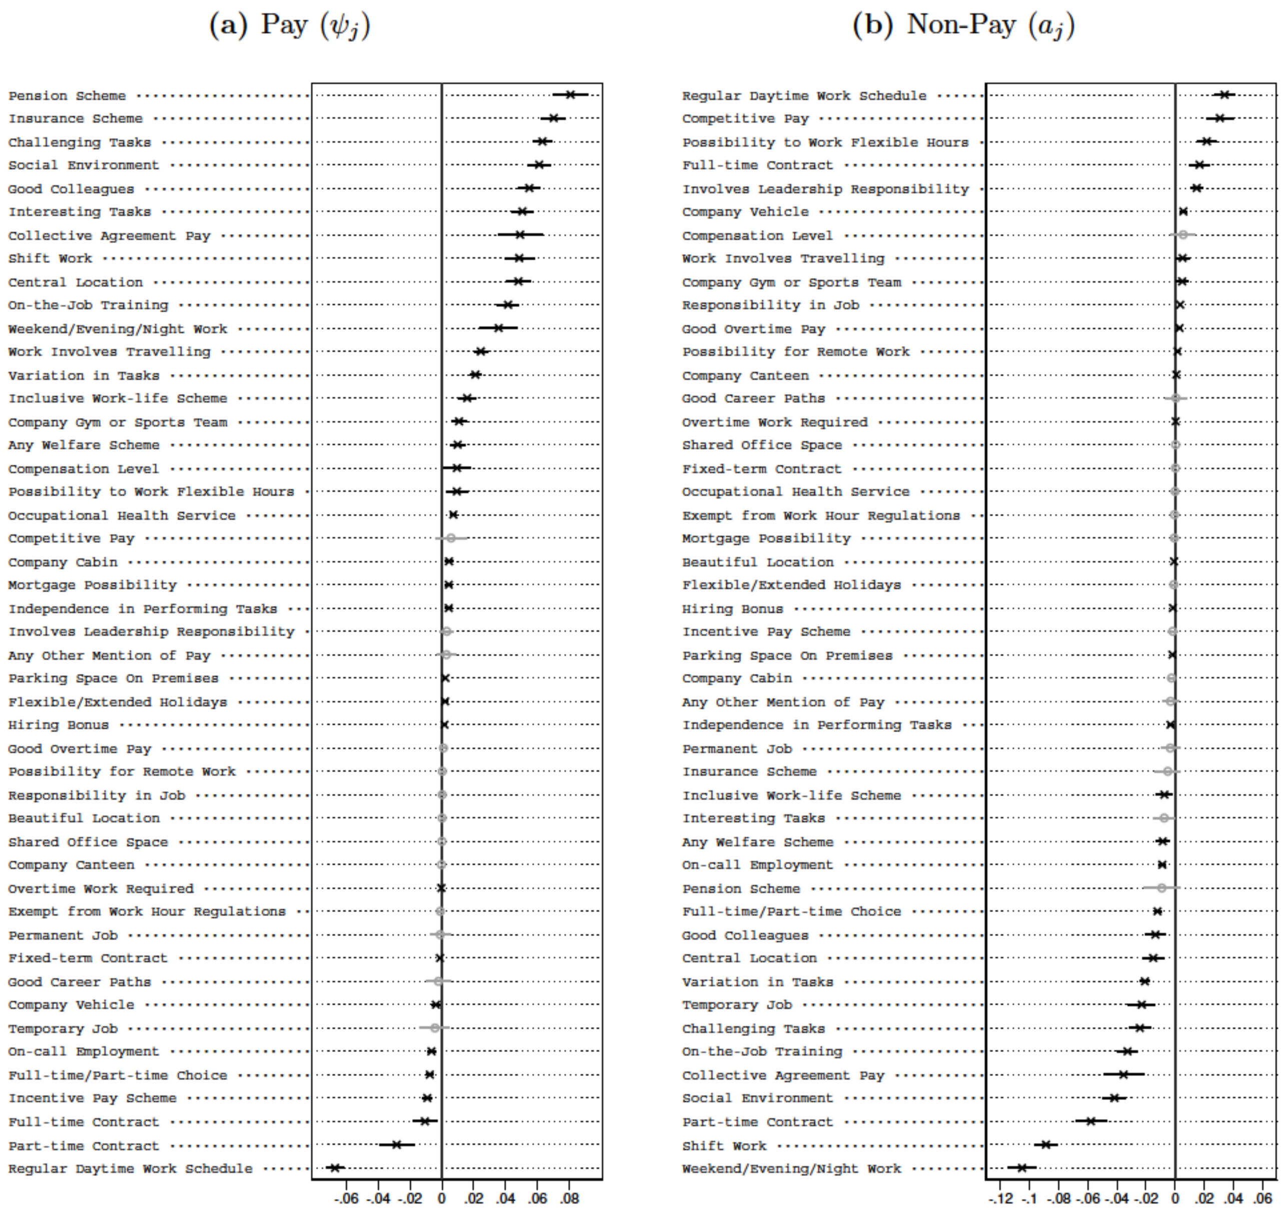

Note: This figure shows parameter estimates and 90% confidence intervals from separate regressions of the share of each ad attribute on employer pay (ψ_j) and non-pay (a_j) values, capturing the change in the fraction of ads with each job attribute associated with a standard deviation increase in employer value. Estimations are done at the employer-cluster level for the five-year period from 2015 to 2019, and are weighted by the number of worker-years. We control for occupation and industry using two-digit occupation and two-digit industry fixed effects and location based on deciles of the number of workers in the municipality of the firm (Oslo as an own group). On average, each cluster has 100 establishments, i.e., G =round($J/100$), while Figure 5 shows the corresponding results in our baseline with G =round($J/50$).

Figure A.8: Decomposing the Predictive Power of Publicly Advertised Job Attributes.

(a) Overall Value (V_j)

(b) Pay (ψ_j)

(c) Non-Pay (a_j)

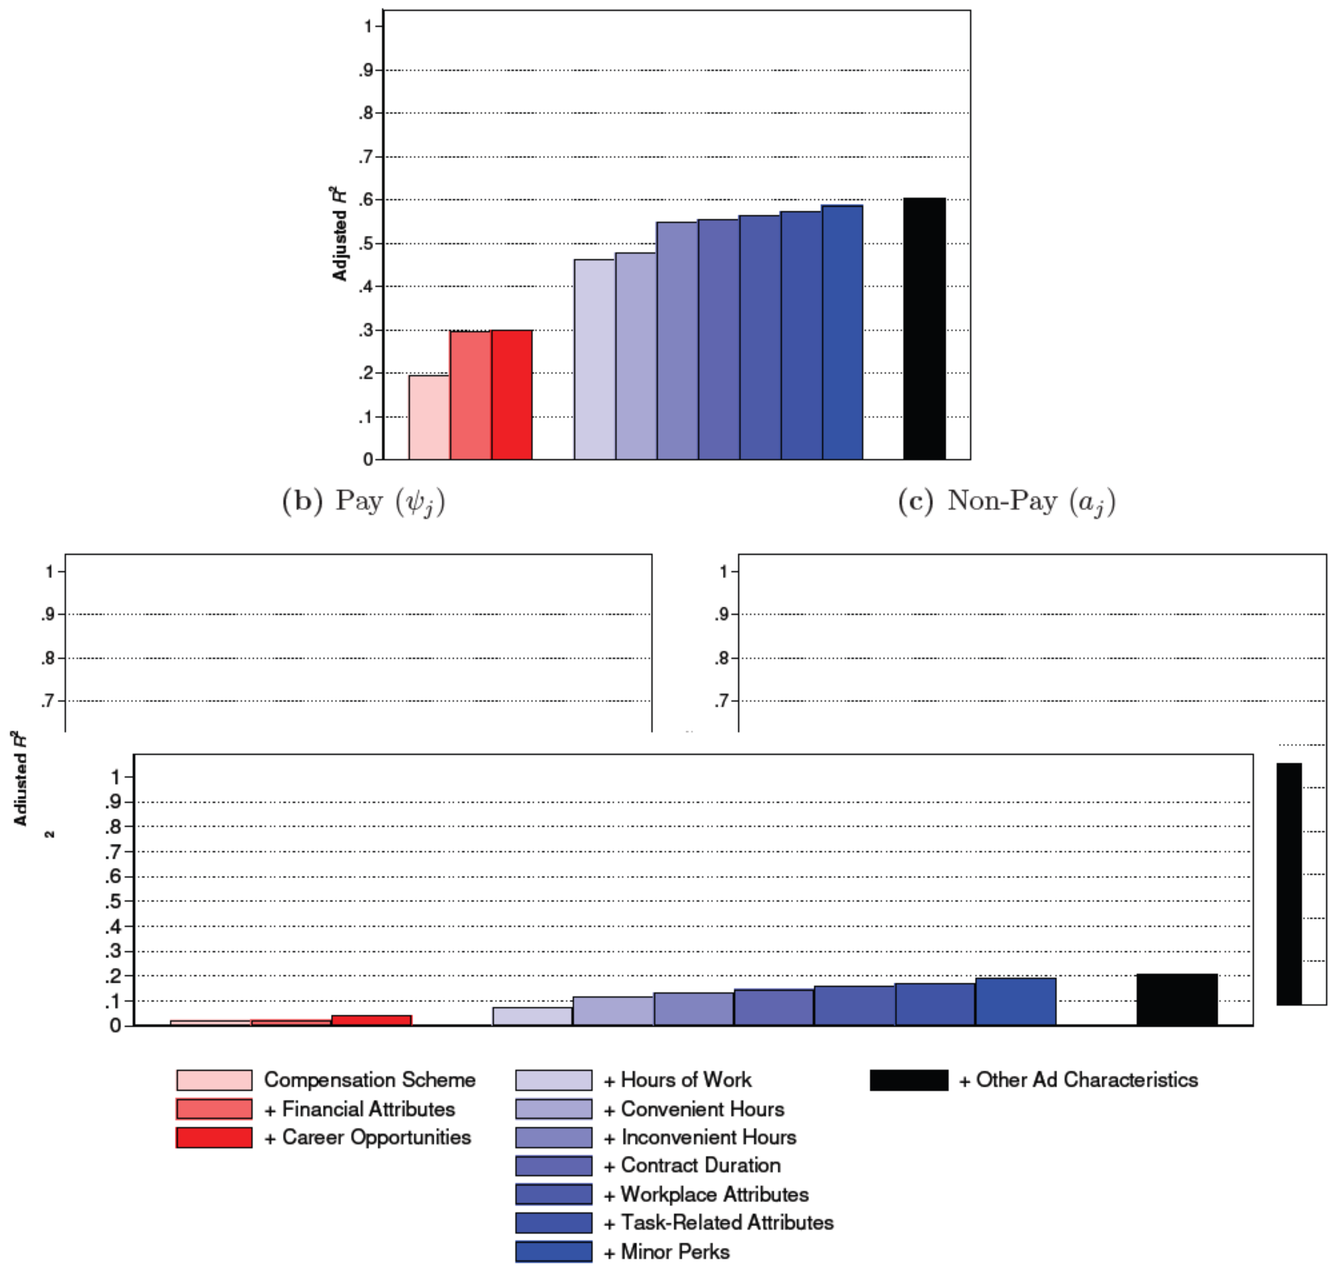

Note: This figure documents adjusted R^2s from regressions of estimated employer values on pay and non-pay job attributes and other ad characteristics detected in the text analysis as in Figure 6, decomposing the overall R^2s into ten broad categories of pay and non-pay attributes. Estimations are done at the employer-cluster level for the five-year period from 2015 to 2019, and are weighted by the number of worker-years in each cluster.

Figure A.9: Predictive Power of Publicly Advertised Job Attributes: By Industry.

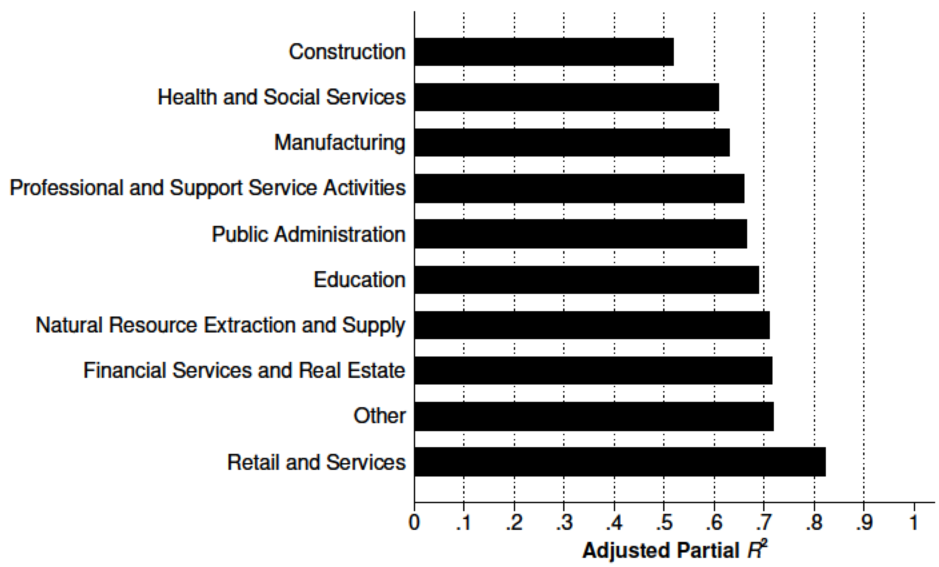

Note: This figure documents adjusted R^2s from regressions of estimated overall employer values on pay and non-pay job attributes and other ad characteristics detected in the text analysis as in Figure 6. Estimations are done separately for the ten main industry groups at the employer-cluster level for the five-year period from 2015 to 2019, and are weighted by the number of worker-years in each cluster. We assign an industry to each cluster g based on the modal industry among the establishments in the cluster.

Figure A.10: Predictive Power of Publicly Advertised Job Attributes: By Cluster Size.

(a) Baseline Model: $G = \text{round}(J/50)$

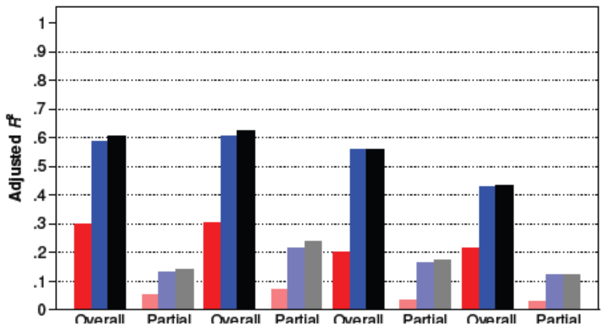

(b) Small (J/100)

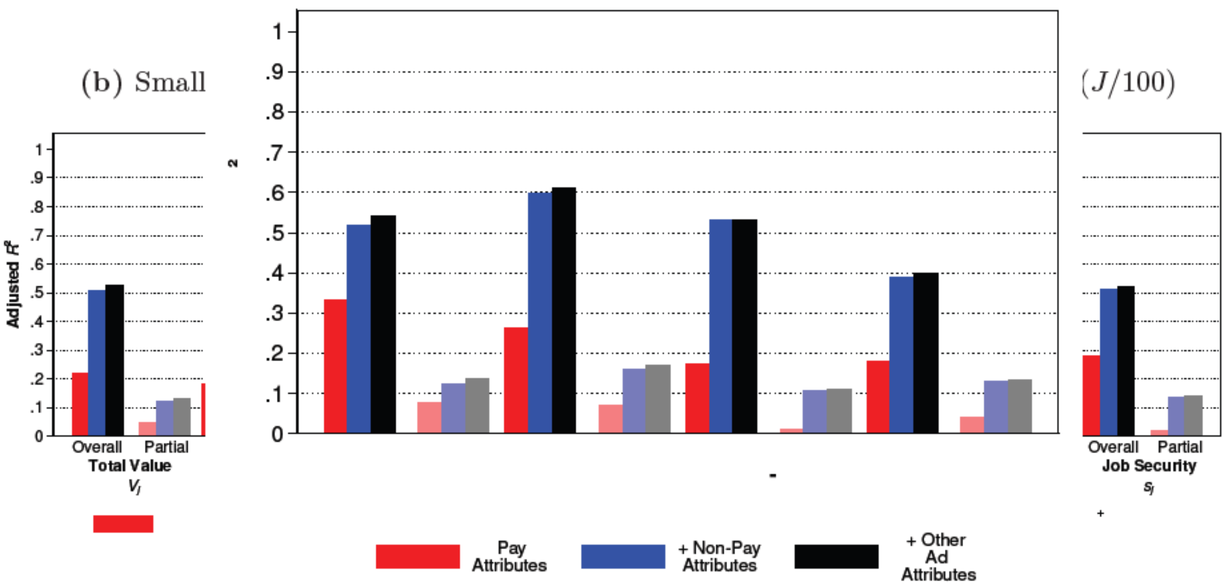

Note: This figure documents adjusted R^2 from regressions of cluster-size specific estimated model parameters on ad attributes and characteristics in the text analysis as in Figure 6. Estimations are done at the employer-cluster level for each of the four five-year period from 2000 to 2019, and are weighted by the total number of worker-years in each cluster and period. Overall R^2s are from regressions that include ad attributes and other ad characteristics, while partial R^2s are from regressions that also control for the composition of industry, occupation and location in each cluster. Industry and occupation controls are defined at the two-digit level. Location indicators split local municipalities into deciles, with municipalities in the same group having similar number of workers (Oslo as an own group).

Figure A.11: Predictive Power of Publicly Advertised Job Attributes Over Time.

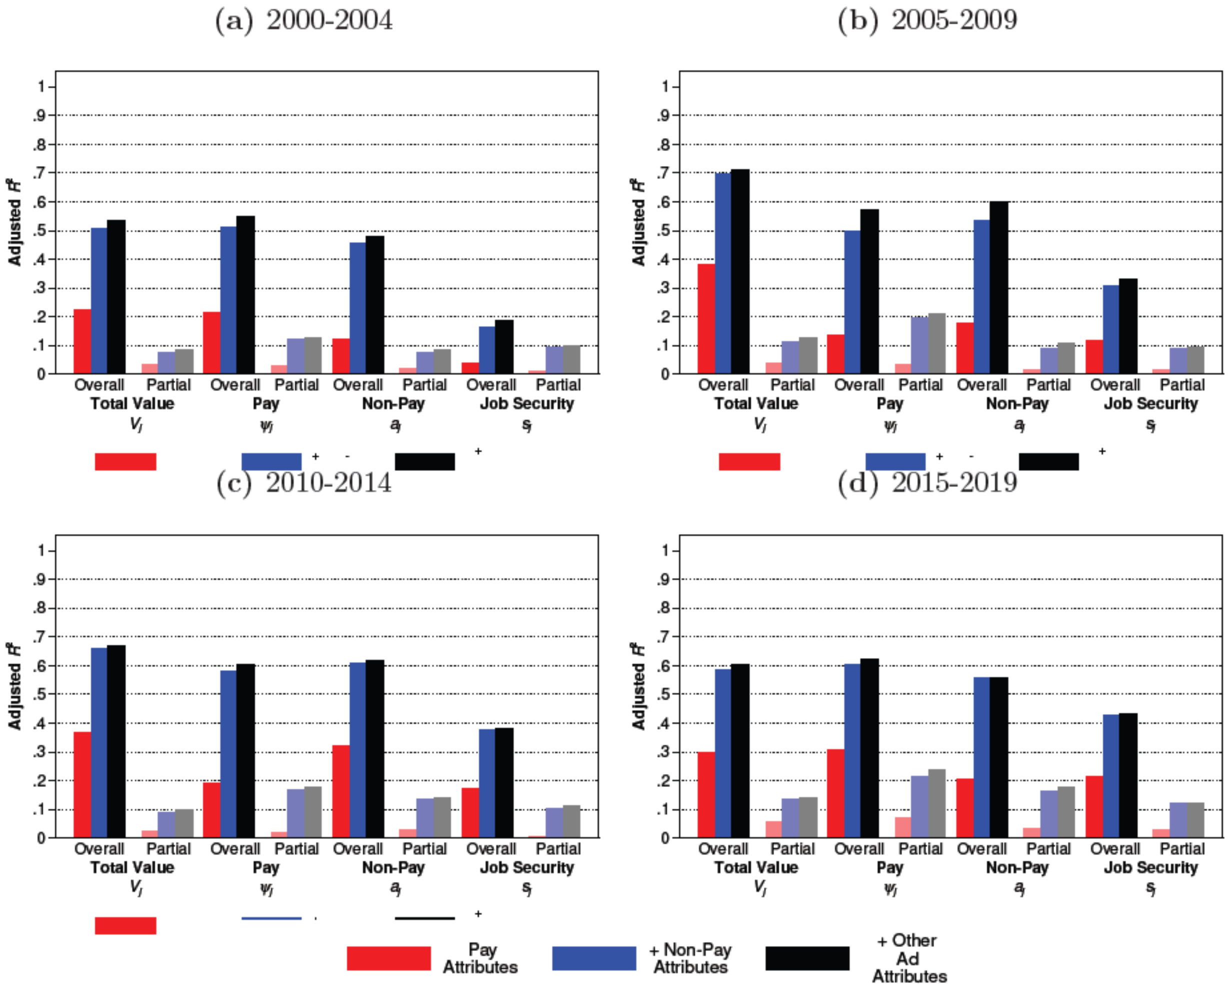

Note: This figure documents adjusted R^2 from regressions of estimated population group-specific model parameters on ad attributes and characteristics in the text analysis as in Figure 6. Estimations are done at the employer-cluster level for each of the four five-year period from 2000 to 2019, and are weighted by the total number of worker-years in each cluster and period. Overall R^2s are from regressions that include ad attributes and other ad characteristics, while partial R^2s are from regressions that also control for the composition of industry, occupation and location in each cluster. Industry and occupation controls are defined at the two digit level. Location indicators split local municipalities into deciles, with municipalities in the same group having similar number of workers (Oslo as an own group). The sample corresponds to the last rows in Panels A-D in Appendix Table A.1

Figure A.12: Heterogeneity in the Predictive Power of Publicly Advertised Job Attributes.

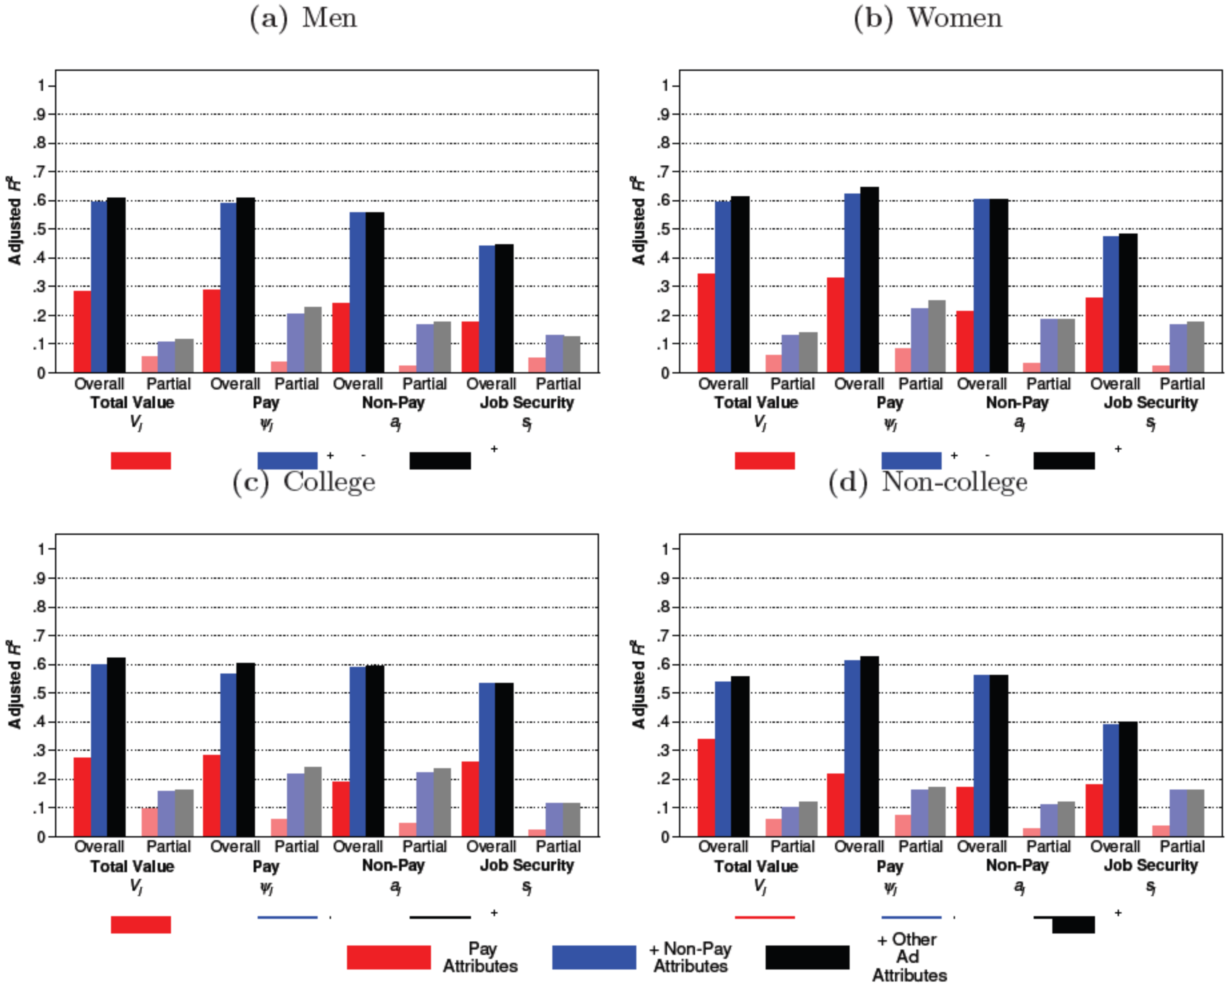

Note: This figure documents adjusted R^2 from regressions of estimated population group-specific model parameters on ad attributes and characteristics in the text analysis as in Figure 7, with fixed weights across sample splits. Estimations are done at the employer-cluster level for the five-year period from 2015 to 2019, separately by gender (Panels A-B) and education (Panels C-D), and are weighted by the total number of worker-years in each cluster, as opposed to by cluster *and* population group as in Figure 7. Overall R^2s are from regressions that include ad attributes and other ad characteristics, while partial R^2s are from regressions that also control for the composition of industry, occupation and location in each cluster. Industry and occupation controls are defined at the two digit level. Location indicators split local municipalities into deciles, with municipalities in the same group having similar number of workers (Oslo as an own group).

B Additional Details on the Text Analysis

Section B.1 documents how we clean the text of job ads before detecting job attributes. Section B.2 describes one of the processes we use to generate the complete lists of attributes we search for in vacancy texts. Section B.3 provides descriptions of detailed attributes and displays some of the phrases used to detect these attributes.

B.1 Text Cleaning of Job Ads

Before attempting to extract attributes from job ads we start with some simple cleaning of the texts contained in vacancies. We start by searching for HTML tags in the vacancies. We do this to avoid considering words that are part of HTML code to indicate job attributes, and to reduce the number of words considered by the model we use to generate phrases for attribute detection. HTML tags always start with “<” and end with “>”. We keep a record of which ads are detected with this pattern and remove the corresponding tags.

Second, we replace numbers in the texts with flags indicating their range and whether the number was followed by a percentage. We do this to be able to group similar phrases containing different numbers. For example, “50 %” and “40 %” are commonly used to indicate part time jobs. By grouping every phrase of the form ‘{a number below 100} %’ we avoid having to search for every number less than a hundred followed by a percentage sign. We start by flagging numbers less than 100 (equal to 100) followed by a percentage sign or the word “percentage”. We then create separate flags for each remaining integer less than or equal to four, remaining numbers between 0 and 100, the number 100, and numbers greater than 100.

Before storing the vacancies, we split the ads into distinct sentences, and sentences into distinct words. This is primarily done because the language model we use to expand the dictionary of target phrases expects input in this form. Splitting sentences also has the advantage of separating the words at the end of one sentence from the words in the beginning of the next. We use the Python package `nltk` to do these splits. The package uses a model that is trained on a corpus in Norwegian, and it avoids splitting at periods that are not at the end of sentences, e.g., as part of a common abbreviation.

B.2 Extracting Common Attributes from "We Offer" Lists

Our choice of which pay and non-pay attributes to include in the text analysis is based partly on commonly advertised attributes extracted from "we offer" sections of job ads. This section documents this process, and shows that attributes related to social environment, on-the-job training, tasks, pension schemes, and insurance schemes are commonly advertised in "we offer" sections of publicly posted vacancies. Our approach has four broad steps: (i) we extract lists from the corpus of job ads; (ii) we sort the lists into four distinct topics using a common model from natural language processing; (iii) we identify the topic that contains "we offer" sections by inspecting words and lists that, according to the model, are closely associated with the "we offer" sections; and (iv) we extract common phrases from the lists that are sorted into the "we offer" category. We assign some of these phrases to attribute categories, and add these categories to the list of attributes to search for in vacancies.

Extracting lists from vacancies To identify "we offer" sections in vacancies, we start by identifying and extracting a large collection of lists from the full corpus of job ads. Similar information in job ads is commonly grouped in lists, such as skill requirements, tasks, and the pay and non-pay content (i.e., "we offer" section). We extract lists by flagging consecutive instances of the html tag "<li>", which is used to indicate list elements in web pages, and plain text lines that start with a centered dot (·), hyphen (-) or a star (*). We combine consecutive sentences with this flag with the previous sentence (the list header) into separate observations. This process allows us to extract sections of texts that are likely to center on a single topic like skill requirements, tasks and pay and non-pay attributes. Although it possibly omits some lists, we are able to extract close to 1.5 million lists from around 50 percent of the vacancies.

Sorting lists into different topics To distinguish "we offer" lists from lists covering other topics, we apply Latent Dirichlet Allocation (LDA) to the collection of lists [Blei et al., 2003]. Latent Dirichlet Allocation is commonly used in natural language processing to partition a collection of texts into different topics. LDA partitions the texts into K categories such that words that commonly occur in topic k are unlikely to occur in other topics. The model performs well when each topic is associated with words that are not frequently used in other topics. However, the model ignores how words are ordered, making the individual words that are used to assign texts to topics less informative. For example, the model splits the phrase "good colleagues" into two observations, "good" and "colleagues." While "good" is commonly used to indicate workplace amenities or pay, it does not contain information about the specific attribute without its context. Additionally, some lists combine "we offer" sections with other topics, such as skill requirements. We therefore expect to lose some of the "we offer" content that is assigned to other topics by the model.

Latent Dirichlet Allocation partitions texts into topics by modelling each individual word as being drawn from a mixture model: Each word is drawn from one of K word-distributions, each corresponding to a different topic. The texts (lists) are associated with separate mixtures over the K topics, meaning that the probability that a word is drawn from topic k varies across texts (lists). That is the probability that the word i in text j (w_{ij}) is the word "offer" is

$$\Pr(w_{ij} = \text{“offer”}) = \sum_{k=1}^{K} \Pr(w_{ij} = \text{“offer”} | \text{topic}_{ij} = k) \Pr(\text{topic}_{ij} = k)$$

We set the number of topics to 4 and estimate the model with maximum likelihood using the implementation provided in the Python package `gensim`.

Labelling topics After estimating the model, we label the topics into four identified topics by considering the word distributions and the text (list) specific topic mixtures. Table B.1 documents the ten words with the highest association with each topic, i.e., the highest probability of being drawn in each topic. Words in topics 1-3 are associated with 1) pay and non-pay job attributes (amenities), 2) skill requirements, and 3) task descriptions, while words in topic 4 contains a mixture of requirements (e.g. education and police certificate), additional information (e.g. contact information), and ambiguous words (numbers and "work"). We label the topics 1-3 according to these observations, and topic 4 as a "Residual Category".

The "we offer" topic contains words commonly used to indicate a good social environment, information about pay, and words used to describe task related attributes that might be attractive to workers (e.g., challenging tasks). We note that some of these words are ambiguous outside of their original context, such as "development", which can be used to indicate personal development or tasks that includes development of projects.

We validate these labels by inspecting the lists with the highest association with each topic in Figure B.1-B.4. These lists serve as stylized examples of each topic and omit only considering one word at the time. These lists largely confirms the labels associated with each topic. The list most strongly associated with the "We Offer"-topic starts with the header "we offer:" and contains information about various non-pay amenities. The list most strongly associated with the "Residual Category"-topic contains various information about the job, such as the hiring process and and information about pay. Overall, these lists suggests that lists in the "We Offer" category contain dense information about amenities, although we are likely to lose some information about pay and non-pay in lists sorted into other categories (e.g., information about pay in lists assigned to the "Residual Category").

Extracting common phrases from "We Offer" sections To identify specific job attributes commonly advertised in vacancies, we now focus on common words and phrases in lists associated with the "we offer" topic identified by the model. We label each list with the topic that has the most mass in the list's associated topic distribution, and extract about 260,000 lists within the "we offer" label (i.e., 17 percent of all lists).

Next, we extract and inspect the 200 most common unigrams, bigrams and trigrams, excluding phrases that includes stopwords (e.g., "and"). We add phrases that indicate pay or non-pay job attributes (amenities) to our dictionary of phrases, and create new attribute categories whenever a phrase indicates an attribute not currently in our list.

Table B.2 documents the 10 most common unigrams, bigrams, and trigrams in "we offer" sections, excluding phrases without a clear reference to a particular job attribute. The table documents that these sections commonly refer to the social environment and pay-related attributes (pay, pension and insurance scheme, and on-the-job training). This procedure does not uncover many attributes related to hours of work, which we suspect is referred to in other sections of ads (e.g., practical information in the "Residual Category").

Table B.1: Words Used to Separate Lists Into Topics

Topic 1: We-Offer		Topic 2: Skill Requirements		Topic 3: Tasks		Topic 4: Residual Category	
Word (1)	Probability (2)	Word (3)	Probability (4)	Word (5)	Probability (6)	Word (7)	Probability (8)
"offer"	0.037	"experience"	0.038	"tasks"	0.025	Number > 99	0.044
"good (plural)"	0.032	"good (plural)"	0.027	"responsibility"	0.011	Number in [5-99]	0.024
"workplace environment"	0.029	"good"	0.024	"to follow up"	0.010	"have to"	0.016
"good (alternative form)"	0.026	"qualifications"	0.021	"contribute"	0.009	"Field of education"	0.009
"pay"	0.019	"ability"	0.016	"cooperation"	0.009	"police certificate"	0.008
"professional"	0.017	"Norwegian"	0.014	"other"	0.008	"contact information"	0.007
"exciting"	0.015	"relevant"	0.014	"tasks"	0.007	"level of education"	0.007
"tasks"	0.014	"written"	0.014	"participate"	0.007	"The position"	0.007
"developement"	0.014	"oral"	0.013	"customers"	0.006	"work"	0.006
"challenging"	0.010	"traits"	0.013	"developement"	0.005	"two"	0.006

Note: This table documents the words that are most strongly associated with the topics identified by applying Latent Dirichlet Allocation to our collection of lists. Columns (2), (4), (6), and (8) contains estimated probabilities of drawing the corresponding word given that the word is drawn from the distribution associated with that topic. The table shows the words with the 10 largest probabilities for each topic. "Number > 99" and "Number in [5-99]" are flags that we have inserted into the vacancy in the cleaning step to group numbers by range. Translation from Norwegian by the authors. Each entry is a single Norwegian word, although the English translations sometimes require more than one word. Clarifications in parentheses are added whenever multiple distinct Norwegian words have identical English translations.

Figure B.1: Sample List Assigned to Topic "We Offer".

We offer:

- a unique opportunity to set foot in Norway's leading exercise company.
- an incredibly exciting job that offers some real challenges that provide a lot of development as a person.
- very competitive conditions.
- we offer the industry's best training program and very good career opportunities for the right person.

Note: This figure shows the list that the model most strongly associates with the "we offer" topic. This means that the models assigns a probability of almost 1 for words in this list to be drawn from the distribution associated with "we offer." Translation from Norwegian by the authors.

Figure B.2: Sample List Assigned to Topic "Skill Requirements".

We are looking for someone who:

- has relevant education in economics, marketing and/or technical subjects.
- has good knowledge of the Norwegian retail.
- has ability to motivate and inspire others.
- communicates effectively and well, written and oral.

[...]

Note: This figure shows the list that the model most strongly associates with the "skill requirement" topic. This means that the models assigns a probability of almost 1 for words in this list to be drawn from the distribution associated with "skill requirement." Translation from Norwegian by the authors.

Figure B.3: Sample List Assigned to Topic "Tasks".

Tasks:

- personnel responsibility for all resources associated with the section.
- financial management and responsibility for results within the area.
- coordination with other sections and departments.
- be the driving force behind the company's proactive approach to security and risk.

[...]

Note: This figure shows the list that the model most strongly associates with the topic "tasks". This means that the models assigns a probability of almost 1 for words in this list to be drawn from the distribution associated with the topic "tasks." Translation from Norwegian by the authors.

Figure B.4: Sample List Assigned to Topic "Residual Category".

Start date, salary placement, and general information about the position:

- the position is a temporary position with a duration until March and is available for immediate entry.
- the position of chief police officer is placed in salary plan [number $>$ 100] to [number $>$ 100].
- our personnel policy is to have a employees from different backgrounds. We therefore encourage women and people with minority background to apply.
- the statutory pension premium is deducted from the salary
- spotless conduct is required. A certificate from the police will be obtained.
- applicants may be called to an interview.

[...]

Note: This figure shows the list that the model most strongly associates with the residual category. This means that the models assigns a probability of almost 1 for words in this list to be drawn from the distribution associated with the residual category. Translation from Norwegian by the authors.

Table B.2: Most Frequent Phrases in "We-Offer" Lists.

Rank	Phrases	English Translation	Occurrences	Assigned Attribute Category
(1)	(2)	(3)	(4)	(5)
Panel A: Unigrams				
3	"arbeidsmiljø"	"workplace environment"	138,459	Social Environment
5	"lønn"	"pay"	91,477	Any Other Mention of Pay
6	"spennende"	"exciting"	88,695	Interesting Tasks
12	"utfordrende"	"challenging"	59,382	Challenging Tasks
21	"opplæring"	"training"	37,351	On-the-Job Training
24	"pensjons"	"pension"	32,520	Pension Scheme
25	"pensjonsordning"	"pension scheme"	32,434	Pension Scheme
26	"utviklingsmuligheter"	"opportunities for development"	31,724	Good Career Paths
28	"forsikringsordninger"	"insurance schemes"	30,366	Insurance Scheme
36	"oslo"	"oslo"	23,336	Central Location
Panel B: Bigrams				
1	"godt arbeidsmiljø"	"good workplace environment"	65,143	Social Environment
2	"konkurransedyktige betingelser"	"competitive conditions"	38,016	Competitive Pay
3	"gode pensjons"	"good pension"	22,344	Pension Scheme
4	"faglig utvikling"	"professional development"	18,415	On-the-Job Traning
5	"varierte arbeidsoppgaver"	"varying tasks"	18,085	Variation in Tasks
7	"utfordrende arbeidsoppgaver"	"challenging tasks"	17,777	Challenging Tasks
8	"fleksibel arbeidstid"	"flexible work hours"	16,700	Possibility to Work Flexible Hours
10	"personlig utvikling"	"personal development"	15,617	On-the-Job Traning
11	"trivelig arbeidsmiljø"	"pleasant working environment"	14,591	Social Environment
12	"god pensjonsordning"	"good pension scheme"	12,854	Pension Scheme
Panel C: Trigrams				
1	"høyt faglig nivå"	"high professional level"	6,021	Good Colleagues
2	"god pensjonsordning gjennom"	"good pension scheme through "	3,716	Pension Scheme
7	"godt faglig miljø"	"good professional environment"	2,134	Good Colleagues
8	"pensjonsordning gjennom klp"	"pension scheme through klp"	2,009	Pension Scheme
9	"godt sosialt miljø"	"good social environment"	1,905	Social Environment
10	"gjennom statens pensjonskasse"	"public pension fund"	1,805	Pension Scheme
17	"faglig dyktige kollegaer"	"professionaly skilles colleagues "	1,195	Good Colleagues
18	"pensjonsordning gjennom statens"	"pension scheme through public"	1,183	Pension Scheme
19	"ulykkesforsikring samt fritidsulykkeforsikring"	"sport and general accident insurance"	1,159	Insurance Scheme
20	"svært godt arbeidsmiljø"	"very good social environment"	1,129	Social Environment

Note: This table documents the words occurring most frequently in the "we offer" lists extracted from the vacancies. Column 3 shows which attribute category we have assigned the word to. These categories are added to the list of attributes we search for in the vacancies. Not every word is assigned to an attribute category. The translations are provided by the authors. Every phrase is originally one word in Norwegian.

B.3 Description of Detailed Pay and Non-Pay Attributes

Tables B.3 to B.5 provides descriptions of each of the attribute we search for in job ads, as well as examples of phrases we use to detect these attributes. In our final dataset of job ads, we label an attribute as detected using an indicator variable whenever the text of the associated job ad contains any of the phrases associated with this attribute.[42]

Table B.3: Description of Pay Related Attributes.

Attribute	Description	Example Phrases (translated)	Example Phrases (Norwegian)
Compensation Scheme:			
– Compensation Level	Indication of pay, pay level, or pay range (e.g., steps in collective bargaining agreements).	wage step, wage range, wage rate	lønnstrinn, lønnsspenn, lønnssats
– Competitive Pay	Indication of good or competative salary.	good pay conditions, competitive wage, favorable pay conditions	gode lønnsbetingelser, konkurransedyktig lønn, gunstige lønnsbetingelser
– Collective Agreement Pay	Explicit reference to a collective bargaining agreement between trade unions and employers' associations.	wage step, tariff agreement, the main tariff agreement [public sector]	lønnstrinn, tariffavtale, hovedtariffavtalen
– Incentive Pay Scheme	Pay scheme dependents on an individual worker's or company's performance.	commission, bonus scheme, sales bonus	provisjon, bonusordning, salgsbonus
– Hiring Bonus	One time bonus at time of hiring, usually to encourage applicants.	recruitment supplement, establishment supplement	rekrutteringstillegg, rekrutteringstilskudd
– Good Overtime Pay	Reference to the level of overtime pay.	overtime pay, paid overtime	overtidsbetalt, betalt overtid
– Any Other Mention of Pay	Any other reference to pay, excluding any of the attributes listed above.	pay, income, fixed salary	lønn, inntekt, fastlønn
Financial Attributes:			
– Pension Scheme	Position covered by a pension scheme. Includes reference to mandatory and additional pension schemes.	pension scheme, occupational pension scheme, contribusion based pension savings	pensionsordning, tjenestepensjonsordning, inskuddspensjon
– Insurance Scheme	Employer-sponsored insurance. Includes mandatory and additional insurance schemes, and coverage of proparty damage, various health insurances and similar.	insurance scheme, group life insurance, life insurance scheme, disability insurance, health insurance	forsikriningsordning, gruppelivsforsikring, livsforsikring, uføreforsikring, helseforsikring
– Mortgage Possibility	Employer provides access to mortgage schemes with favorable conditions.	mortgage scheme, favorable mortgage	boliglånsordning, gunstig lånetilbud
Career Opportunities:			
– Good Career Paths	Suggests that the position is a good step in advancing a career, or that there are good opportunities to advance within the employer.	career opportunities, promotions, internal career	karrieremuligheter, forfremmelse, intern karriere
– On-the-Job Training	The position includes a training program or internship or opportunities to learn new skills.	training program, professional developement, internship	etterutdaningsmuligheter, faglig utvikling, praksisplass

Note: This table describes each of the pay-related attributes detected in vacancies and documents some of the phrases used to detect these attributes. The phrases in Norwegian are used in the analysis, and the translations to English are for illustration. Comments to individual phrases are displayed in brackets.

[42]The exception to this rule is the attribute "Any Other Mention of Pay", which we construct by first searching for a set of generic words associated with pay and labelling the attribute as detected in the associated job ad, but later on set the indicator for whether this attribute is detected to zero if any of the other attributes listed under "Compensation Scheme" were already detected by the same set of word.

Table B.4: Description of Non-Pay Related Attributes (Part 1).

Attribute	Description	Example Phrases (translated)	Example Phrases (Norwegian)
Hours of Work:			
– Full-time Contract	Full time position, usually 37.5 hours per week.	full-time position, full-time employee, 100%	fulltidsstilling, fulltidsansatt, 100%
– Part-time Contract	Part-time position.	part time, part-time position, 50% [includes all numbers less than 100]	deltidsstilling, deltidsjobb, 50%
– Full-time/Part-time Choice	Employer open for both full-time and part-time hires.	full-time or part-time, full-time/part-time	heltid eller deltid, fulltid/deltid
Convenient Hours:			
– Possibility to Work Flexible Hours	Flexible work hours, commonly reflecting workers' discretion to choose when to start or stop work during a shift.	flexible working hours, flexible hours scheme	fleksibel arbeidstid, fleksibel arbeidstidsordning
– Regular Daytime Work Schedule	Regular, day-time work hours, i.e. working on evenings, nights or weekends is not required.	regular working hours	normal arbeidsdag
– Exempt from Work Hour Regulations	The job is exempt from the work hours regulation covering most workers, usually leaders, academic staff or similar.	exempt from the working environment act, leader position [legal term], doctoral fellow [always exempted]	untatt arbeidsmiljøloven, ledende stilling, doktorgradsstilling
Inconvenient Hours:			
– Shift Work	Position requires shift work where employees usually switch between working day, evening, and night shifts.	shift work, rotation scheme, night shift	skiftarbeid, turnusordning, natturnus
– Weekend/Evening/Night Work	Position requires shift work involving either working on weekends, evening, and/or night.	night shift, evening work,late night shift	nattursnus, kveldsarbeid, senvakt
– On-call Employment	Positions where the employer needs additional on-call employees, often hired temporarily.	on-call help, on-call substitute, calling help [norwegian expression],	ekstrahjelp, tilkallingsvikar, ringevikar
– Overtime Work Required	Indication that some overtime work is required in the position.	overtime work, some overtime expected, overtime in periods	overtidsarbeid, påregnes noe overtid, overtid i perioder
Contract Duration:			
– Permanent Job	Permanent job.	permanent position, permanent employment	fast stilling, fast ansettelse
– Temporary Job	Temporary job, including substitute positions.	temporary job, substitute position, one year substitute	midlertidig stilling, vikariat, ettårsvikariat
– Fixed-term Contract	This position is restricted to a fixed period, and is exempted from Norwegian restriction of the length of temporary spells.	fixed-term contract [legel term]	åremål
Workplace Attributes:			
– Social Environment	Indication of good social environment.	good social environment, good working community, great colleagues	godt arbeidsmiljø, godt arbeidsfellessakap, flotte kolleger
– Good Colleagues	Professional, knowledgable, helpful, or in other ways good colleagues.	skilled colleagues, good colleagues, knowledgeable colleagues, helpful colleagues	dyktige kolleger, gode kolleger, kunnskapsrike kolleger, hjelpsomme kolleger
– Possibility for Remote Work	Possibility to work remotely.	home office, home office scheme, work from home	hjemmekontor, hjemmekontorløsning, arbeide hjemmefra
– Shared Office Space	Shared office space.	shared office space, open office space	kontorlandskap, åpent landskap
– Inclusive Work-life Scheme	Indication that the employer provides an inclusive work-life scheme, which is a formal work-life scheme that provides more generous sick-leave arrangements.	inclusive work-life scheme, IA [abbrevation]	inkluderende arbeidsliv, IA

Note: This table and the next describe each of the non-pay-related attributes detected in vacancies and document some of the phrases used to detect these attributes. The phrases in Norwegian are used in the analysis, and the translations to English are for illustration. Comments to individual phrases are displayed in brackets.

Table B.5: Description of Non-Pay Related Attributes (Part 2).

Attribute	Description	Example Phrases (translated)	Example Phrases (Norwegian)
Task-Related Attributes:			
– Interesting Tasks	Indicates interesting, exciting, or meaningful tasks. Includes description of company, sector, or industry.	interesting, exciting, meaningful	interessante, spennende, meninsfylt
– Challenging Tasks	Indicates challenging tasks.	challenging, demanding	utfordrende, krevende
– Variation in Tasks	Varying, nonmonotonous tasks.	varied tasks, versatile tasks	varierte arbeidsoppgaver, allsidige oppgaver
– Responsibility in Job	Involves responsibilities. Includes responsibility as leader or project manager.	full responsibility, a lot of responsibilities, leadership responsibilities	totalansvar, mye ansvar, lederansvar
– Independence in Performing Tasks	Freedom to choose how to approach and solve tasks.	manage the working day, influence your own, shape your own	styre arbeidsdagen, påvirke egen, utforme egen
– Involves Leadership Responsibility	The job involves leadership responsibility.	looking for CEO, leadership role, communications director	søker daglig leder, lederrolle, kommunikasjonsdirektør
– Work Involves Travelling	The job requires traveling.	business trips, travel days, travel activity	arbeidsreiser, reisedøgn, reiseaktivitet
Minor Perks:			
– Beautiful Location	Describes the job location's environment as beautiful.	beautiful nature, magnificent nature, great tracking environment	flott natur, praktfull natur, flott turterreng
– Central Location	Central location, including mentions of major cities.	oslo [capial], centrally located, close to the city center	oslo, sentral beliggenhet, sentrumsnært
– Company Gym or Sports Team	The company offers the opportunity for paid exercise during work hours, access to fitness center/equipment, and/or has a sports team or the like.	company sports team, exercise during work-hours, fitness room, fitness center membership	bedriftsidrettslag, trening i arbeidstiden, treningsrom, treningsavtale
– Parking Space On Premises	Available/free parking spaces on premises.	parking facilities, free parking	parkeringsmuligheter, gratis parkering
– Company Vehicle	Access to company vehicle during work hours, sometimes involves company vehicle available full-time.	company vehicle, company vechicle scheme, leasing vechicle	firmabil, bilordning, leasingbil
– Any Welfare Scheme	Generic description of company welfare scheme.	welfare schemes, personnel schemes, employee benefits	velferdsordninger, personalordninger, personalgoder
– Company Cabin	The company has access to a cabin usable for employees. These cabins can be accessable for company arrangements or for employees' access during holidays or vacations.	company cabins, personnell cabin, vacation apartments	firmahytte, personalhytte, ferieleiligheter
– Occupational Health Service	Access to health professionals and doctors. Involves preventions and treatment of injuries/sickness.	occupational health service, company doctor	bedriftshelsetjeneste, bedriftslege
– Company Canteen	Access to canteen.	canteen scheme, personnell canteen	kantineordning, personalkantine
– Flexible/Extended Holidays	Description of the total length of vacation, or of extended length compared to mandatory vacation.	x week vacation [where x is a number], extra vacation	x ukers ferie, ekstra feie

Note: This table and the previous describe each of the non-pay-related attributes detected in vacancies and document some of the phrases used to detect these attributes. The phrases in Norwegian are used in the analysis, and the translations to English are for illustration. Comments to individual phrases are displayed in brackets.

C Additional Details on the Model

Here we describe the transitions to and from non-employment in our structural model. The flow measure of workers from employer j to non-employment N is the sum of a relocation part M_{jN}^R and a voluntary part M_{jN}^V. The relocation flow to non-employment is simply the fraction of employment at firm j hit by an exogenous separation shock δ_j

$$M_{jN}^R = L_j \cdot \delta_j. \tag{B.1}$$

The voluntary flow to non-employment is given by the fraction of workers who decide to leave employer j

$$M_{jN}^V = L_j \cdot (1 - \delta_j - \rho_j) \cdot (1 - \lambda_1) \cdot \Pr(N \succ j), \tag{B.2}$$

where $1 - \delta_j - \rho_j$ is the fraction of workers who are not forced to reallocate and $1 - \lambda_1$ is the fraction of workers who do not have an alternative offer in the current period. Similarly to the case of voluntary flows between employers (2), the probability that a worker employed at j moves to non-employment is

$$\Pr(N \succ j) = \Pr(V_N + \varepsilon_N \geq V_j + \varepsilon_j) = \frac{\exp(V_N)}{\exp(V_j) + \exp(V_N)}, \tag{B.3}$$

given the maintained assumption of i.i.d. idiosyncratic taste shocks $(\varepsilon_j, \varepsilon_N)$ from a Extreme Value Type I distribution.

All transitions from non-employment N to employer j are voluntary. The decision to accept a job opportunity is again subject to i.i.d. taste shocks $(\varepsilon_N, \varepsilon_j)$. The flow measure of workers moving from non-employment to employer j is given by

$$M_{Nj} = L_N \cdot \lambda_0 \cdot f_j \cdot \frac{\exp(V_j)}{\exp(V_N) + \exp(V_j)}. \tag{B.4}$$

In Equation (B.4), L_N is the measure of workers in non-employment, λ_0 is the probability that a non-employed worker gets an offer, and f_j is the chance the probability that offer is from employer j. For the identification result, it is required that the offer distribution f is the same for employed and non-employed worker, which we assume.

With the notation introduced in Equation (B.4), the value function of non-employment follows directly as

$$V_N = u_N + \beta \left\{ \lambda_0 \sum_k f_k \mathbf{E}\big[\max\{V_k + \varepsilon_k, V_N + \varepsilon_N\}\big] + (1 - \lambda_0)\mathbf{E}\big[V_N + \varepsilon_N\big] \right\}, \qquad \text{(B.5)}$$

where u_N is the flow value of non-employment and β the discount factor. The continuation value is made of two terms. With probability λ_0, non-employed workers get a draw from the offer distribution f, which they decide to accept or reject given their draw of taste shocks. With probability λ_0, they remain in non-employment. All expectations $\mathbf{E}[.]$ are taken over the i.i.d. taste shocks $(\varepsilon_k, \varepsilon_N)$.

D Predictive Power of Publicly Posted Job Attributes

Which job attributes hold the highest predictive power about the actual employer pay and non-pay values? In Section 5.2, we decompose the overall R^2 from regressions of employer values on ad attributes by sequentially adding *groups* of attributes. This approach abstracts away from the 47 *individual* attributes and four other ad characteristics, and the results of this approach may thus depend on the grouping and ordering of attributes. This section documents that the main patterns in Section 5.2 and Appendix Figure A.8 hold when we relax the dependency on the ordering and focus on individual attributes and ad characteristics.

To do this, we implement a method that sequentially identifies the attributes with the highest predictive power. We start with a regression containing only a constant term, and then add the attribute with the largest adjusted (partial) R^2 among the 51 attributes as an explanatory variable (with additional controls). We continue by sequentially adding attributes that increase the explanatory power further, adding one variable at a time. This procedure delivers a tractable sequence of explanatory variables where variables with the most explanatory power are added first. It reveals, at each step, the variable with the largest predictive power given the variables already added. However, it does not identify the set of n variables with the largest *joint* predictive power at each step n.

Figure D.1: Explanatory Power as a Function of the Number of Regressors.

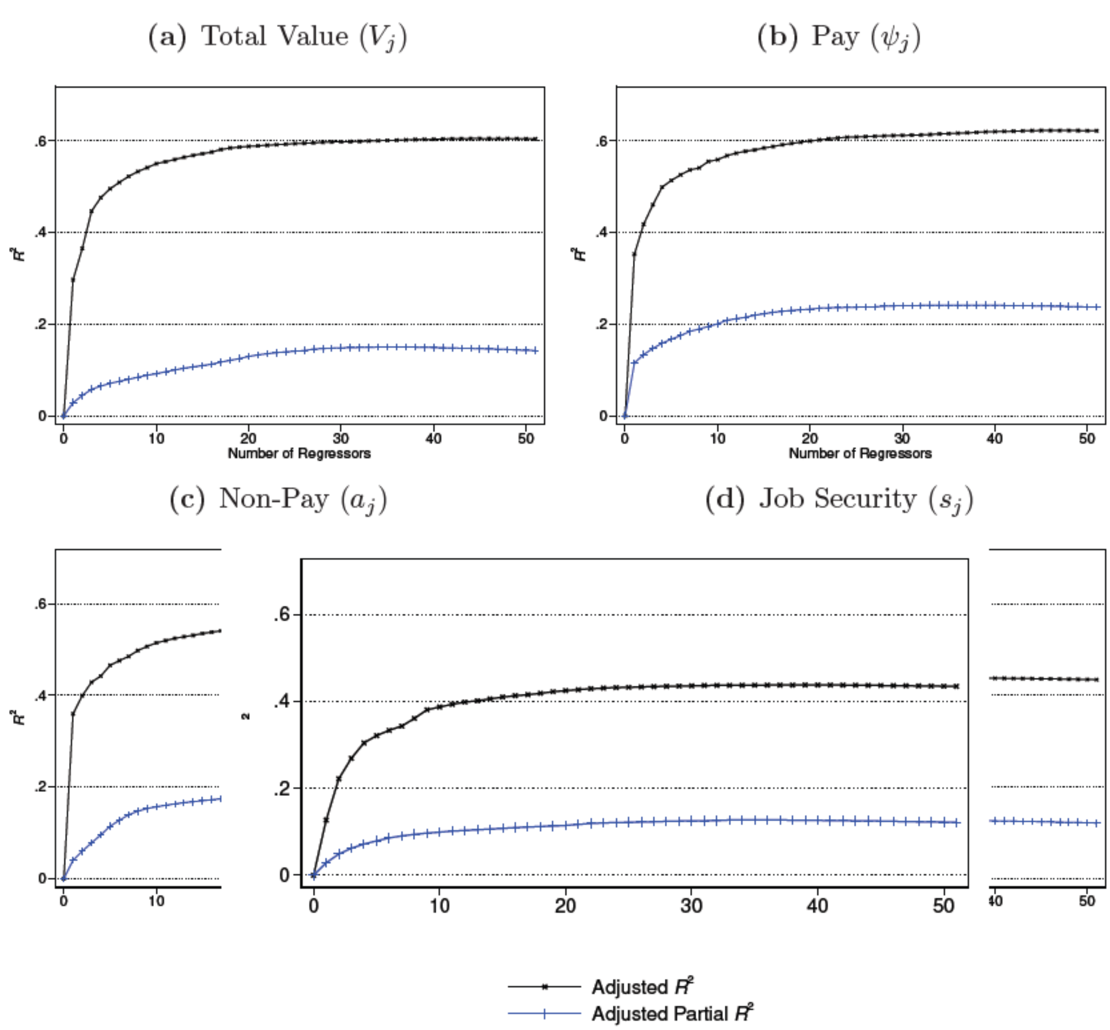

Note: This figure documents the change in explanatory power in regressions models when adding variables representing individual job attributes one at a time, in terms of adjusted R^2 for regressions without additional controls (black line) and in terms of partial R^2 for regressions with additional controls for industry, occupation and location (blue line). In each case, we start with a regression containing only a constant term, and iteratively add the variable that increases the adjusted R^2 the most.

Figure D.1 shows that 10-20 attributes contain most of the explanatory power in regressions of each model parameter on job attributes. The black lines show sequential adjusted R^2 measures in regressions without additional controls, while the blue lines show sequential partial R^2 measures in regressions with additional controls. Table D.1 summarizes the first 5, 10, and 20 attributes added to these regressions, with and without additional controls.

For regressions without additional controls, we find that attributes related to convenient or inconvenient hours are important predictors (i.e., among the first five added variables) for all model components. This confirms the results in Section 5.2. For pay values, we find that attributes related to financial attributes and hours of work are important, but also "regular daytime work schedule" and "shift work". For non-pay values, we also find that "regular daytime work schedule", "shift work" and "weekend/evening/night work" are important, but also "social environment" and "permanent job". This confirms the observation that non-pay values are related to particular types of attributes posted in job ads (Section 5.2), and uncovers that "regular daytime work schedule" and "shift work" carry a significant portion of the prediction power among convenient and inconvenient hours attributes. Finally, for job security, the presence of "any welfare scheme" also appears to be important, which may signal that the company has in place specific schemes to value their workers' welfare.

For regressions with additional controls, we find broadly similar patterns. The various attributes related to convenient and inconvenient hours again feature prominently among those with the highest predictive power, while "permanent job" is now reassuringly the first attribute added in regressions for job security with the additional controls.

Table D.1: The Most Predictive Publicly Advertised Job Attributes.

	No Additional Controls (Adjusted R^2)				With Additional Controls (Adjusted PartialR^2)			
	Total Value (V_j)	Pay (ψ_j)	Non-Pay (a_j)	Job Security (s_j)	Total Value (V_j)	Pay (ψ_j)	Non-Pay (a_j)	Job Security (s_j)
Detected Job Attribute:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Compensation Scheme								
– Compensation Level		*		*	*			*
– Competitive Pay	*			*				
– Collective Agreement Pay		*				*		
– Incentive Pay Scheme	**		*	***	**		*	**
– Hiring Bonus	*	**					*	
– Good Overtime Pay								*
– Any Other Mention of Pay				*				
Financial Attributes								
– Pension Scheme	**	*		***				
– Insurance Scheme	***	***		*	***	***		
– Mortgage Possibility				**	*			***
Career Opportunities								
– Good Career Paths		*	*		**	**		
– On-the-Job Training								
Hours of Work								
– Full-time Contract				**	*	*	*	*
– Part-time Contract		***	*	**		*	**	***
– Full-time/Part-time Choice	***		*	**	*	*	**	**
Convenient Hours								
– Possibility to Work Flexible Hours		**	**	*		***	***	
– Regular Daytime Work Schedule	***	***	***	***	***	***	***	***
– Exempt from Work Hour Regulations	***	*	*		***	**		
Inconvenient Hours								
– Shift Work	**	***	***	*		*	***	
– Weekend/Evening/Night Work	***		***	***	***	*		
– On-call Employment			*				*	
– Overtime Work Required				*				*
Contract Duration								
– Permanent Job		*	***		**	**	***	***
– Temporary Job	*	*		**				*
– Fixed-term Contract								
Workplace Attributes								
– Social Environment	**		***					*
– Good Colleagues		***	*		*	*	*	
– Possibility for Remote Work	**		*					
– Shared Office Space					**		*	
– Inclusive Work-life Scheme								
Task-Related Attributes								
– Interesting Tasks	*			*	*			**
– Challenging Tasks						**	*	
– Variation in Tasks		*	**					**
– Responsibility in Job			*			*		*
– Independence in Performing Tasks	*			*	*			
– Involves Leadership Responsibility	*	**			***		*	*
– Work Involves Travelling								
Minor Perks								
– Beautiful Location					*		*	
– Central Location		*	**			***		**
– Company Gym or Sports Team		**	*					*
– Parking Space On Premises	*				**		**	
– Company Vehicle								
– Any Welfare Scheme			**	***	*		**	***
– Company Cabin	*					**		*
– Occupational Health Service	*							
– Company Canteen		*						
– Flexible/Extended Holidays						*	*	
Other Ad Characteristics								
– Number of Words	*	**	**			***	***	
– Posted Online						*	**	
– Employer Name Disclosed								
– Number of Attributes				*	*			

*** Model with 5 regressors, ** Model with 10 regressors, * model with 20 regressors

Note: This table documents the variables selected in restricted models with 5, 10 and 20 attributes. Columns (1)-(4) show the first variables selected in regressions without including additional controls, while columns (5)-(8) show the first variables selected in regressions including controls for occupation, industry and location. We start with a regression on a constant term, and iteratively add variables that increase the adjusted (partial) R^2 the most. Regressions with 5 regressors include the first 5 variables added using this procedure. Regression with 10 regressors also includes variables selected in models with 5 regressors, etc. Industry and occupation controls are defined at the two-digit level. Location indicators split local municipalities into deciles, with municipalities in the same group having a similar number of workers.