Well-Being Salience in Corporate Disclosures and Worker Health*

Riccardo Di Francesco† Peter Hull‡ Seetha Menon§¶

September 8, 2026

Abstract

We study how firms' orientation toward employee well-being shapes the health of their workers. Because this orientation is not directly observable, we measure it from the language Danish firms use in their public disclosures, applying language models to construct a firm-level well-being salience score. We link this score to population-wide administrative registers of health and sickness absence over 2013–2022, and exploit workers' mobility between firms to estimate how workers respond to a change in the salience of their employer. We find an asymmetry: workers who move to higher-salience firms take substantially more sickness absence, while their use of health care does not change. Because worsening health would raise both, the absence response suggests a behavioral change rather than deteriorating health. Broadly, our results show that firms shape not only worker health but how workers act on it, and that the language of routine corporate disclosure can reveal economically meaningful features of the workplace that administrative data leave unrecorded.

Keywords: Worker health, mental health, firms and workers, sickness absence, mover design, text as data.

JEL Codes: C45, D22, I12, J63, J81

* The work presented here is supported by the Carlsberg Foundation, grant CF23-0987. We would like to thank Julius Koschnick and Alexander Schandlbauer for their feedback and suggestions. We are also grateful to seminar participants at VIVE, University of St. Gallen, and University of Fribourg, as well as conference participants at CIML 2026, EAYE 2026, IAAE 2026, and EEA 2026, for helpful comments and discussions. We also thank Mads Fristed Navne and Rikke Engedal Lau for invaluable research assistance.

† Department of Economics, University of Southern Denmark, Odense, Denmark, email: rdif@sam.sdu.dk.

‡ Department of Economics, Brown University, email: peter_hull@brown.edu.

§ Department of Economics, University of Southern Denmark, Odense, Denmark, email: smr@sam.sdu.dk.

¶Corresponding author.

1 Introduction

Firms and workplaces are central to people's lives. Most adults spend a large share of their waking hours at work, and a growing body of evidence shows that where people work shapes their wages, careers, and health (Ahammer et al., 2024). The relationship runs in both directions: Rettl et al. (2022) show that the profitability of Danish firms declines in years when their workers are more severely affected by seasonal influenza, highlighting that firms have a direct financial stake in the health of their workforce. Yet the features of a workplace likely to matter most for workers are precisely those that are hardest to observe. Administrative registers, however rich, offer no direct view of firms' priorities and internal practices. Survey-based measures of workplace culture, though informative (e.g., Amore et al., 2025), are costly to collect, often confined to particular sectors or countries, and difficult to scale. As a result, while we know that firms differ in ways that affect their workers, we know much less about which aspects of the workplace matter.

In this paper, we study how firms' orientation toward employee well-being relates to the health of their workers, using firm-level text and population-wide administrative data from Denmark over 2013–2022. Because this orientation is not directly observable, we measure it from the language Danish firms use in their annual reports—the statutory financial statements they file with the Danish Business Authority. Applying deep-learning language models to this text, we construct a firm-level *well-being salience score* by measuring how prominently the theme of employee well-being is discussed in a firm's disclosures. We then link our score to administrative registers covering the entire population, which record health care use and sickness absence for every employed individual, to estimate how workers' health responds to the salience their employer gives to well-being.

The well-being salience score is constructed in two steps. First, we use validated psychological dictionaries (Tausczik and Pennebaker, 2010; Boyd et al., 2022) to define a set of reference terms for well-being and mental health. Second, we map both these terms and each report's sentences into a common semantic space using a pretrained language model (Reimers and Gurevych, 2019). We then score each report by the semantic similarity between its language and the reference terms, assigning higher values to firms whose reports contain passages linguistically close to well-being and mental-health themes. The resulting measure is purely semantic: it captures linguistic proximity to well-being themes, not the tone or sentiment with which they are discussed.

We establish the construct validity of our salience score in three ways. First, we ask whether the score reflects industry- or sector-specific boilerplate rather than firm choices. We find that the distributions of salience scores overlap heavily across industries, and industry fixed effects explain essentially none of the score's variation, so the variation we exploit is genuinely cross-firm. Second, we ask whether the score merely proxies for firm attributes already recorded in administrative data. We show that the score is largely unrelated to standard measures of workplace structure and workforce composition, indicating that it captures a dimension of firms that registers do not. Third, we ask whether the score is linguistic noise unrelated to anything workers experience. We document that the score consistently predicts worker health-care use and sickness absence, confirming that it carries information about workers' actual health.

We then turn to our main question: does a firm's emphasis on well-being affect its workers' health? To answer it, we adapt a mover design (Finkelstein et al., 2016; Badinski et al., 2023; Ahammer et al., 2024), tracking workers who change firms and relating their outcome trajectories to the salience differential between their origin and destination firms.[1] The design uses the fact that a worker's employer—and with it its well-being salience—changes at the move while their time-invariant characteristics do not, such that within-worker changes in outcomes can be attributed to the change in workplace. The identifying assumption is straightforward: that the timing of a worker's move is unrelated to any sudden changes in their health or other outcomes. We provide several pieces of evidence to support this assumption: workers do not sort systematically toward higher- or lower-salience employers, the direction of a move is unrelated to observable worker traits, and outcomes show no differential pre-trends that would signal deteriorating health before a move.[2]

Our central finding is an asymmetry. Workers who move to higher-salience firms take substantially more sickness absence afterward on both the intensive and extensive margins, yet their use of health care does not change. A genuine deterioration in health would be expected to raise both. That absence responds while care does not is hard to reconcile with workers becoming less healthy at higher-salience firms, and points instead to a behavioral channel: where well-being is salient, workers take and report sick leave more readily, without a corresponding change in their measured health.

We look more closely at this absence response in several ways. The first is a heterogeneity analysis. The response appears throughout the workforce, with similar estimates across gender, age, occupation, and commuting status. Splitting movers by their pre-move use of

[1] Our design exploits the same variation as the literature on firm wage effects (Abowd et al., 1999; Card et al., 2013, 2018; Bonhomme et al., 2026), which compares movers' earnings across employers to separate the roles of firms' pay policies and workers' earning capacity in wage determination.

[2] The Danish setting is well suited to a mover design. Health care is universal and free at the point of use, which removes the income- and access-related channels through which firm characteristics could mechanically affect care use. Differences across firms therefore more plausibly reflect the working environment than workers' ability to afford treatment. Moreover, the labor market's *flexicurity* model, with its high rates of job mobility, provides the large number of movers the design requires.

mental-health care, however, separates the two margins: both groups become more likely to take some leave, while the additional days accrue to those with a documented history of care. Notably, even these workers' use of medical care never moves, ruling out worsening health precisely where it would have been easiest to detect.

Second, we find that the additional absence does not consist of casual sick days alone. Most of the additional spells are short, lasting at most one working week, yet the few additional long spells are so day-intensive that they carry roughly half of the additional days. Together with the heterogeneity results, this suggests that higher-salience firms do not simply induce more absence: they may make workers more willing to take the leave their health requires.

Third, we ask whether the response originates with the firms or with the workers, turning to workers who never change firms during the sample window, whom we refer to as the *stayers*. Among stayers, any variation in the salience a worker faces originates entirely with the firm, providing a complementary source of variation to that of the movers. The asymmetry replicates in this sample: stayers' absence rises with the salience of their employer, while none of the medical outcomes moves. Because the two samples share neither workers nor the source of variation in salience, the same pattern emerging in both strengthens the firm-side interpretation of our main results.

A final concern is compensating wage differentials. If well-being is an amenity, workers who move to higher-salience firms may accept lower wages in exchange for it, so that the change in compensation drives the response (Rosen, 1986). However, we show that the difference in pay between a mover's origin and destination firms is uncorrelated with their salience differential, and conditioning on it leaves the absence response unchanged. We thus find no evidence that salience is priced in workers' pay.

Our paper makes three key contributions. First, we add to a growing literature showing that firms matter for worker health. Recent work documents substantial firm-level heterogeneity in health-care utilization and well-being outcomes (Nekoei et al., 2024), and finds that firms are responsible for a substantial share of the cross-worker variation in health-care expenditures (Ahammer et al., 2024). Other work shows that worker health is adversely affected by specific firm-level events, such as financial distress (Kárpáti and Renneboog, 2024) and corporate acquisitions (Bach et al., 2021). These studies, however, largely treat the firm as a black box, leaving the features of the workplace responsible for these effects unobserved.[3] We contribute by measuring one such feature of the workplace—the salience

[3] An exception is Amore et al. (2025), who use a large-scale Danish survey to construct firm-level measures of the working environment along its physical, psychological, and social dimensions, and show that firms with better environments have longer employee tenures and lower rates of hospitalization and workplace accidents.

firms give to employee well-being—from text firms already disclose, and by identifying its effect on health outcomes by exploiting workers' mobility between firms.

Second, we complement a broader literature on occupational health, which has documented systematic differences in health across occupations, industries, and broad job characteristics. A long tradition links lower-status occupations and occupational grades to worse mental health outcomes (Marmot et al., 1997, 2013; Rugulies et al., 2013), and a related literature emphasizes the protective role of job characteristics such as task autonomy, stability, and working conditions against depression and psychological distress (Belloni et al., 2022). Within Denmark, elevated mental health risks have been documented among social workers, healthcare professionals, and educators (Wieclaw et al., 2005; Madsen et al., 2010). We complement these studies by investigating how firms shape worker mental health, speaking to differences across workplaces that analyses at the occupation or industry level cannot capture.

Third, we make a methodological contribution to a body of work in economics that uses unstructured data to measure latent dimensions of economic and political life. Prior studies have used text to quantify sentiment (Kosmidis et al., 2019) and emotionality (Gennaro and Ash, 2022) in political communication, to study deliberation within the Federal Open Market Committee (Hansen et al., 2018), and to detect stereotyping in job advertisements (Burn et al., 2022). More recent work extends this logic to non-verbal signals such as voice tone (Gorodnichenko et al., 2023) and facial expressions (Barry et al., 2025). We show how the text of corporate disclosures can be used to measure latent dimensions of firms and workplaces and integrated with administrative data, providing a template for empirical work that seeks to incorporate firm-level soft information into standard analyses. Importantly, our approach is portable across themes: the construction of the salience score relies only on a set of reference terms, so replacing them with terms for workplace safety, diversity, or environmental responsibility, for instance, yields analogous firm-level measures.

The rest of the paper is organized as follows. Section 2 describes the annual report corpus and the administrative data, and details the construction of the analytical samples. Section 3 presents the construction of our well-being salience score from corporate text. Section 4 provides an anatomy of the salience score, documenting its sources of variation, its relationship with observable firm characteristics, and its association with worker health and sickness absence. Section 5 turns to the mover analysis, where we exploit workers' mobility between firms to identify how well-being salience shapes worker outcomes. Section 6 concludes.

2 Data

2.1 Firms' Annual Reports

Our text data consist of the universe of annual reports filed by Danish firms with the Danish Business Authority over 2013–2022. Annual reports are the primary written record through which firms communicate their financial standing and operational performance to external stakeholders. In Denmark, filing is mandatory and submission is strictly enforced, which yields near-universal coverage of the firm population.[4] Reports vary in length and scope with firm size and legal form. At minimum they include standard financial statements (management statement, income statement, balance sheet, and notes), and most also contain a management review in which firms discuss non-financial priorities such as corporate social responsibility, sustainability, and workplace practices. Figure A.1 in Appendix A presents a representative example.

We obtain the reports from the Danish Business Authority's public API for financial statements, which provides programmatic access to current and historical filings. We query the API for calendar years 2013–2022 and download all returned documents. We then translate all documents into English using the open-source machine translation model M2M-100 (Fan et al., 2021).[5] Translation enables the use of well-validated English-language tools such as dictionaries and language models (e.g., Tausczik and Pennebaker, 2010; Reimers and Gurevych, 2019) and places the analysis in a language widely accessible to researchers and readers.[6]

Table 1 reports descriptive statistics for the resulting corpus. The corpus comprises more than 2.5 million documents, 150 million sentences, and 4 billion words. Reports vary widely in length, with a median of 1,463 words and 56 sentences and a long right tail of exceptionally long filings.[7] Lexical diversity is more limited, with a median type–token ratio of 0.341.[8] This reflects the standardized nature of financial-statement language, with extensive boilerplate,

[4] Firms must submit reports digitally no later than six months after the end of the fiscal year. Non-compliance triggers fines and, eventually, compulsory dissolution. Some entities, including sole proprietorships, personally owned small firms, and firms in bankruptcy, are exempt from this obligation.

[5] M2M-100 is a multilingual encoder–decoder Transformer trained to translate directly among 100 languages. We use the publicly released `facebook/m2m100_418M` model.

[6] Translating text data into English is standard practice in the social sciences (e.g., Lucas et al., 2015; Koschnick, 2025). Empirically, machine translation yields downstream measurements comparable to those from human translation (De Vries et al., 2018), and open-source models such as M2M-100 perform comparably to commercial services such as *Google Translate* or *DeepL* while avoiding usage costs and the lack of transparency of closed-source systems (Licht et al., 2024).

[7] The shortest documents are filings in which firms submit an exemption from reporting obligations rather than a complete annual report.

[8] The type–token ratio is the ratio of unique words to total words in a document; lower values indicate more repetition. Values close to one arise in very short documents, where repetition is mechanically limited.

Table 1 — Corpus size, document length, and lexical diversity.

Panel A. Corpus size				
N. documents				2,557,796
N. sentences				150,017,499
N. words				4,021,515,009
Panel B. Document length and lexical diversity				
	Words	*Sentences*	*Unique words*	*TTR*
Mean	1,572	59	503	0.365
SD	952	34	221	0.105
Min	2	1	1	0.005
P10	596	24	257	0.288
P25	1,051	41	391	0.312
P50	1,463	56	496	0.341
P75	1,972	72	613	0.378
P90	2,539	91	739	0.444
Max	78,697	2,592	9,327	1.000

Notes. Statistics computed on the translated corpus. Panel A displays statistics aggregated across all firms' reports. Panel B displays report-level statistics. TTR is the share of unique tokens (here words) in a document; lower values indicate more repetition. Values are rounded to the nearest integer, except TTR (three decimals).

accounting terminology, and recurring institutional phrases. Figure A.2 in Appendix A confirms this pattern, showing that domain-specific words such as *accounting*, *annual*, and *tax* appear in nearly all reports.

2.2 Administrative Data

The remainder of our data comes from Danish administrative registers maintained by Statistics Denmark, which provide the outcomes and individual-level characteristics used throughout the analysis. We combine registers covering mental-health-related drug use, hospitalizations for mental and physical conditions, outpatient mental-health visits, and sickness absence, as well as employment conditions and sociodemographics, all linked at the individual level via the personal identification number assigned to every Danish resident. To harmonize exposure to workplace conditions, we restrict attention to workers in full-time employment and retain only their primary job in each year. The resulting individual-year panel covers the population of full-time Danish workers in their main employment over 2013–2022.

We measure mental-health-related pharmaceutical uptake from the Danish national prescription database, which records every prescription filled at a pharmacy. We focus on the medication classes commonly prescribed in connection with mental-health conditions (anti-migraine preparations, antipsychotics, anxiolytics, hypnotics and sedatives, and antidepressants), identified using the third-level Anatomical Therapeutic Chemical (ATC) classification. For each worker-year we construct three complementary measures of use: total expenditure in DKK, the number of defined daily doses (DDDs) filled-out, and the number of prescriptions redeemed.[9] These three margins capture different aspects of use: expenditure is sensitive to price, product, and package-size choices, whereas DDDs measure treatment intensity and duration on a standardized scale, and prescriptions capture the frequency of dispensations.

We complement the prescription-based measures with two register-based measures of more severe and resource-intensive mental-health care: the annual number of inpatient days attributable to mental or behavioural disorders (ICD-10 classification), which we can measure only through 2018, and the annual number of visits to psychiatrists and psychologists. Both measures capture care provided through the public health-care system only, which is unlikely to be material for hospitalizations, as public hospitals account for more than 95% of all hospitalizations in Denmark (Amore et al., 2025).

We further construct two measures of physical health: the annual number of inpatient days attributable to respiratory conditions and to musculoskeletal conditions (ICD-10 classification), which we can measure only through 2018. These measures allow us to look for changes in physical health, which the mental-health measures above would not capture.

Finally, we measure sickness absence using register-based information on the annual number of days an individual is absent from work due to own or a child's illness. Absence captures an economically relevant margin of health-related impairment that extends beyond formal treatment uptake and may reflect both physical and mental health conditions.

We obtain employment conditions from the Danish register of labour market research, including occupation codes, industry codes, hourly wages, and job tenure.[10] We complement these data with sociodemographic characteristics (age, gender, municipality of residence, and immigration status) from the population register.

2.3 Sample Construction

Our empirical analysis uses two distinct samples built from the individual-year panel of the employed population described in the previous subsection: a firm-year sample, used to

[9] The ATC/DDD system (World Health Organization, 2024) is the international standard for measuring pharmaceutical utilization and is widely used in drug utilization research (e.g., Hollingworth and Kairuz, 2021; Hoff et al., 2025). The ATC classification groups substances by the organ or system on which they act and by their therapeutic, pharmacological, and chemical properties. The DDD is the assumed average maintenance dose per day for a drug's main indication in adults and is linked to the corresponding ATC code.

[10] Occupations are classified using DISCO-08, the Danish version of the International Standard Classification of Occupations (ISCO-08). Industries are classified using DB07, the Danish version of the European statistical classification of economic activities (NACE).

characterize and validate the salience score in Section 4, and a mover sample, used to identify its effect on worker outcomes in Section 5.

2.3.1 Firm-Year Sample

For the analysis in Section 4, we aggregate the individual-year panel to the firm-year level. For each firm-year we compute intensive-margin means for each outcome: mean drug use (expenditure, DDDs, and prescription counts), mean inpatient days, mean psychiatrist/psychologist visits, and mean absence days. We also compute two extensive-margin measures for each outcome: the share of workers with any positive value, and an indicator for whether the firm has at least one such worker.[11] We additionally compute three groups of firm characteristics: firm size (number of workers and number of establishments), workforce composition (mean age, the shares of female, Danish, non-commuting, and unionized workers, and an occupational concentration index based on the Herfindahl–Hirschman index), and employment conditions (mean hourly wage and mean tenure).

To mitigate mechanical noise in firm-level rates and to comply with Statistics Denmark's disclosure rules, we restrict attention to firm-year observations with at least ten workers. We further drop the few firm-years whose mean hourly wage exceeds 1,000 DKK. We then merge the firm-year aggregates with the annual reports using firm identifiers, retaining only firms with a filed report. This yields a sample of 171,982 firm–year observations, with 31,314 unique firms.

Table 2 compares firms with and without a linked salience score. All industry sections are represented in both groups, which are broadly similar in size, structure, and employment conditions. They differ in two related respects. Linked firms have substantially lower shares of female workers, slightly younger workforces, and somewhat lower shares of unionized and non-commuting workers. Their workers also use less health care and take less absence, on both the intensive and extensive margins, though differences in hospitalization are modest.

From this linked group, we further drop the two industry sections (*Private households with employed assistants* and *Extraterritorial organisations and bodies*) that contain fewer than ten linked firms each, jointly accounting for five firm-year observations, to comply with Statistics Denmark's disclosure rules.

2.3.2 Mover Sample

For the analysis in Section 5, we construct a second sample from the same individual-year panel, focusing on workers who change firms during the sample window. We retain workers

[11] This second measure is mechanically increasing in headcount, since larger firms more easily contain at least one affected worker.

Table 2 — Descriptive statistics for firms with and without a linked salience score.

	Not linked	**Linked**	**N. diff.**	**Log SD**
Sample size				
Firm–years	75,413	171,982	–	–
Unique firms	21,684	31,314	–	–
N. industries	21	21	–	–
Outcomes				
Avg. drug exp.	60.02	39.73	-0.110	-0.475
Avg. drug cons.	26.40	18.04	-0.287	-0.262
Avg. # prescriptions	0.45	0.30	-0.300	-0.269
Share drug users	0.10	0.07	-0.394	-0.195
Any drug user	0.83	0.77	-0.155	0.120
Avg. hosp. days (mental health)	0.0018	0.0013	-0.019	-0.675
Share hosp. (mental health)	0.0008	0.0008	0.007	0.016
Any hospitalization (mental health)	0.05	0.03	-0.070	-0.170
Avg. psych visits	0.18	0.11	-0.247	-0.256
Share psych users	0.03	0.02	-0.295	-0.226
Any psych visit	0.52	0.39	-0.257	-0.025
Avg. hosp. days (respiratory)	0.0124	0.0112	-0.017	0.041
Share hosp. (respiratory)	0.0041	0.0039	-0.012	-0.008
Any hospitalization (respiratory)	0.16	0.13	-0.091	-0.092
Avg. hosp. days (musculoskeletal)	0.0131	0.0121	-0.013	0.193
Share hosp. (musculoskeletal)	0.0057	0.0050	-0.046	-0.033
Any hospitalization (musculoskeletal)	0.20	0.16	-0.125	-0.104
Avg. absence days	2.00	0.64	-0.467	-0.726
Share absent	0.18	0.07	-0.480	-0.492
Any absence	0.59	0.43	-0.307	0.006
Firm structure				
Share large firms	0.50	0.49	-0.017	-0.000
Share multi-establishment firms	0.22	0.21	-0.037	-0.025
Workforce composition				
Avg. worker age	41.45	39.85	-0.199	-0.045
Share female workers	0.55	0.33	-0.823	-0.029
Share Danish workers	0.90	0.89	-0.084	0.059
Share non-commuting workers	0.56	0.50	-0.196	0.052
Share union members	0.65	0.60	-0.244	-0.051
Occupational HHI	0.50	0.51	0.067	0.097
Employment conditions				
Avg. hourly wage (DKK)	210.48	221.45	0.161	0.091
Avg. tenure (years)	4.39	4.46	0.022	0.096

Notes. Counts are reported first, followed by means. "Linked" indicates firm-years for which a salience score can be assigned from an observed annual report. "Share large firms" is the share of firms with an above-median worker count in a given year. Hospitalization measures are computed over the years through 2018 for which the register is available. "N. diff." denotes the normalized difference in means; "Log SD" reports the logarithm of the ratio of standard deviations.

who change firms exactly once, so that each worker has a single, unambiguous origin and destination firm. To observe outcomes on both sides of the move, we further require that a worker be observed for at least three years at the origin firm before the move and three years

Table 3 — Descriptive statistics for the full panel and the mover sample.

	Full panel	**Movers**	**N. diff.**	**Log SD**
Sample size				
N. worker–years	23,299,210	684,432	–	–
N. unique workers	3,579,085	71,714	–	–
Outcomes				
Annual drug expenditure	51.09	33.40	-0.033	-0.473
Annual drug consumption	23.08	12.98	-0.097	-0.335
N. prescriptions	0.39	0.24	-0.088	-0.308
Any drug use	0.09	0.06	-0.110	-0.177
Hospitalization days (mental health)	0.001	0.001	-0.007	-0.404
Any hospitalization (mental health)	0.001	0.000	-0.011	-0.244
N. psychiatrist visits	0.15	0.08	-0.063	-0.353
Any psychiatrist visit	0.02	0.01	-0.072	-0.256
Hospitalization days (respiratory)	0.012	0.009	-0.010	-0.364
Any hospitalization (respiratory)	0.004	0.003	-0.012	-0.097
Hospitalization days (musculoskeletal)	0.013	0.010	-0.008	-0.388
Any hospitalization (musculoskeletal)	0.006	0.004	-0.016	-0.116
Absence days	4.60	2.41	-0.183	-0.452
Any absence	0.39	0.28	-0.238	-0.086
Worker characteristics				
Age	41.28	41.44	0.013	-0.235
Female	0.49	0.29	-0.411	-0.094
Danish	0.89	0.92	0.102	-0.142
Employment conditions				
Hourly wage (DKK)	223.70	243.27	0.116	-0.089
Tenure (years)	4.31	5.81	0.250	-0.040
Union member	0.69	0.70	0.028	-0.012
Same residence	0.51	0.44	-0.157	-0.008

Notes. "Full panel" refers to all worker-year observations in the linked worker–firm panel introduced in Section 2.2. "Movers" refers to workers who change firms exactly once and are observed for at least three years at the origin firm before the move and three years at the destination firm afterwards. Means for movers are computed over pre-move worker-years only. Hospitalization measures are computed over the years through 2018 for which the register is available. "N. diff." denotes the normalized difference in means; "Log SD" reports the logarithm of the ratio of standard deviations.

at the destination firm afterward, in addition to the move year itself. This yields a sample of 684,432 worker–year observations, with 71,714 unique movers. We show in Appendix C.2 that our estimates are nearly unchanged under alternative tenure windows.

Table 3 compares the mover sample to the full panel of workers. The two are similar along most dimensions: movers and the broader workforce have comparable age, nationality, and union membership, and differ only modestly in baseline health and mental-health outcomes. Movers are, however, disproportionately male and have somewhat longer tenure, the latter is a mechanical consequence of requiring several years of observation at both the origin and destination firms. These differences bear on the external validity of our estimates—the extent to which they generalize beyond the population of movers—rather than on their internal validity, which we discuss in Section 5.1.

3 Constructing the Salience Score

This section describes the construction of our well-being salience score from corporate text. The score measures how prominently the theme of employee well-being figures in a firm's annual report, based on whether the report contains passages semantically aligned with well-being and mental-health language.

Our approach has two ingredients. The first is a dictionary of words and phrases that capture the idea of well-being, which fixes the meaning we are looking for. The second is a language model that represents the meaning of any passage of text as a point in a shared space, placing passages with similar meaning close together. We measure a firm's salience as how close the language of its report comes to the dictionary's terms in that space. The three subsections that follow take up the dictionary, the language model, and the scoring procedure in turn.

3.1 Salience Dictionary

The first ingredient of our measure is a dictionary that defines the semantic target—the set of terms against which we score the language of firms' reports. We build it from the validated word lists of the *Linguistic Inquiry and Word Count* (LIWC) (Tausczik and Pennebaker, 2010; Boyd et al., 2022), which provide thematically grouped sets of words and expressions, developed and refined by linguistic psychologists, that capture emotional, cognitive, and social dimensions of language. LIWC lists are widely used to measure latent concepts in text across economics and political science (e.g., Kosmidis et al., 2019; Gennaro and Ash, 2022).

From LIWC we select the two lists that jointly span our target concept of employee well-being: the *Mental Health* list, comprising 126 words and wildcard expressions[12] (e.g., *coping*, *depression*, *mental health*), and the *Wellness* list, comprising 118 words and wildcard expressions (e.g., *mindfulness*, *healthful*, *well-being*).[13] We fix this choice in advance, on conceptual grounds, without reference to firm outcomes. We then expand the wildcard expressions against the lemmas in WordNet's vocabulary of English to obtain an initial candidate dictionary.[14]

[12] A wildcard expression is a truncated string that matches all words sharing the same prefix—for example, *delusion** matches *delusion*, *delusional*, and *delusions*.

[13] Table A.1 in Appendix A reports the original entries in both lists.

[14] WordNet is a curated lexical database of English that organizes words into sets of synonyms and

Figure 1 — Salience dictionary word cloud.

Notes. The figure displays the terms in our well-being salience dictionary, with each term sized in proportion to its frequency in the corpus of translated annual reports.

Because our reports are machine-translated from Danish and drawn from a financial-reporting register, the candidate list requires two adaptations, which we validate manually and document in full in Appendix B. A research assistant independently read a random sample of 50 translated reports and flagged instances of poor machine translation, which we cross-checked against the candidate entries. This led us to remove four terms traced to a recurring mistranslation of a Danish accounting term, and to replace a set of pure acronyms, prone to false positives in financial text, with their full forms. The final dictionary comprises 293 terms. Figure 1 visualizes it as a word cloud. Figure A.3 in Appendix A reports document frequencies for dictionary terms.

3.2 Sentence Embeddings

The second ingredient is a model that places report text and dictionary terms in a common semantic space. A natural starting point is to score each report by counting occurrences of dictionary words, an approach widely used in the social sciences (e.g., Kosmidis et al., 2019; Osnabrügge et al., 2021). Dictionary counts, however, have two well-known limitations (Gennaro and Ash, 2022). First, because they rely on exact matches to a fixed list, any gap in coverage (e.g., missing synonyms) translates mechanically into measurement error. Second, they treat every listed term as equally diagnostic of the target concept, and every unlisted

provides lemmatized vocabularies.

term as equally uninformative, at odds with how language conveys meaning in context. These concerns are acute in our setting, where well-being and mental-health language is rare in otherwise standardized financial statements: a count-based score would be zero for most reports and would fail to register the graded differences across firms that we seek to measure.[15]

We therefore adopt an embedding-based approach increasingly used in economic research (e.g., Gennaro and Ash, 2022; Burn et al., 2022; Ash and Hansen, 2023) that captures meaning in context. An *embedding model* maps each sentence to a dense numerical vector, its *embedding*, positioned so that semantically related sentences lie close together and unrelated sentences lie far apart. In our analysis, we use a pretrained Sentence–BERT model (Reimers and Gurevych, 2019) that produces embeddings at the level of whole sentences rather than individual words.[16] Working at the sentence level, rather than the word level more common in earlier work, offers two advantages.[17] First, sentence embeddings encode compositional meaning. Many terms are polysemous, and their meaning depends on modifiers, negation, and local syntax; word-level approaches treat words largely out of context, so similarity can reflect surface lexical overlap rather than the meaning of the full proposition.[18] Second, sentence-level representations let us evaluate reports *locally*, scoring each sentence individually rather than collapsing the document to a single vector.[19] This suits our setting, where well-being and mental-health themes are typically conveyed in a small number of focused passages rather than uniformly throughout a firm's report.

3.3 Scoring Annual Reports

We now combine the dictionary and the embedding model to score each report by the prominence of well-being and mental-health language. The procedure has four steps.

First, we split each report into sentences and embed every sentence with Sentence–BERT, yielding a collection of sentence vectors for the report. Second, we embed each entry in our salience dictionary and average the resulting vectors into a single *dictionary vector*,

[15] Figure A.3 in Appendix A reports document frequencies for dictionary terms and confirms their rarity in the corpus.

[16] Sentence–BERT adapts BERT (Devlin et al., 2019) to sentence-level representations by pooling token-level outputs, mapping sentences to fixed-length vectors that place semantically similar sentences close together. We use the publicly released `all-MiniLM-L6-v2` model.

[17] Much of the earlier economics literature embeds words—typically using word2vec (Mikolov et al., 2013a,b) or GloVe (Pennington et al., 2014)—and aggregates them to the document level, for instance by averaging (e.g., Gennaro and Ash, 2022).

[18] Averaging word embeddings within a sentence is an alternative way to incorporate context, but Reimers and Gurevych (2019) show it performs poorly for sentence-level semantic similarity.

[19] Related local approaches in economics embed short word sequences such as n-grams by averaging the embeddings of words within the sequence (e.g., Burn et al., 2022).

Figure 2 — Distribution of alternative report-level salience scores.

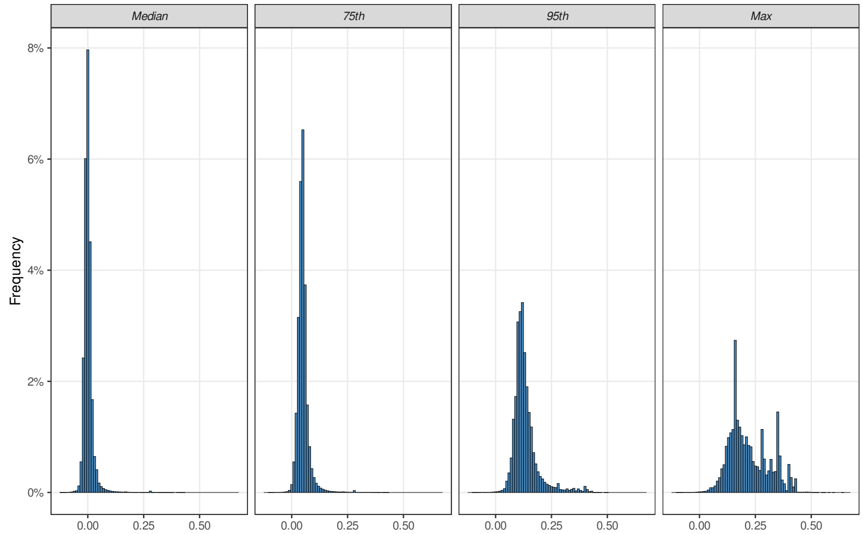

Notes. For each report, we compute the similarity between every sentence embedding and the salience dictionary vector, and summarize the resulting distribution of sentence-level similarities by a single statistic. The figure plots the cross-report distribution of this statistic under four alternative choices: the median, the 75th percentile, the 95th percentile (our baseline), and the maximum. Higher values indicate reports whose sentences are more similar to the well-being dictionary.

which represents the location in the embedding space of the well-being and mental-health concepts. Third, for each report we compute the *cosine similarity* between every sentence embedding and the dictionary vector, producing a report-specific distribution of sentence-level similarities.[20] Sentences in the upper tail of this distribution are, by construction, those most closely aligned with well-being and mental-health language.

The final step reduces each report's distribution of sentence-level similarities to a single summary statistic. In principle any percentile could serve, but different percentiles capture different features of a report and are not equally suited to our objective (e.g., Burn et al., 2022). The median reflects a report's typical sentence, which in our setting is standardized financial text and is therefore uninformative about our semantic target. Higher percentiles, by construction, isolate the passages most closely aligned with the target. The maximum, however, places all weight on a single sentence, making it highly sensitive to idiosyncratic matches; it further tends to increase mechanically with document length, since longer reports offer more opportunities for an extreme match. We therefore seek a percentile high enough to capture a report's most aligned passages, yet below the maximum, where sensitivity to individual sentences and document length sets in.

This leads us to define a firm's *well-being salience score* as the 95th percentile of its

[20] Cosine similarity is a standard measure of semantic proximity in natural language processing (e.g., Clark, 2015). For two vectors $\boldsymbol{u}$ and $\boldsymbol{v}$, it is defined as $CS(\boldsymbol{u},\boldsymbol{v}) := (\boldsymbol{u}\cdot\boldsymbol{v})/(\|\boldsymbol{u}\|\,\|\boldsymbol{v}\|)$ and ranges from -1 to 1, with higher values indicating greater semantic similarity.

report's sentence-similarity distribution.[21] Figure 2 supports this choice, illustrating empirically the trade-off it resolves. The median yields a distribution tightly concentrated near zero, consistent with the typical sentence being orthogonal to our target, while the maximum yields a far more dispersed distribution with a pronounced right tail and visible bunching, reflecting its sensitivity to rare matches. The 95th percentile captures substantially more cross-report variation than the median or the 75th percentile, yet avoids domination by extreme realizations. Figure A.4 in Appendix A reinforces the point: the maximum-based score rises sharply with document length, whereas the 95th-percentile score shows only a mild gradient.

4 Anatomy of the Salience Score

Because our salience score is derived from unstructured text, we examine its properties before using it to study worker health.

We begin by showing qualitative evidence that our approach separates language about employee well-being from routine corporate disclosure. We sample annual reports from the top and bottom deciles of the salience distribution and extract the sentences most and least aligned with our well-being dictionary. Table A.2 in Appendix A displays the most extreme examples. The contrast is sharp: the highest-salience sentences speak directly to the physical and psychological working environment, employee well-being, and sickness absence, whereas the lowest-salience are anchored entirely by financial and accounting disclosure.[22]

We then examine the score along three dimensions. First, we ask whether the score reflects industry-specific reporting language rather than genuine cross-firm variation, examining its sources of variation across firms, industries, and time. Second, we ask whether it merely re-labels firm characteristics already recorded in administrative registers, or instead captures a dimension of firms that registers do not. Third, we ask whether the score carries information about workers' actual circumstances, relating it to health-care use and sickness absence.

[21] This choice is not consequential for our findings: the association between salience and worker health is stable across alternative percentiles, as we show in Section 4.3.

[22] For example, one of the highest-salience sentences reads "The Group focuses on a healthy and secure physical and mental work environment for all employees and professional entrepreneurs on the workplace," while one of the lowest reads "Any tax is calculated by the applicable corporate tax percentage of the difference between accounting and tax values."

4.1 Variation in Salience Score

The first question is whether the salience score reflects genuine differences across firms or merely systematic differences in reporting language across industries. Some sectors may use language closer to our dictionary themes for reasons unrelated to firm priorities—firms in health-related activities, for example, may employ more mental-health terminology in their disclosures regardless of how they treat their workers. If the score primarily tracked such sectoral patterns, it would capture industry composition rather than the firm-level priorities we seek to measure.

Figure 3 plots the distribution of the salience score by industry and shows little support for this concern. The distributions overlap heavily: industries share a similar central mass, within-industry dispersion is substantial and comparable in magnitude across sectors, and a right tail of high-salience reports appears in nearly all industries rather than being confined to a few. Variation in the score is thus largely within, not between, industries.

We confirm this quantitatively in Table 4, which reports adjusted R^2 values from separate regressions of the salience score on industry fixed effects, the occupational HHI, year fixed effects, and firm fixed effects. Industry fixed effects account for 1.5% of the variation, occupational concentration for essentially none, and year fixed effects for less than 1%. Firm fixed effects, by contrast, account for roughly 65%. Variation in the salience score is therefore overwhelmingly across firms and persistent over time, rather than a reflection of industry composition, workforce specialization, or aggregate trends.

4.2 Well-Being Salience and Firm Characteristics

The second question is whether the salience score captures something registers do not. Administrative data record a firm's organizational structure and workforce composition, but not its internal policies, managerial priorities, or how it frames employee well-being. If salience were merely a re-labeling of recorded characteristics—if, say, firms with younger workforces mechanically scored higher—it would add nothing to what registers already measure. We therefore ask whether the score is spanned by standard observables or instead reflects a more latent dimension of firms' stated priorities.

Figure 4 plots the salience score against eight standard firm characteristics, each with a fitted linear relationship. The slopes are essentially flat throughout, with substantial dispersion in salience conditional on every characteristic. Firms with similar workforce composition and pay structures thus differ widely in salience, and high- and low-salience firms appear throughout the support of each observable. The score therefore does not re-label registry-based characteristics, but captures a dimension of firms' stated priorities that structure and

Figure 3 — Salience score distribution by industry.

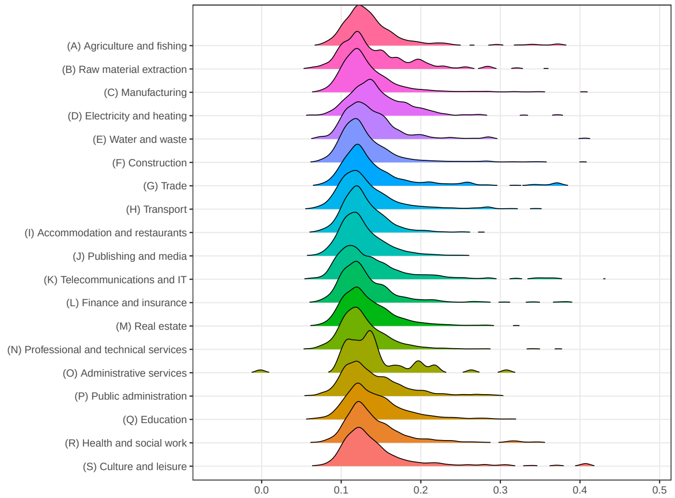

Notes. The figure shows the distribution of the firm-level salience score across industries. Industries are identified by their letter code in the DB07 classification.

Table 4 — Explained variation in the salience score.

	Salience score			
	(1)	(2)	(3)	(4)
Adjusted R^2	0.015	0.001	0.006	0.651
Industry FEs	✓			
Occupational HHI		✓		
Year FEs			✓	
Firm FEs				✓
Observations	171,977	171,977	171,977	171,977

Notes. Column (1) regresses the salience score on industry fixed effects only. Column (2) regresses the salience score on the occupational Herfindahl–Hirschman Index (HHI) only. Columns (3) and (4) regress the salience score on year fixed effects and firm fixed effects, respectively. The table reports adjusted R^2 to summarize the share of variation accounted for by each specification.

composition measures cannot proxy.

4.3 Well-Being Salience and Employee Health

The third question is whether the salience score carries information about how workers fare, rather than being linguistic noise. We relate each of the worker health outcomes to the salience score by fitting a series of linear regressions at the firm-year level. Throughout, we condition on a common set of firm characteristics capturing organizational structure (number of establishments, employment size), workforce composition (mean worker age, share of female workers, share of Danish workers, occupational concentration), and employment conditions (mean hourly wage, mean tenure, share of union members, and share of

FIGURE 4 — Well-being salience and standard firm-level characteristics.

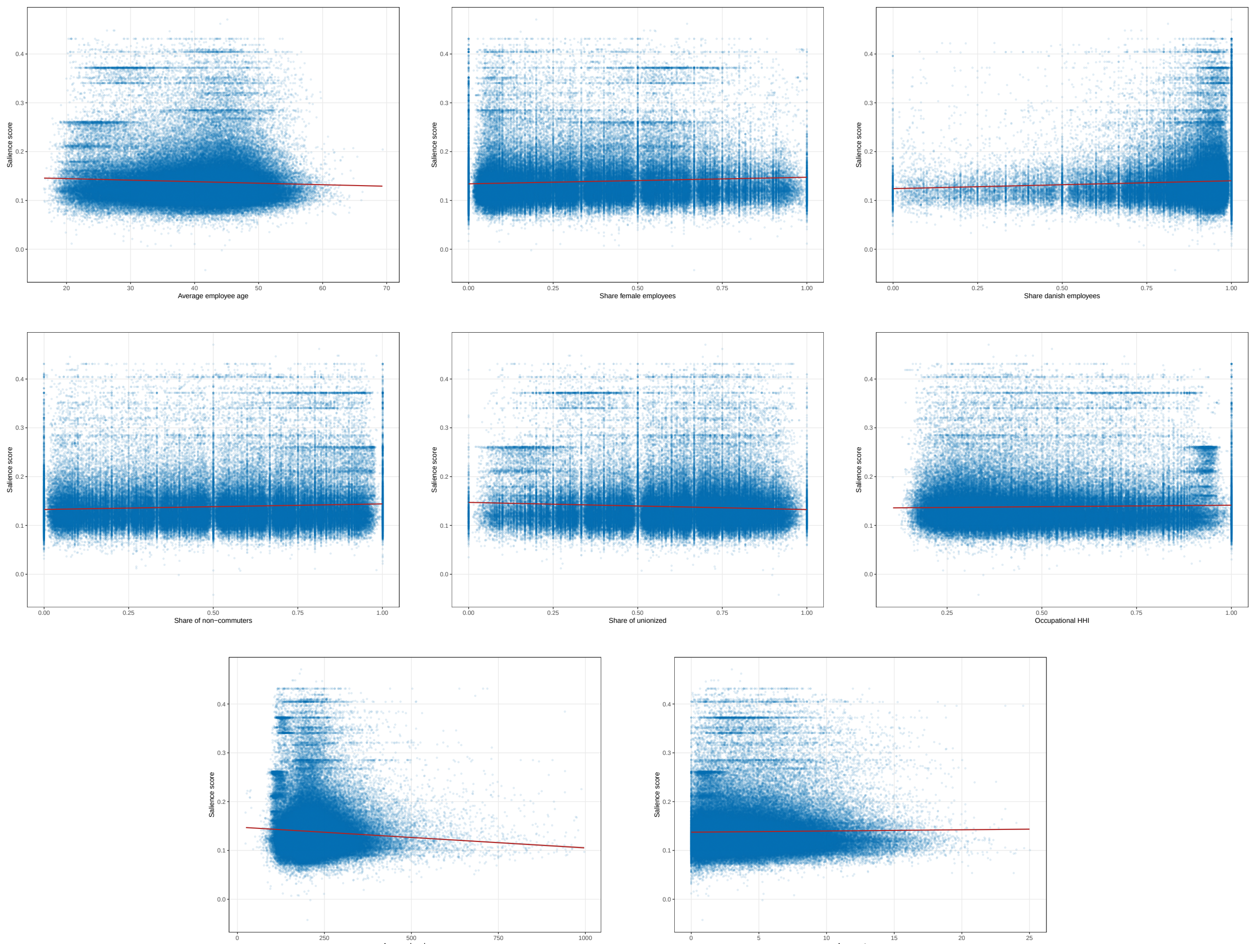

Notes. Each panel is a scatterplot of the firm-level salience score against a single firm characteristic, with a linear fit superimposed. The panels show, from top left to bottom right: mean worker age, share of female workers, share of Danish workers, share of non-commuting workers, share of union members, occupational concentration (Herfindahl index), mean hourly wage, and mean tenure.

non-commuting workers), together with industry and year fixed effects. We interpret these associations as descriptive and make no causal claim; whether well-being salience affects worker health is a question we take up in Section 5.

Table 5 reports the associations between salience and prescription drug outcomes. The salience score is positively and significantly associated with all five measures. Firms whose disclosures place greater emphasis on well-being thus have workforces with higher mental-health-related drug use, on both the intensive and extensive margins.

Table 6 extends the analysis to more severe and resource-intensive mental-health care: mental-health hospitalizations and outpatient psychiatrist/psychologist visits. For hospitalizations, salience is positively associated with all margins, mirroring the drug-use pattern. For psychiatrist/psychologist visits the pattern is mixed: the firm-level extensive margin is strongly positively associated with salience, while the share of users is negatively associated

Table 5 — Salience and worker health: prescription drug use.

	Avg. drug exp. (1)	Avg. drug cons. (2)	Avg. # prescriptions (3)	Share drug users (4)	Any drug user (5)
Salience	25.58***	6.43***	0.122***	0.008*	0.388***
	(8.60)	(1.91)	(0.035)	(0.005)	(0.031)
Outcome mean	39.73	18.04	0.304	0.073	0.774
Observations	171,977	171,977	171,977	171,977	171,977
Industry FEs	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓
Firm controls	✓	✓	✓	✓	✓

Notes. Each column reports estimates from a separate OLS regression of the indicated firm-level outcome on the salience score and a common set of firm-level controls (organizational structure, workforce composition, and employment conditions), including industry and year fixed effects. Standard errors clustered at the firm level are reported in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 6 — Salience and worker health: mental-health hospitalizations and specialist visits.

	Mental-health hospitalizations			Psych. visits		
	Avg. days (1)	Share (2)	Any (3)	Avg. visits (4)	Share (5)	Any (6)
Salience	0.002*	0.001**	0.095***	-0.024	-0.005***	0.529***
	(0.001)	(0.000)	(0.016)	(0.014)	(0.002)	(0.042)
Outcome mean	0.001	0.001	0.032	0.109	0.018	0.389
Observations	98,856	98,856	98,856	171,977	171,977	171,977
Industry FEs	✓	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓	✓
Firm controls	✓	✓	✓	✓	✓	✓

Notes. Each column reports estimates from a separate OLS regression of the indicated firm-level outcome on the salience score and a common set of firm-level controls (organizational structure, workforce composition, and employment conditions), including industry and year fixed effects. The hospitalization register is available through 2018, so columns (1)–(3) are estimated on firm-years through that year. Standard errors clustered at the firm level are reported in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

and the average number of visits is statistically indistinguishable from zero.

Table 7 reports the results for hospitalizations for respiratory and musculoskeletal conditions. The intensive margins are indistinguishable from zero for both condition groups, while the firm-level extensive margin is strongly positive.

Finally, Table 8 relates salience to sickness absence. Salience is positively and precisely associated with all three measures.

We assess the robustness of these associations in two ways. First, we re-estimate our specifications separately for male- and female-dominated firms, motivated by documented gender differences in health-care utilization (e.g., Hoff et al., 2025). The pattern from our main specification holds in both subsamples. The two groups differ mainly in magnitude: the associations with prescription drug use are markedly larger in female-dominated firms, while those with absence are larger in male-dominated firms. We report the full estimates

Table 7 — Salience and worker health: respiratory and musculoskeletal hospitalizations.

	Respiratory			Musculoskeletal		
	Avg. days (1)	Share (2)	Any (3)	Avg. days (4)	Share (5)	Any (6)
Salience	0.002	-0.000	0.294***	0.002	0.001	0.351***
	(0.004)	(0.001)	(0.030)	(0.005)	(0.001)	(0.032)
Outcome mean	0.011	0.004	0.130	0.012	0.005	0.157
Observations	98,856	98,856	98,856	98,856	98,856	98,856
Industry FEs	✓	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓	✓
Firm controls	✓	✓	✓	✓	✓	✓

Notes. Each column reports estimates from a separate OLS regression of the indicated firm-level outcome on the salience score and a common set of firm-level controls (organizational structure, workforce composition, and employment conditions), including industry and year fixed effects. The hospitalization register is available through 2018, so all columns are estimated on firm-years through that year. Standard errors clustered at the firm level are reported in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 8 — Salience and worker health: sickness absence.

	Avg. absence days (1)	Share absent (2)	Any absence (3)
Salience	2.23***	0.227***	0.391***
	(0.19)	(0.019)	(0.040)
Outcome mean	0.64	0.070	0.434
Observations	171,977	171,977	171,977
Industry FEs	✓	✓	✓
Year FEs	✓	✓	✓
Firm controls	✓	✓	✓

Notes. Each column reports estimates from a separate OLS regression of the indicated firm-level outcome on the salience score and a common set of firm-level controls (organizational structure, workforce composition, and employment conditions), including industry and year fixed effects. Standard errors clustered at the firm level are reported in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

in Table C.1 of Appendix C.1.

Second, we assess robustness to alternative definitions of the salience score, re-estimating our baseline specifications with the score built from the median, 75th-percentile, and maximum of the sentence-similarity distribution. Table C.2 in Appendix C.1 reports the results. The pattern from our main specification is stable across all four. The maximum performs least well, with several intensive-margin associations losing significance, consistent with its greater sensitivity to idiosyncratic wording and outliers.

Finally, we ask whether our findings reflect the well-being content of reports specifically, rather than report length or firm verbosity. We re-estimate the specifications replacing the salience score with the number of sentences in a firm's annual report, reporting the estimates in Table C.3 in Appendix C.1. Report length does not reproduce the salience pattern: although its coefficients are often statistically significant in our large sample, they are economically negligible. This indicates that what firms write about, not how much, drives the score's association with worker health.

5 Mover Analysis

The cross-sectional patterns documented in Section 4.3 show that well-being salience predicts worker health, with firms whose disclosures emphasize well-being employing workforces with higher health-care use and sickness absence. This association admits two very different explanations. One is that the workplace shapes worker behavior, so that cultures foregrounding well-being change how workers use care and report absence. The other is compositional sorting, whereby high-salience firms employ workers who already use more care and take more absence, and discuss well-being because they have more reason to. The two explanations carry opposite implications. The first attributes the association to firms, implying that the workplace itself affects worker health; the second attributes it to the workers a firm employs, so that the same workers would behave no differently elsewhere.

The cross-sectional design cannot separate these explanations. Because each firm is observed with a single workforce, an association between salience and worker outcomes is equally consistent with firms shaping worker behavior and with workers sorting across firms. To break this tie we turn to workers who change firms, adapting a class of *mover designs* now well established in health economics. These designs exploit individual movement across contexts—typically firms or regions—to separate person-specific from context-specific contributions to observed outcomes (Finkelstein et al., 2016; Badinski et al., 2023; Ahammer et al., 2024), akin to the use of job mobility in disentangling the roles of workers and firms in wage determination(Abowd et al., 1999; Card et al., 2013, 2018; Bonhomme et al., 2026).[23]

5.1 Empirical Strategy

Our goal is to estimate how worker health and mental-health outcomes respond to a change in firm salience. To this end, we track workers who change firms during the sample window and relate their outcome trajectories to the salience differential between their destination and origin firms. The design works because a worker's employer—and with it the salience of

[23] Finkelstein et al. (2016) use Medicare beneficiary moves across hospital referral regions to decompose geographic variation in health-care utilization into patient- and place-specific components. Ahammer et al. (2024) transport this logic to firms, decomposing health expenditures into firm- and worker-specific components using Austrian worker moves. Badinski et al. (2023) exploit both patient and physician moves to further decompose the place component.

the workplace—changes at the move while their own time-invariant characteristics do not, so within-worker changes in outcomes around the move can be attributed to the change in workplace.

We formalize the design as an event study around a worker's move. Consider a worker i who changes firms during the sample window, and let $o(i)$ and $d(i)$ denote their origin and destination firms, t_i^* the year of the move, and $r_{it} := t - t_i^*$ the relative event time. We summarize the change in workplace each worker experiences by the *salience differential*

$$\delta_i := \bar{S}_{d(i)} - \bar{S}_{o(i)}, \tag{1}$$

where $\bar{S}_j$ is the salience score of firm j averaged over the years it is observed. A positive δ_i indicates a move to a firm whose disclosures place greater emphasis on well-being; a negative value, the reverse. We then estimate the event-study specification

$$y_{it} = \alpha_i + \sum_{r \neq -1} \theta_r \mathbf{1}(r_{it} = r)\delta_i + \tau_t + x_{it}'\beta + \varepsilon_{it}, \tag{2}$$

where y_{it} is a worker health outcome, α_i are worker fixed effects, τ_t are year fixed effects, and x_{it} collects additional controls (age-by-year, industry fixed effects, and an indicator for blue-collar occupation), with the pre-move coefficient θ_{-1} normalized to zero.[24]

The coefficients of interest are the θ_r. Each measures the within-worker change in y_{it} at event time r, per unit of the salience differential, relative to the year before the move. If the workplace did not shape outcomes, a worker's trajectory would not change when the salience of their employer changes, and the post-move θ_r would be zero; a jump at the move is evidence that the workplace, and not only the worker, matters.[25]

Identification rests on the assumption that the timing of a worker's move is unrelated to contemporaneous changes in their health. The assumption would fail if, for instance, a worker moved to a higher-salience firm precisely because their health had begun to deteriorate, so that the move and the health change shared a common cause. We do not require that moves be random—only that, conditional on moving, when a worker moves is not driven by an unfolding health shock.

We offer three pieces of evidence consistent with this assumption. First, workers do not sort systematically toward higher- or lower-salience firms. Figure 5 plots the distribution of

[24] We deliberately exclude post-move outcomes of the worker, such as wages, which would be bad controls.

[25] Our specification differs from the share-of-variance designs of Finkelstein et al. (2016) and Ahammer et al. (2024), who regress an outcome on the origin–destination differential in that same outcome and obtain post-move coefficients bounded in $[0, 1]$, interpretable as the share of cross-context variation attributable to the firm or place. Because our regressor is the salience differential rather than the differential in the outcome itself, θ_r is a slope rather than a variance share.

Figure 5 — Distribution of the salience differential.

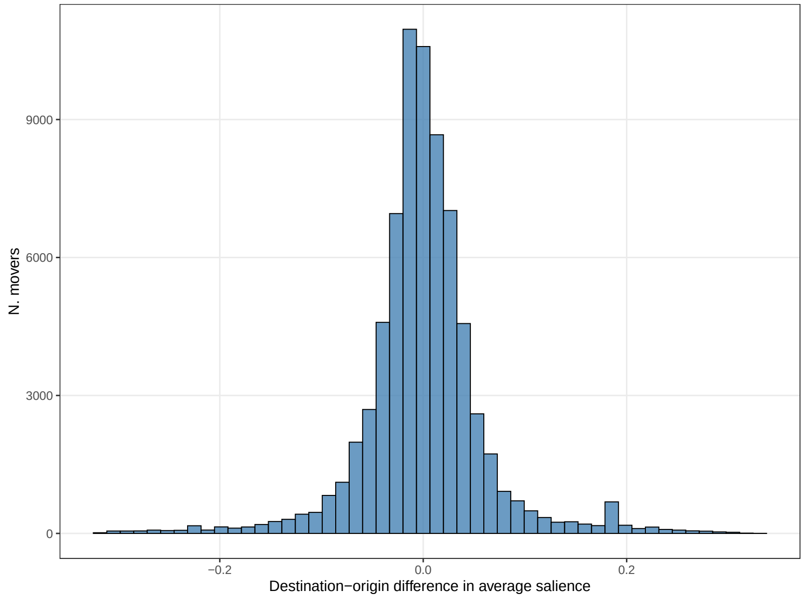

Notes. The figure plots the distribution of the salience differential δ_i across workers in the mover sample.

the salience differential δ_i across movers. It is centered near zero and roughly symmetric, with as many workers moving toward lower-salience firms as toward higher-salience ones and no aggregate tendency in either direction. Whatever drives the decision to move, it does not systematically push workers toward employers that emphasize well-being.

Second, the direction of a move is unrelated to worker characteristics. Table A.3 in Appendix A compares workers who move toward higher-salience firms ($\delta_i > 0$) with those who move toward lower-salience firms ($\delta_i < 0$). The two groups are nearly indistinguishable: normalized differences are small across demographic and job characteristics and across baseline health and mental-health outcomes, the one exception being a modest difference in baseline absence.[26] Workers moving up and down the salience gradient thus look alike before they move, as they would if the direction of a move were unrelated to the health trajectories that concern us.

Third, a deterioration in health preceding a move would surface in the event study itself. Health worsening in the run-up to a move would appear as a pre-trend in the years before $r = 0$, and an acute shock at the move would produce a jump that subsequently fades rather than persists. We find neither pattern: the pre-move coefficients are flat, and the post-move response is stable rather than transitory. We return to this evidence when we present the estimates below.

[26] We address this difference below and show that our main results are robust to controlling for each worker's baseline absence.

5.2 Main Results

Figure 6 presents the event-study estimates for our two absence outcomes over the three years before and after a move. The pattern is the same in both panels and has three features. Before the move, the estimates show no systematic trend and lie close to zero, so workers bound for higher- and lower-salience firms were on parallel absence paths, with no sign of health deteriorating in the run-up to the move. At the move, both outcomes jump sharply, rise further over the first full post-move year, and then stabilize.[27] That the effect grows and persists, rather than spiking and fading, is difficult to reconcile with a one-time health shock prompting the move and reinforces the reading of these estimates as a response to the new workplace. The magnitudes are economically meaningful: at $r = 1$, a one-standard-deviation increase in the salience differential raises absence by roughly 0.56 days—about a quarter of the pre-move mean—and the probability of any absence by about five percentage points, roughly a fifth of its pre-move mean.

We next ask whether the same workers show any corresponding change in their use of medical care. Figures 7–9 repeat the exercise for prescription drug use, mental-health hospitalizations, psychiatric and psychological visits, and hospitalizations for respiratory and musculoskeletal conditions. The contrast with absence is stark: none of these outcomes responds to the salience differential. The coefficients are small and statistically indistinguishable from zero at every event time and on every margin, with no jump at the move and no trend thereafter.

This asymmetry is the central finding of our mover analysis. When workers move to higher-salience firms, their recorded sickness absence rises sharply and persistently, while their use of health care does not change. A clinically significant deterioration in health would be expected to move both margins together: a worker whose condition worsens enough to raise their sick leave by a quarter would also fill more prescriptions, see specialists more often, or be hospitalized. The medical outcomes, however, remain flat throughout the post-move window. That absence rises while care does not therefore makes worsening health an unlikely explanation for the absence response, and points instead toward a behavioral channel: workers at higher-salience firms take and report more absence not because they have become measurably less healthy, but because the workplace makes doing so more acceptable.

[27] The estimate in the move year, $r = 0$, lies below the stabilized post-move level because moves are observed only at annual frequency: a worker who switches partway through the year is recorded at the destination from $r = 0$, but their absence that year reflects months spent at both firms. The first year in which a worker is observed entirely at the destination is $r = 1$, where the response settles.

Figure 6 — Event study for sickness absence.

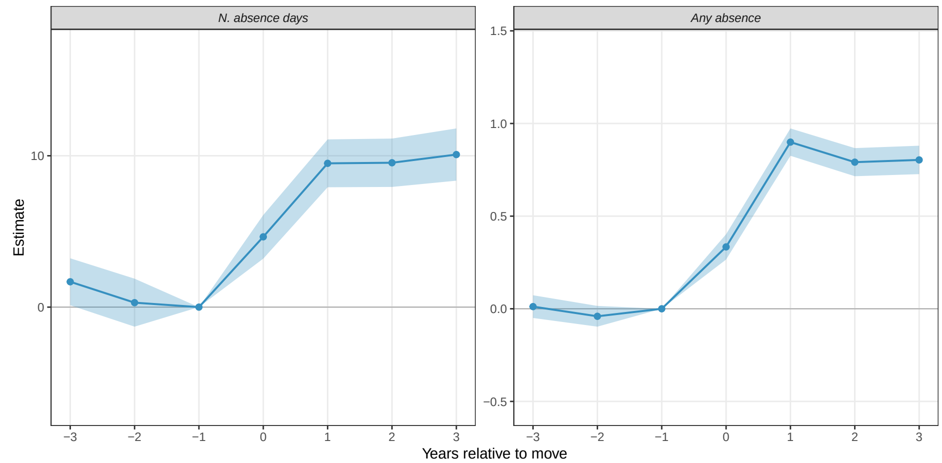

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero. The outcome is the number of absence days in the left panel and an indicator for any absence in the right panel. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Figure 7 — Event study for prescription drug use.

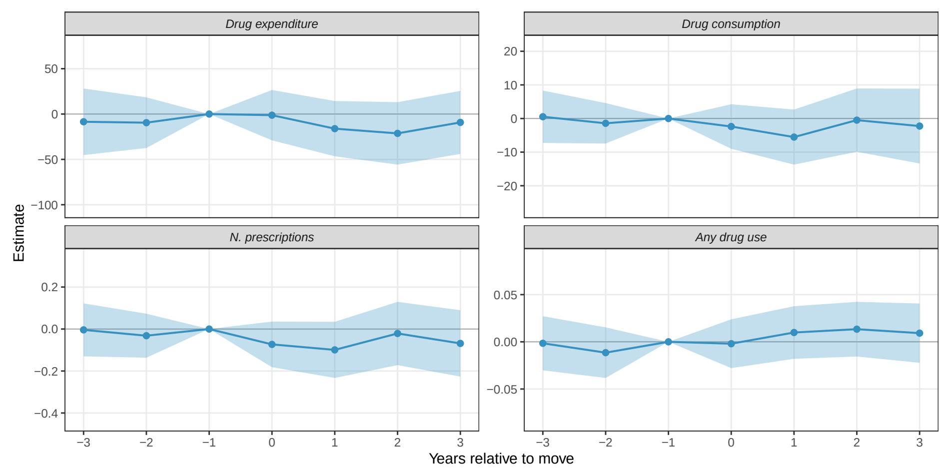

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, for four measures of prescription drug use: annual drug expenditure, annual drug consumption, the number of prescriptions, and an indicator for any drug use. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

5.3 Understanding the Absence Response

The previous subsection established that moving to a higher-salience firm raises workers' absence without any change in their use of medical care, suggesting a change in behavior rather than in health. In this subsection, we look more closely at this behavior.

We begin with who takes the additional leave. Splitting the sample by gender, age, occupation, or commuting status yields responses of similar magnitude in every subgroup (Figures A.5–A.6 in Appendix A). Splitting movers into workers with and without any pre-

Figure 8 — Event study for mental-health hospitalizations and specialist visits.

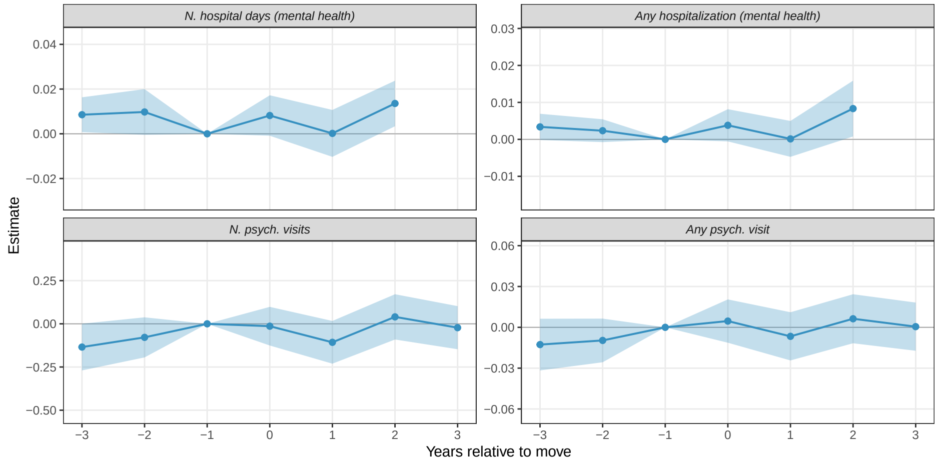

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, for mental-health hospitalizations (the number of hospital days and an indicator for any hospitalization) and psychiatric and psychological visits (the number of visits and an indicator for any visit). The hospitalization register is available only through 2018, so the hospitalization estimates are reported up to $r = 2$; the visit outcomes run through $r = 3$. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Figure 9 — Event study for respiratory and musculoskeletal hospitalizations.

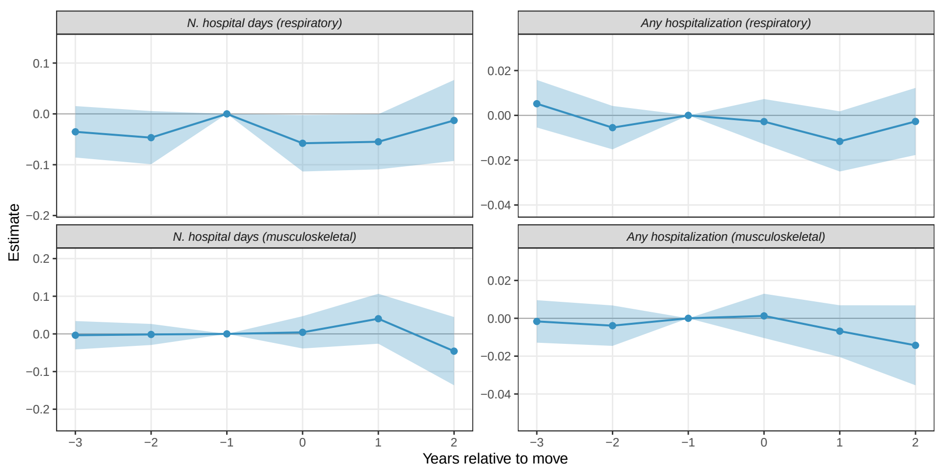

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, for hospitalizations attributable to respiratory conditions (top row) and musculoskeletal conditions (bottom row). The hospitalization register is available only through 2018, so estimates are reported up to $r = 2$. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

move mental-health care use, by contrast, separates the two margins: both groups become more likely to take some leave by similar amounts, while the additional days are roughly twice as large among workers with prior use (Figure A.7).

Notably, Figures A.8–A.9 show that the mental-health outcomes do not respond in either group, so the asymmetry replicates within each. For workers with prior use, these flat estimates are a crucial piece of evidence for the behavioral channel. These workers already fill prescriptions and see specialists, so responding to worsening health would mean intensifying care already under way, not initiating it. Had their health deteriorated at the new firm, the deterioration would have appeared in their care use. Instead, their care use remains flat, ruling out worsening health precisely where it would have been easiest to detect.

We turn next to what the additional absence consists of, examining absence spells and their duration. We classify a spell as *short* if it lasts at most five working days—one full working week—and as *long* otherwise.[28] Figure A.10 in Appendix A shows that the number of spells rises sharply at the move, but almost entirely through additional short spells, with only a few additional long ones. The same holds at the extensive margin, where Figure A.11 shows the probability of any short spell rising markedly and that of any long spell only modestly. The additional long spells, however, are long enough to matter. Figure A.12 shows that the days they carry match those from the many additional short spells, with each type accounting for roughly half of the total absence response. The additional absence thus does not consist only of scattered sick days—workers also take the sustained leave that recovery from illness requires, which a physician must certify.

We further ask whether the response originates with the firms or with the workers. To this end, we turn to workers who never change firms during the sample window, whom we call *stayers*, and test whether their absence changes as the salience of their own employer changes over time. Though descriptive, the exercise exploits a source of variation complementary to that of the movers: among stayers, any variation in the salience a worker faces originates entirely with the firm, so a firm-driven response should leave a trace in this sample. Appendix D shows that the asymmetry we found for movers is replicated among the stayers: absence rises with the employer's salience, while none of the medical outcomes moves. Because the two samples share neither workers nor the source of variation in salience, the same pattern emerging in both is hard to attribute to anything about the workers.

Last, we ask whether the response reflects compensating wage differentials rather than salience itself. Firms may price their orientation toward employee well-being, so that workers who move to higher-salience firms accept lower pay in exchange for the amenity (Rosen, 1986). If so, the change in pay, rather than in salience, could drive the absence response. However, we find that the origin–destination differential in mean wages is uncorrelated with the salience differential (Figure A.13), and conditioning on it leaves the absence response unchanged (Figures C.11–C.13). We thus find no evidence of compensating wage differentials:

[28] The distinction has institutional content in the Danish setting. The first days of an absence are effectively self-certified: the employer decides whether to require documentation, and typically does not. Longer absences must be certified by a general practitioner, and beyond a few weeks the municipality assumes its own follow-up obligations.

salience does not appear to be priced in workers' pay, and the absence response does not reflect changes in compensation.

5.4 Robustness

We conclude by assessing the robustness of the estimated absence response, with full details in Appendix C.2.

A first set of checks concerns the construction of the mover sample. Our estimates are remarkably stable when we vary the minimum-tenure requirement used to define movers from two to four years at both the origin and destination firm (Figure C.1), when we restrict the sample to movers whose entire event window predates the COVID-19 pandemic (Figure C.2), and when we exclude workers with children at home, for whom recorded absence may reflect a child's illness rather than their own (Figure C.3).

A second check concerns the design itself. Replacing each worker's salience differential with a random draw from its empirical distribution—as if origin and destination firms were assigned at random—yields estimates flat and close to zero throughout, confirming that the estimated absence response reflects the actual salience gap between origin and destination firms rather than an artifact of the event-study specification (Figure C.4).

A related concern is the timing of the differential itself: because δ_i is built from firm averages over all observed years, it incorporates reports filed after the move and is therefore partly post-treatment. Two features of our setting limit this concern: the score is constructed from report text rather than from worker outcomes, so a mover's own behavior cannot enter it, and the score is highly persistent (Table 4). We nonetheless address it directly by reconstructing the differential using only pre-move reports and re-estimating the event study, which yields the same pattern for every outcome, though the absence response is attenuated (Figures C.5–C.8). The attenuation is expected, since the new differential is constructed from fewer reports and is therefore noisier.

Our inference is also robust to the level of clustering: clustering standard errors at the origin–destination firm pair, at either firm, or two-way on both leaves every post-move estimate significant, with the exception of the move year on the extensive margin (Figure C.9).

We also address the difference in baseline absence between up- and down-movers documented in Table A.3, augmenting (2) with each worker's average pre-move absence interacted with event time. Figure C.10 shows that the absence response is unchanged at the extensive margin, while at the intensive margin it attenuates but follows the same pattern as the baseline estimates.

A final set of checks concerns what the salience differential captures. A worker moving to a higher-salience firm may simultaneously move to a larger, better-paying, or compositionally different one, and the absence response could reflect those changes rather than salience itself. Two pieces of evidence rule this out. First, the salience differential is essentially orthogonal to the analogous origin–destination differentials in the firm characteristics we observe, with all correlations being negligible (Figure A.13). Second, augmenting (2) with these differentials, each interacted with event time, leaves the estimates essentially unchanged for every outcome (Figures C.11–C.14).

6 Conclusion

This paper develops a scalable, text-based measure of how prominently firms foreground employee well-being in their corporate disclosures, and links it to population-wide administrative data on worker health. Applying sentence-embedding models to the universe of Danish annual reports, we construct a firm-level well-being salience score that captures a dimension of the workplace that administrative registers do not record. The score varies overwhelmingly across firms rather than industries, is largely unrelated to standard measures of firm structure and workforce composition, and reflects the content of disclosures rather than their length.

We use this measure to ask whether a firm's emphasis on well-being is reflected in the health of its workers. Exploiting the mobility of workers across firms, we track those who change employers and relate their outcome trajectories to the salience differential between origin and destination. The estimates reveal an asymmetry: moving to a higher-salience firm raises sickness absence sharply and persistently, while leaving health-care use unchanged. Because a genuine deterioration in health would raise both, the absence response appears behavioral rather than clinical.

Looking more closely, we find that the additional days do not fall evenly across workers, since they accrue to those with a documented history of care; that the response is not confined to scattered sick days, since half of those days come in spells lasting more than a working week; that it does not originate with the workers, since the same pattern appears among workers who never move; and that it does not reflect compensation. Higher-salience firms therefore appear to lower the threshold at which a worker converts poor health into recorded absence, rather than to change how healthy workers are or merely to hand out additional days off. Several channels are consistent with this reading: reduced stigma around taking leave, better information about the leave available, and greater managerial tolerance.

More broadly, our paper illustrates how the language firms use in routine disclosures can be turned into economically meaningful, firm-level measures of otherwise unobserved priorities, and integrated with administrative data at population scale. The approach is portable: because the salience score is built from a set of reference terms, replacing them yields analogous measures of other workplace dimensions—safety, diversity, environmental responsibility—across other corpora and countries. It offers a practical template for bringing soft information about firms into empirical work on work, health, and well-being.

A Additional Figures and Tables

FIGURE A.1 — Example annual report.

Deloitte
Statsautoriseret
Revisionspartnerselskab
CVR-nr. 33963556
Tværkajen 5
Postboks 10
5100 Odense C

Telefon 63 14 66 00
Telefax 63 14 66 12
www.deloitte.dk

Komplementarselskabet
Ugeavisen Esbjerg A/S
Kongensgade 110-114
6700 Esbjerg
CVR-nr. 10000211

Årsrapport 2017

Godkendt på selskabets generalforsamling, den 03.05.2018

Dirigent

Navn: Holger Refslund

Medlem af Deloitte Touche Tohmatsu Limited

Komplementarselskabet Ugeavisen Esbjerg A/S

Indholdsfortegnelse

majacobsen/2.05.2018 - 11:50/W.6.6.0/MStC_C Selskaber/E.16.2018 Status II: 0

Komplementarselskabet Ugeavisen Esbjerg A/S 3

Den uafhængige revisors revisionspåtegning

Til kapitalejerne i Komplementarselskabet Ugeavisen Esbjerg A/S
Konklusion
Vi har revideret årsregnskabet for Komplementarselskabet Ugeavisen Esbjerg A/S for regnskabsåret 01.01.2017 - 31.12.2017, der omfatter resultatopgørelse, balance, egenkapitalopgørelse og noter, herunder anvendt regnskabspraksis. Årsregnskabet udarbejdes efter årsregnskabsloven.

Det er vores opfattelse, at årsregnskabet giver et retvisende billede af selskabets aktiver, passiver og finansielle stilling pr. 31.12.2017 samt af resultatet af selskabets aktiviteter for regnskabsåret 01.01.2017 - 31.12.2017 i overensstemmelse med årsregnskabsloven.

Grundlag for konklusion
Vi har udført vores revision i overensstemmelse med internationale standarder om revision og de yderligere krav, der er gældende i Danmark. Vores ansvar ifølge disse standarder og krav er nærmere beskrevet i revisionspåtegningens afsnit "Revisors ansvar for revisionen af årsregnskabet". Vi er uafhængige af selskabet i overensstemmelse med internationale etiske regler for revisorer (IESBA's Etiske regler) og de yderligere krav, der er gældende i Danmark, ligesom vi har opfyldt vores øvrige etiske forpligtelser i henhold til disse regler og krav. Det er vores opfattelse, at det opnåede revisionsbevis er tilstrækkeligt og egnet som grundlag for vores konklusion.

Ledelsens ansvar for årsregnskabet
Ledelsen har ansvaret for udarbejdelsen af et årsregnskab, der giver et retvisende billede i overensstemmelse med årsregnskabsloven. Ledelsen har endvidere ansvaret for den interne kontrol, som ledelsen anser for nødvendig for at udarbejde et årsregnskab uden væsentlig fejlinformation, uanset om denne skyldes besvigelser eller fejl.

Ved udarbejdelsen af årsregnskabet er ledelsen ansvarlig for at vurdere selskabets evne til at fortsætte driften, at oplyse om forhold vedrørende fortsat drift, hvor dette er relevant, samt at udarbejde årsregnskabet på grundlag af regnskabsprincippet om fortsat drift, medmindre ledelsen enten har til hensigt at likvidere selskabet, indstille driften eller ikke har andet realistisk alternativ end at gøre dette.

Revisors ansvar for revisionen af årsregnskabet
Vores mål er at opnå høj grad af sikkerhed for, om årsregnskabet som helhed er uden væsentlig fejlinformation, uanset om denne skyldes besvigelser eller fejl, og at afgive en revisionspåtegning med en konklusion. Høj grad af sikkerhed er et højt niveau af sikkerhed, men er ikke en garanti for, at en revision, der udføres i overensstemmelse med internationale standarder om revision og de yderligere krav, der er gældende i Danmark, altid vil afdække væsentlig fejlinformation, når sådan findes. Fejlinformation kan opstå som følge af besvigelser eller fejl og kan betragtes som væsentlige, hvis det med rimelighed kan forventes, at de enkeltvis eller samlet har indflydelse på de økonomiske beslutninger, som regnskabsbrugerne træffer på grundlag af årsregnskabet.

Som led i en revision, der udføres i overensstemmelse med internationale standarder om revision og de yderligere krav, der er gældende i Danmark, foretager vi faglige vurderinger og opretholder professionel skepsis under revisionen. Herudover:

Notes. The figure displays three pages from an example firm annual report: the cover, the table of contents, and the auditor's report.

FIGURE A.2 — Most frequent words in the analysis corpus.

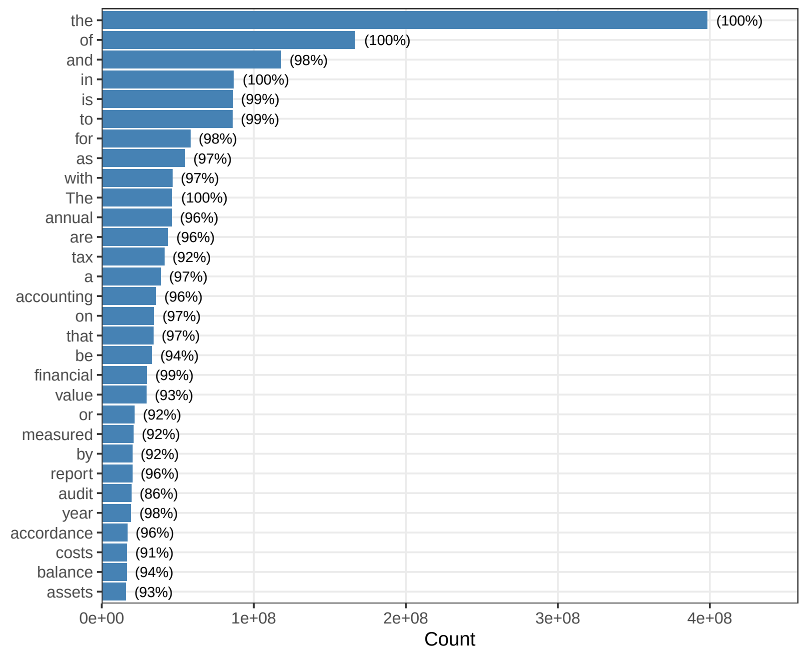

Notes. The figure shows the thirty most frequent words in the analysis corpus, ordered by total token count. Bars report total token counts; parenthetical labels report the share of reports containing each word.

Figure A.3 — Document frequency of salience dictionary entries.

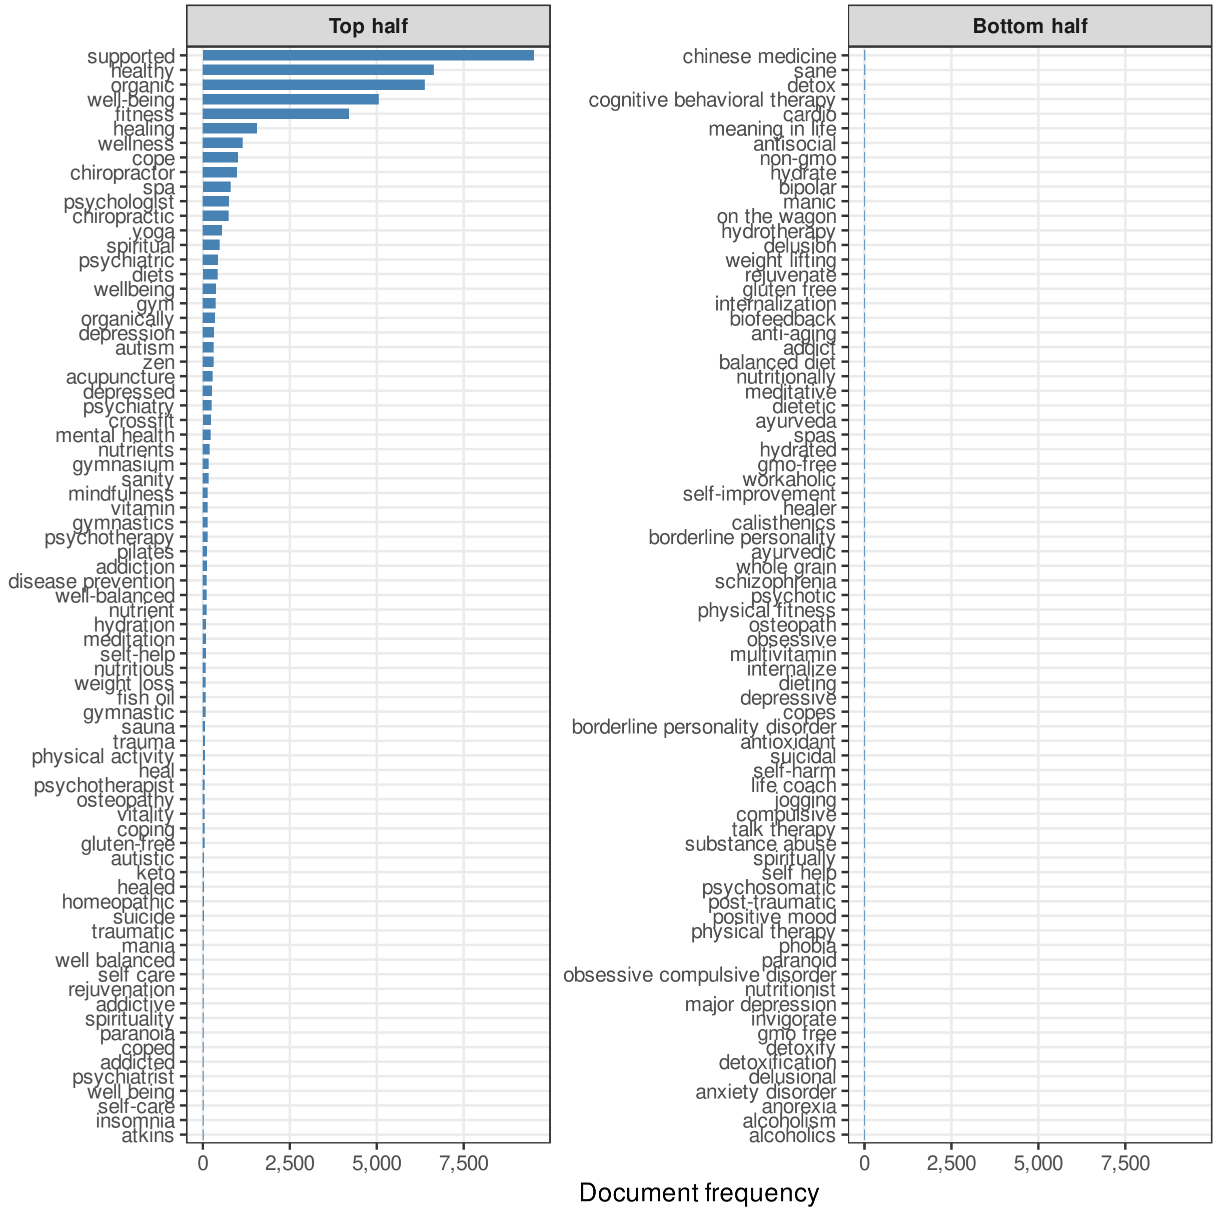

Notes. Each bar reports the number of translated annual reports containing at least one occurrence of the corresponding salience dictionary term. Entries with zero corpus frequency are omitted.

Figure A.4 — Salience scores against document length.

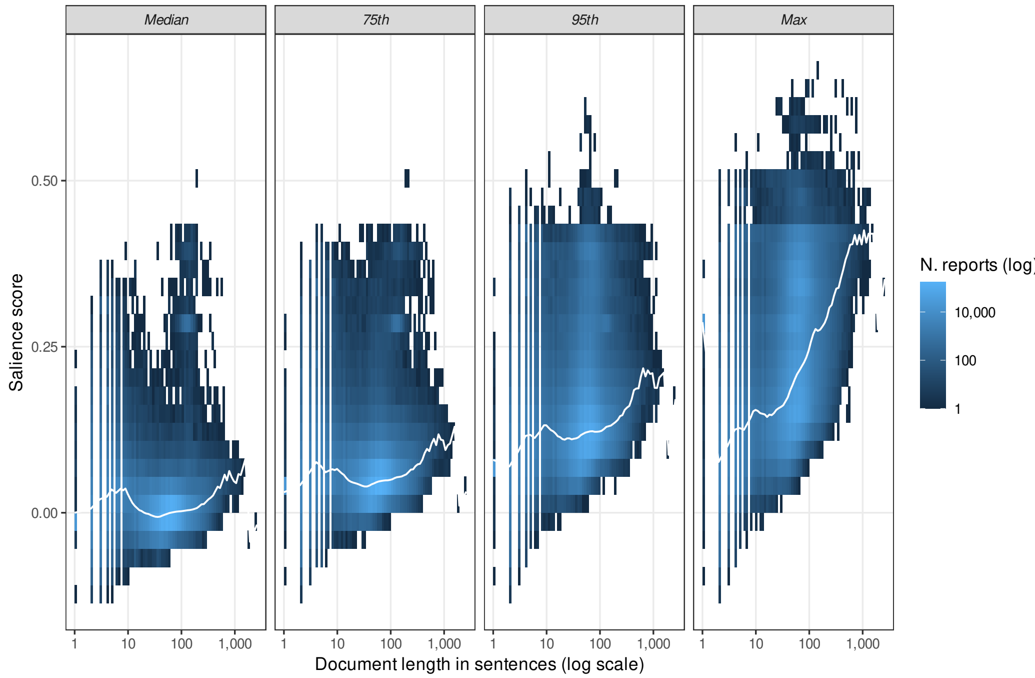

Notes. The figure plots report-level salience scores against document length, measured by the number of sentences. Columns use alternative within-report summaries of the sentence–dictionary similarity distribution. To mitigate overplotting, each panel is a binned heatmap: the length–score plane is partitioned into rectangular bins, with shading reflecting the number of reports in each bin. The white curve traces the median score within each length bin.

Figure A.5 — Event study of absence days, by worker subgroup.

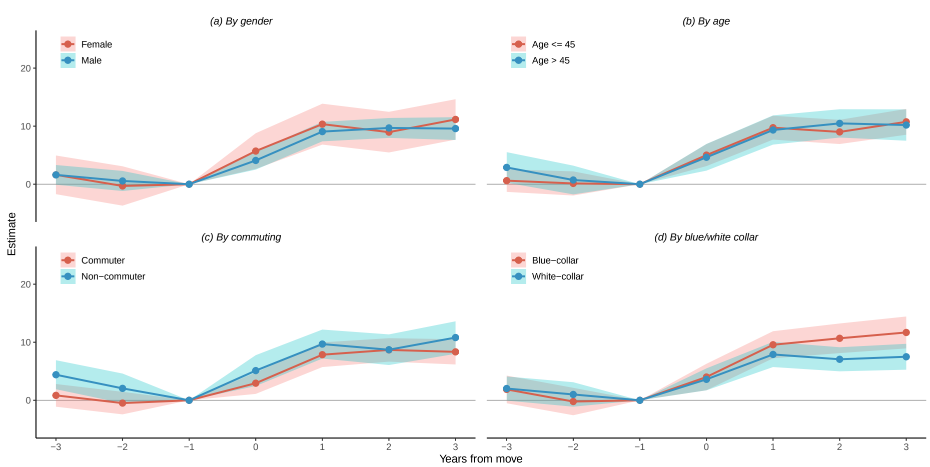

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2) for the number of absence days, estimated separately within subgroups defined by (a) gender, (b) age above or below 45, (c) commuting status, and (d) blue- versus white-collar occupation. The specification includes worker, year, age-by-year, and industry fixed effects, together with a blue-collar indicator except in panel (d), where it is dropped as the split is defined by occupation. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Figure A.6 — Event study of any absence, by worker subgroup.

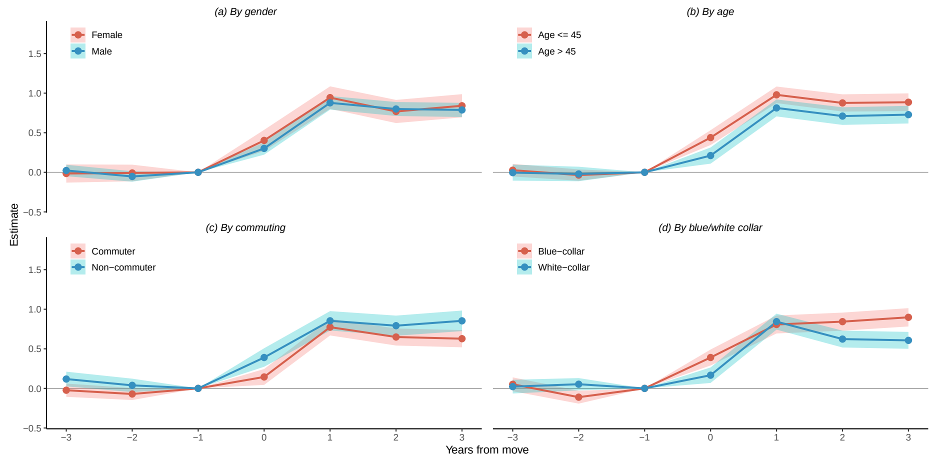

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2) for the indicator of any absence, estimated separately within subgroups defined by (a) gender, (b) age above or below 45, (c) commuting status, and (d) blue- versus white-collar occupation. The specification includes worker, year, age-by-year, and industry fixed effects, together with a blue-collar indicator except in panel (d), where it is dropped as the split is defined by occupation. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Figure A.7 — Event study for absence outcomes, by pre-move mental-health care use.

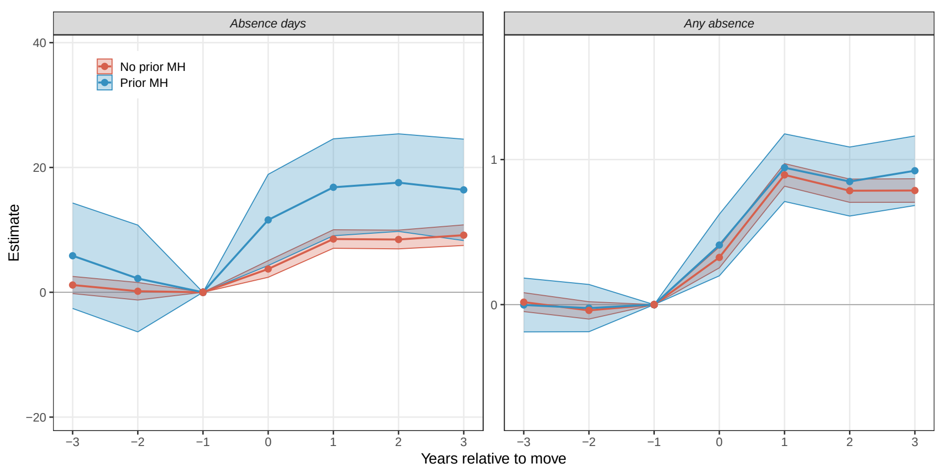

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, obtained separately for workers with and without pre-move mental-health care use. The outcome is the number of absence days in the left panel and an indicator for any absence in the right panel. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals. A worker is classified as having prior mental-health care use if they had any psychotropic drug use, any mental-health hospitalization, or any psychiatric or psychological visit in at least one pre-move year.

Figure A.8 — Event study for prescription drug outcomes, by pre-move mental-health care use.

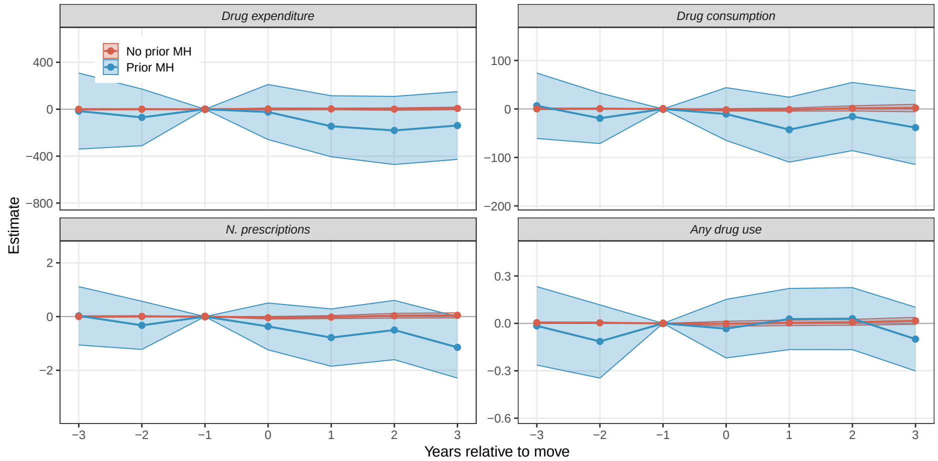

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, obtained separately for workers with and without pre-move mental-health care use, for four measures of prescription drug use: annual drug expenditure, annual drug consumption, the number of prescriptions, and an indicator for any drug use. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals. A worker is classified as having prior mental-health care use if they had any psychotropic drug use, any mental-health hospitalization, or any psychiatric or psychological visit in at least one pre-move year.

Figure A.9 — Event study for mental-health hospitalizations and specialist visits, by pre-move mental-health care use.

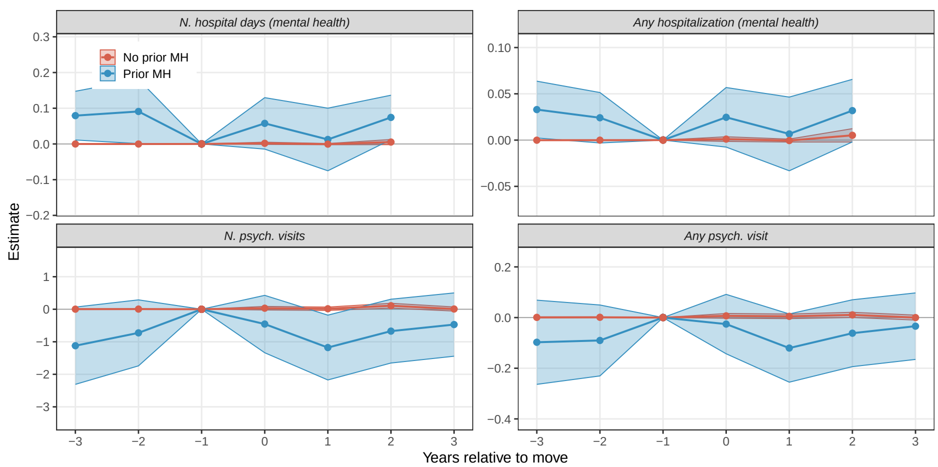

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, obtained separately for workers with and without pre-move mental-health care use, for mental-health hospitalizations (the number of hospital days and an indicator for any hospitalization) and psychiatric and psychological visits (the number of visits and an indicator for any visit). The hospitalization register is available only through 2018, so the hospitalization estimates are reported up to $r = 2$; the visit outcomes run through $r = 3$. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals. A worker is classified as having prior mental-health care use if they had any psychotropic drug use, any mental-health hospitalization, or any psychiatric or psychological visit in at least one pre-move year.

Figure A.10 — Event study for the number of absence spells, by spell duration.

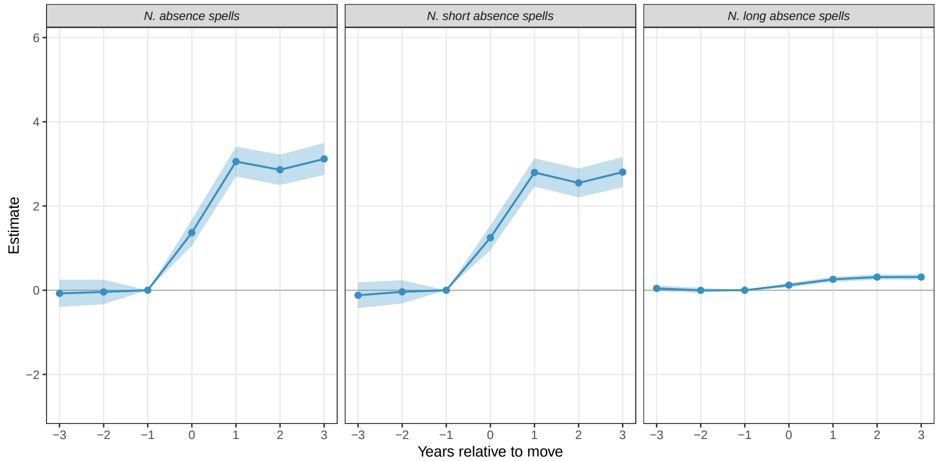

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero. The outcomes are the total number of absence spells in the year (left panel), the number of short spells (middle panel), and the number of long spells (right panel); the latter two outcomes sum to the first by construction, so the estimates in the middle and right panels decompose those in the left panel exactly. A spell is classified as short if it lasts at most five working days and as long otherwise. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Figure A.11 — Event study for the incidence of absence spells, by spell duration.

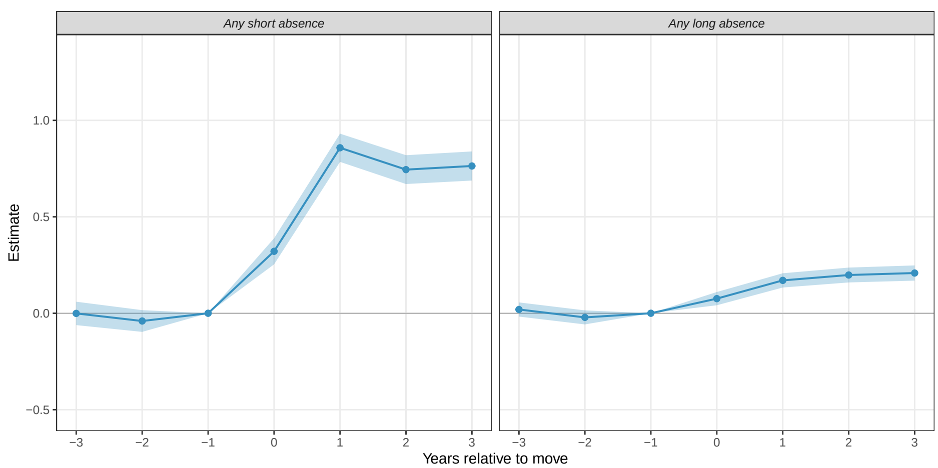

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero. The outcome is an indicator for any short absence spell in the year in the left panel and for any long spell in the right panel. A spell is classified as short if it lasts at most five working days and as long otherwise. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

FIGURE A.12 — Event study for absence days, by spell duration.

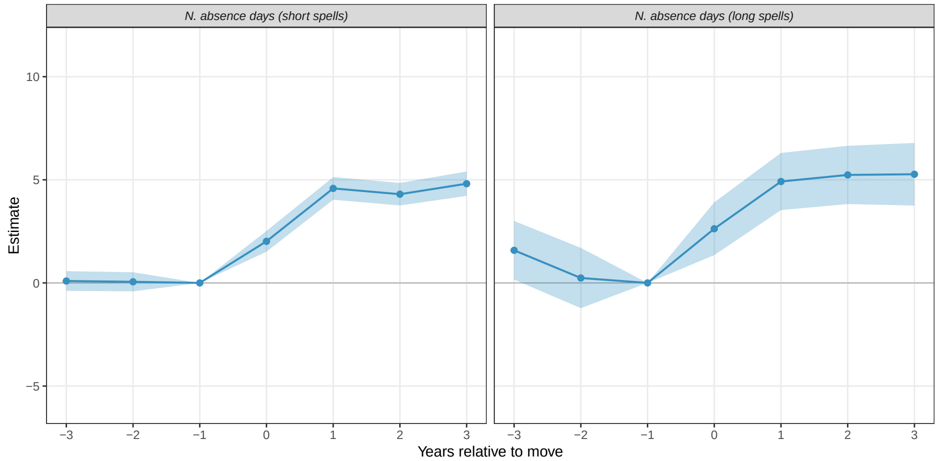

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero. The outcome is the number of absence days accumulated in short spells in the left panel and in long spells in the right panel; the two outcomes sum to the total number of absence days by construction, so the estimates decompose the response in Figure 6 exactly. A spell is classified as short if it lasts at most five working days and as long otherwise. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

FIGURE A.13 — Correlations between the salience differential and firm-characteristic differentials.

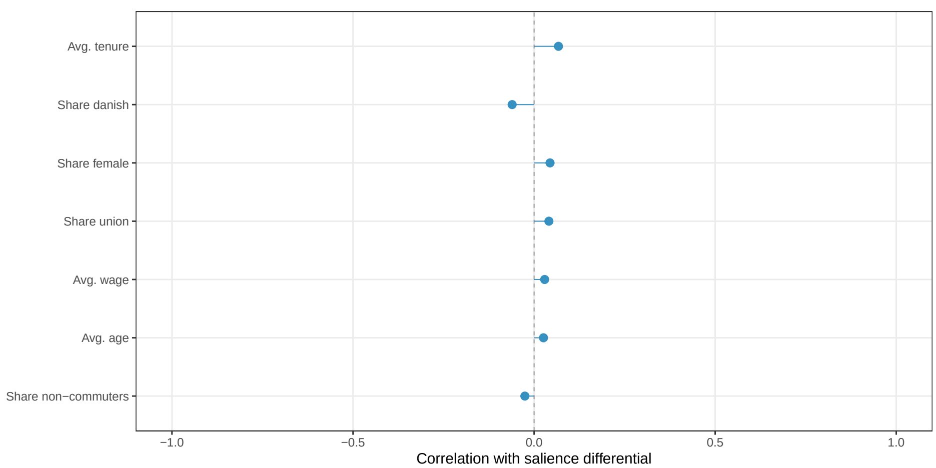

Notes. The figure plots the correlation between the salience differential δ_i and the analogous origin–destination differentials in firm characteristics, computed across the movers in our sample. Each differential is constructed as in (1), replacing the salience score with the corresponding firm-level characteristic averaged over the years the firm is observed.

Table A.1 — LIWC *Mental Health* and *Wellness* entries.

Panel A. Unigrams (N = 156)

*acupunct**, *addict*, *addicted*, *addicting*, *addiction*, *addictive*, *adhd*, *aerobics*, *agoraphobi**, *alcoholics*, *alcoholism*, *anhedon**, *anorexi**, *antidepressant**, *antioxidant**, *antipsychotic*, *antisocial*, *aromatherapy*, *asperger**, *aspie*, *aspies*, *atkins*, *autism*, *autistic*, *ayurveda*, *ayurvedic*, *biofeedback*, *bipolar*, *bpd*, *bulimi**, *calisthenics*, *cardio*, *cbt*, *chakra*, *chiropract**, *compulsive*, *cope*, *coped*, *copes*, *coping*, *crossfit*, *delusion**, *depressed*, *depression*, *depressive*, *detox**, *diet*, *dietary*, *dietetic**, *dietician**, *dieting*, *dietitian**, *diets*, *dsm*, *dysmorphi**, *fitbit**, *fitness**, *gym**, *hallucinat**, *heal*, *healed*, *healer**, *healing*, *healthful*, *healthily*, *healthy*, *histrionic*, *homeopath**, *hydrat**, *hydrotherapy*, *hypochondria**, *insanit**, *insomnia**, *internaliz**, *invigor**, *jogger**, *jogging*, *keto*, *ketogen**, *kleptomania**, *malinger**, *mania*, *manic*, *masseur*, *masseuse*, *meditat**, *mindfulness*, *multivitamin**, *munchausen*, *nami*, *neurodiver**, *neuroses*, *neurosis*, *neurotic*, *neuroticism*, *neurotypical*, *nutrient*, *nutrients*, *nutrition**, *nutritionist*, *nutritious*, *obsessive*, *ocd*, *organic**, *osteopath**, *paranoi**, *paranoid*, *phobi**, *pilates*, *prozac*, *psychiatric*, *psychiatrist*, *psychiatry*, *psychoanalysis*, *psychoanalyst*, *psychologist*, *psychopath*, *psychopathology*, *psychopathy*, *psychosis*, *psychosomatic*, *psychotherapist*, *psychotherapy*, *psychotic*, *ptsd*, *pyromania*, *pyromaniac*, *rejuvenat**, *ritalin*, *sane*, *sanity*, *sauna*, *schizo*, *schizo**, *scizophrenia*, *sertraline*, *sociopath**, *spa*, *spas*, *spiritual**, *ssri*, *suicid**, *supported*, *thorazine*, *trauma**, *trichotilloman**, *tw*, *vitality*, *vitamin**, *wellbeing*, *wellness*, *workahol**, *xanax*, *yoga*, *zen*, *zoloft*

Panel B. Bigrams (N = 81)

*alternative medicine**, *anti aging*, *anti depressant**, *anti-aging*, *anti-depressant**, *anxiety disorder*, *balanced diet*, *behav* therapy*, *binge eat**, *binge-eat**, *borderline personality*, *breathing exercise**, *chinese medicine*, *clinical depression*, *coconut water*, *conduct disorder*, *crisis intervention**, *day spa**, *disease prevention*, *eating disorder**, *essential oil**, *expressive writing*, *fish oil*, *gluten free*, *gluten-free*, *gmo free*, *gmo-free*, *herbal medicine**, *herbal tea**, *life coach*, *lift* weights*, *major depression*, *mental health*, *mental illness**, *mental institution**, *mentally ill*, *nervous breakdown*, *nervous disorder*, *non gmo*, *non-gmo*, *omega 3*, *panic attack**, *panic disorder**, *personality disorder**, *physical activity*, *physical fitness*, *physical therapy*, *positive mood*, *post traumatic*, *post-traumatic*, *postpartum depression*, *preventative care*, *preventative medicine*, *quit* smoking*, *self care*, *self compassion*, *self harm*, *self help*, *self improvement*, *self-care*, *self-compassion*, *self-harm*, *self-help*, *self-improvement*, *strength train**, *substance abuse*, *support group**, *sweat lodge**, *t'ai chi*, *tai chi*, *tai-chi*, *talk therapy*, *talk-therapy*, *weight gain*, *weight lifting*, *weight loss*, *well balanced*, *well being*, *well-balanced*, *well-being*, *whole grain*

Panel C. 3+ word phrases (N = 6)

attention deficit disorder, *attention deficit hyperactivity disorder*, *meaning in life*, *on the wagon*, *seasonal affective disorder*, *substance use disorder*

Notes. Entries are taken from the LIWC *Mental Health* and *Wellness* lists and shown as written in the source. Terms are alphabetically ordered within each panel.

Table A.2 — Example sentences from high- and low-salience annual reports.

Panel A. High-salience reports (top decile)
In other words to attribute people with special needs properties that make them self-help in everyday life.
Louis Poulsen A/S wants to create a healthy and desirable physical and psychological working environment with focus on the well-being of the employees including sickness absenteeism.
It is the management's view that both the physical and mental work environment is satisfying.
The Group focuses on a healthy and secure physical and mental work environment for all employees and professional entrepreneurs on the workplace.
The joy of the small, spontaneous dialogue with colleagues and by "investing" ourselves in meeting.
Panel B. Low-salience reports (bottom decile)
Any tax is calculated by the applicable corporate tax percentage of the difference between accounting and tax values.
The reason for the change in practice is that there is no requirement under the annual accounting law that the company must register leasing contracts as financial leasing.
The revenue requirement is determined taking into account the location of the property, age, maintenance status and lease rate, including the terms and conditions in lease contracts.
Differences between the exchange rates at the balance sheet date and the transaction date rates are recognized under financial income and expenses in the income statement.
When determining the income ratio of the property, the entering lease contracts, including the current lease and other relevant circumstances, are taken into account.

Notes. Each row displays a sentence from a sampled annual report. We randomly draw 1,000 annual reports from each of the top and bottom deciles of the salience score distribution. After applying standard text-quality filters to remove non-prose content such as addresses and headers, we extract the sentence with the highest cosine similarity to the well-being dictionary vector for reports in the top decile and the sentence with the lowest similarity for reports in the bottom decile. Each panel displays the five most extreme sentences across the sampled reports.

Table A.3 — Descriptive statistics for up- and down-movers.

	Down-movers	**Up-movers**	**N. diff.**	**Log SD**
Sample size				
N. unique workers	37,370	34,344	–	–
Outcomes				
Annual drug expenditure	34.35	32.34	-0.005	-0.081
Annual drug consumption	12.92	13.04	0.001	0.012
N. prescriptions	0.24	0.24	0.001	0.097
Any drug use	0.06	0.06	-0.004	-0.007
Hospitalization days (mental health)	0.0008	0.0006	-0.004	-0.835
Any hospitalization (mental health)	0.0004	0.0004	-0.002	-0.040
N. psychiatrist visits	0.08	0.08	-0.004	-0.070
Any psychiatrist visit	0.01	0.01	0.000	0.002
Hospitalization days (respiratory)	0.0097	0.0088	-0.003	-0.046
Any hospitalization (respiratory)	0.0034	0.0032	-0.003	-0.027
Hospitalization days (musculoskeletal)	0.0104	0.0102	-0.001	0.014
Any hospitalization (musculoskeletal)	0.0043	0.0045	0.003	0.024
Absence days	2.79	2.00	-0.087	-0.212
Any absence	0.31	0.24	-0.161	-0.082
Worker characteristics				
Age	41.61	41.25	-0.033	-0.001
Female	0.29	0.29	0.002	0.001
Danish	0.92	0.92	-0.022	0.034
Employment conditions				
Hourly wage (DKK)	245.24	241.08	-0.026	-0.138
Tenure (years)	5.89	5.72	-0.029	-0.015
Union member	0.71	0.70	-0.011	0.005
Same residence	0.44	0.44	-0.003	-0.000

Notes. "Down-movers" are workers whose salience differential is negative ($\delta_i < 0$), i.e. who move to a lower-salience firm; "Up-movers" are workers with $\delta_i > 0$. Means are computed over pre-move worker-years ($r < 0$) only. Hospitalization measures are computed over the years through 2018 for which the register is available. "N. diff." denotes the normalized difference in means; "Log SD" reports the logarithm of the ratio of standard deviations.

B Dictionary Validation and Refinement

This section documents the manual validation procedure and the resulting dictionary refinements summarized in Section 3.1.

Before finalizing the dictionary, we subject the candidate list to a manual validation step. A research assistant independently read a random sample of 50 translated reports drawn uniformly from the full corpus and flagged instances of poor machine translation by comparing each translated document against the original Danish source. We then reviewed each flagged instance and cross-checked it against the candidate dictionary entries.

This exercise identified one recurring mistranslation artifact. The Danish compound *kostpris*—a standard accounting term meaning *cost price*—was occasionally rendered by the M2M-100 translation model using English words that overlap with LIWC entries. This occurs because *kost* is polysemous: it carries the meaning *cost* in an accounting register but *diet* or *nutrition* in everyday usage, and the translation model resolves the ambiguity inconsistently. Consistent with this interpretation, the four affected terms—*diet*, *dietary*, *nutrition*, and *nutritional*—exhibit document frequencies that are implausibly high relative to the rest of the dictionary, confirming that their occurrences are driven by the artifact rather than genuine content. We therefore remove all four from the candidate list.

We additionally handle nine pure acronyms: *adhd*, *bpd*, *cbt*, *dsm*, *nami*, *ocd*, *ptsd*, *ssri*, and *tw*. Two considerations motivate this choice. First, machine translation models handle acronyms inconsistently—depending on context, a given acronym may be passed through unchanged, expanded into its full form, or silently dropped—making such terms unreliable signals of the underlying concept in translated text. Second, short letter sequences are inherently prone to false positives in financial documents, where they routinely serve as abbreviations for legal entities, financial instruments, and regulatory bodies unrelated to mental health. Rather than discarding the underlying concepts entirely, we substitute each acronym with its full form, preserving the original conceptual coverage of the LIWC lists without introducing any arbitrary deviation from the source vocabulary. Note that *attention deficit hyperactivity disorder* was already present in full form in the LIWC lists and is therefore not added again.

After these operations, the final dictionary comprises 293 terms. Figure 1 in Section 3.1 visualizes the final dictionary as a word cloud, while Figure A.3 in Appendix A displays document frequencies for dictionary words.

C Robustness

This appendix collects the robustness analyses referenced in the main text. Section C.1 examines the cross-sectional associations of Section 4.3, considering heterogeneity by workforce gender composition and sensitivity to alternative constructions of the salience score. Section C.2 turns to the mover analysis of Section 5, reporting sensitivity to alternative sample constructions, a placebo test, and checks that condition on origin–destination differentials in firm characteristics.

C.1 Cross-Sectional Analysis

Table C.1 — Salience and worker health, by gender composition of firms.

	Male-dominated firms	Female-dominated firms
Panel A: Prescription drug use		
Avg. drug exp.	10.60 (9.452)	56.58*** (16.04)
Avg. drug cons.	5.060** (2.068)	9.537*** (3.654)
Avg. # prescriptions	0.0878** (0.0375)	0.2026*** (0.0655)
Share drug users	0.0106** (0.0051)	0.0093 (0.0082)
Any drug user	0.3903*** (0.0390)	0.3604*** (0.0452)
Panel B: Mental-health hospitalizations and specialist visits		
Avg. hosp. days	0.0017 (0.0014)	0.0024* (0.0014)
Share hosp.	0.0012** (0.0006)	0.0010 (0.0007)
Any hospitalization	0.0936*** (0.0201)	0.0850*** (0.0232)
Avg. psych visits	0.0064 (0.0167)	-0.0690** (0.0273)
Share psych users	-0.0001 (0.0020)	-0.0132*** (0.0036)
Any psych visit	0.6072*** (0.0518)	0.3315*** (0.0614)
Panel C: Respiratory and musculoskeletal hospitalizations		
Avg. hosp. days (respiratory)	-0.0043 (0.0050)	0.0125 (0.0080)
Share hosp. (respiratory)	-0.0013 (0.0011)	0.0021 (0.0014)
Any hospitalization (respiratory)	0.2606*** (0.0384)	0.3287*** (0.0454)
Avg. hosp. days (musculoskeletal)	0.0013 (0.0055)	0.0019 (0.0081)
Share hosp. (musculoskeletal)	0.0002 (0.0013)	0.0028* (0.0016)
Any hospitalization (musculoskeletal)	0.3452*** (0.0412)	0.3217*** (0.0463)
Panel D: Sickness absence		
Avg. absence days	2.444*** (0.2476)	1.613*** (0.2276)
Share absent	0.2527*** (0.0248)	0.1596*** (0.0219)
Any absence	0.5012*** (0.0490)	0.1357** (0.0558)
Observations	124,749	47,228
Observations (hosp.)	71,988	26,868
Industry FEs	✓	✓
Year FEs	✓	✓
Firm controls	✓	✓

Notes. Each cell reports the estimated coefficient on the salience score from a separate regression of the indicated firm-level outcome on the salience score and a common set of firm-level controls (organizational structure, workforce composition, and employment conditions), including industry and year fixed effects. The specification is identical to that of Tables 5–8, estimated separately on each subsample. The hospitalization register is available through 2018; "Observations (hosp.)" gives the corresponding count for the hospitalization rows of Panels B and C. Male-dominated firms are those with a share of female employees below one half; female-dominated firms are those above one half. Standard errors clustered at the firm level are reported in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table C.2 — Alternative salience score definitions and worker health.

	Median	75th-percentile	95th-percentile	Maximum
Panel A: Prescription drug use				
Avg. drug exp.	99.68*** (25.35)	46.47*** (17.98)	25.58*** (8.60)	9.36* (5.39)
Avg. drug cons.	13.83*** (5.36)	10.03** (4.22)	6.43*** (1.91)	1.17 (1.03)
Avg. # prescriptions	0.298*** (0.091)	0.175** (0.071)	0.122*** (0.035)	0.029 (0.019)
Share drug users	0.017 (0.014)	0.012 (0.010)	0.008* (0.005)	0.000 (0.003)
Any drug user	1.349*** (0.110)	0.986*** (0.077)	0.388*** (0.031)	0.234*** (0.017)
Panel B: Mental-health hospitalizations and specialist visits				
Avg. hosp. days	0.004 (0.003)	0.003 (0.002)	0.002* (0.001)	0.001 (0.001)
Share hosp.	0.002* (0.001)	0.002** (0.001)	0.001** (0.000)	0.000 (0.000)
Any hospitalization	0.434*** (0.081)	0.311*** (0.057)	0.095*** (0.016)	0.048*** (0.009)
Avg. psych visits	-0.015 (0.044)	-0.001 (0.037)	-0.024 (0.014)	-0.002 (0.008)
Share psych users	-0.005 (0.006)	-0.001 (0.004)	-0.005*** (0.002)	-0.001 (0.001)
Any psych visit	2.704*** (0.182)	1.970*** (0.130)	0.529*** (0.042)	0.341*** (0.023)
Panel C: Respiratory and musculoskeletal hospitalizations				
Avg. hosp. days (respiratory)	0.021 (0.016)	0.014 (0.013)	0.002 (0.004)	-0.007*** (0.003)
Share hosp. (respiratory)	0.002 (0.003)	0.002 (0.002)	-0.000 (0.001)	-0.001 (0.001)
Any hospitalization (respiratory)	1.510*** (0.161)	1.096*** (0.108)	0.294*** (0.030)	0.179*** (0.017)
Avg. hosp. days (musculoskeletal)	0.015 (0.013)	0.005 (0.009)	0.002 (0.005)	-0.002 (0.003)
Share hosp. (musculoskeletal)	0.006 (0.003)	0.003 (0.002)	0.001 (0.001)	-0.001 (0.001)
Any hospitalization (musculoskeletal)	1.821*** (0.166)	1.284*** (0.112)	0.351*** (0.032)	0.200*** (0.018)
Panel D: Sickness absence				
Avg. absence days	12.32*** (0.98)	8.86*** (0.68)	2.23*** (0.19)	1.33*** (0.11)
Share absent	1.283*** (0.097)	0.899*** (0.067)	0.227*** (0.019)	0.134*** (0.010)
Any absence	2.092*** (0.173)	1.524*** (0.122)	0.391*** (0.040)	0.285*** (0.022)
Observations	171,977	171,977	171,977	171,977
Observations (hosp.)	98,856	98,856	98,856	98,856
Industry FEs	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓
Firm controls	✓	✓	✓	✓

Notes. Each cell reports the estimated coefficient on the salience score from a separate regression of the indicated firm-level outcome on the salience score and a common set of firm-level controls (organizational structure, workforce composition, and employment conditions), including industry and year fixed effects. The specification is identical to that of Tables 5–8; the columns differ only in the summary statistic used to construct the salience score, namely the median, the 75th percentile, the 95th percentile (our baseline), and the maximum of the sentence-level similarity distribution. The hospitalization register is available through 2018, so the hospitalization rows of Panels B and C are estimated on the observation count reported in "Observations (hosp.)"; all other rows use the full sample. Standard errors clustered at the firm level are reported in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table C.3 — Report length and worker health.

Panel A: Prescription drug use	
Avg. drug exp.	-0.017** (0.008)
Avg. drug cons.	-0.007*** (0.002)
Avg. # prescriptions	-0.0001*** (0.00003)
Share drug users	-0.00002*** (0.000004)
Any drug user	0.001*** (0.00004)
Panel B: Mental-health hospitalizations and specialist visits	
Avg. hosp. days	-0.0000 (0.0000)
Share hosp.	-0.0000* (0.0000)
Any hospitalization	0.0003*** (0.00005)
Avg. psych visits	-0.00005*** (0.00001)
Share psych users	-0.00001*** (0.000002)
Any psych visit	0.002*** (0.00008)
Panel C: Respiratory and musculoskeletal hospitalizations	
Avg. hosp. days (respiratory)	-0.00001** (0.000005)
Share hosp. (respiratory)	-0.000002** (0.0000008)
Any hospitalization (respiratory)	0.001*** (0.0001)
Avg. hosp. days (musculoskeletal)	-0.00001** (0.000005)
Share hosp. (musculoskeletal)	-0.000003*** (0.0000009)
Any hospitalization (musculoskeletal)	0.001*** (0.0001)
Panel D: Sickness absence	
Avg. absence days	0.006*** (0.0004)
Share absent	0.001*** (0.00004)
Any absence	0.002*** (0.00008)
Observations	171,977
Observations (hosp.)	98,856
Industry FEs	✓
Year FEs	✓
Firm controls	✓

Notes. Each cell reports the estimated coefficient on report length from a separate regression of the indicated firm-level outcome on report length and a common set of firm-level controls (organizational structure, workforce composition, and employment conditions), including industry and year fixed effects. The specification is identical to that of Tables 5–8, with the salience score replaced by the number of sentences in a firm's annual report. The hospitalization register is available through 2018, so the hospitalization rows of Panels B and C are estimated on the observation count reported in "Observations (hosp.)"; all other rows use the full sample. Standard errors clustered at the firm level are reported in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

C.2 Mover Analysis

Sensitivity to the tenure window. Figure C.1 reports the post-move estimate $\hat{\theta}_1$ as the minimum tenure required at the origin and destination firm varies from two to four years. The point estimates are remarkably stable across the range for both outcomes, with confidence intervals that widen as the window lengthens—an expected consequence of the smaller sample the stricter requirement implies. Our results do not depend on the particular window used to define the mover sample.

FIGURE C.1 — Sensitivity of the absence response to the tenure window.

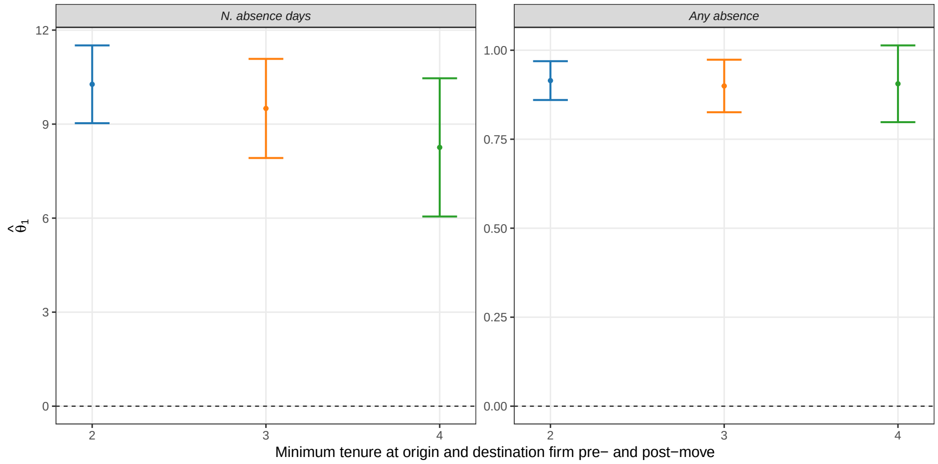

Notes. Each point is the post-move estimate $\hat{\theta}_1$ from (2), estimated on a separate mover sample defined by the minimum tenure required at both the origin and destination firm, which we vary from two to four years. The outcome is the number of absence days in the left panel and an indicator for any absence in the right panel. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Vertical bars denote 95% confidence intervals, with standard errors clustered at the worker level.

Pre-pandemic sample. Figure C.2 re-estimates the event study on movers whose entire event window falls before 2020. Doing so under our baseline three-year tenure requirement would confine moves to 2016, leaving too few workers for estimation. We therefore shorten the window to two years on each side of the move, which admits move years from 2015 to 2017 and retains enough movers to estimate the response. The point estimates track those from the full sample for both outcomes. The absence response is therefore not an artifact of the pandemic period.

Figure C.2 — Event study for absence outcomes, pre-pandemic sample.

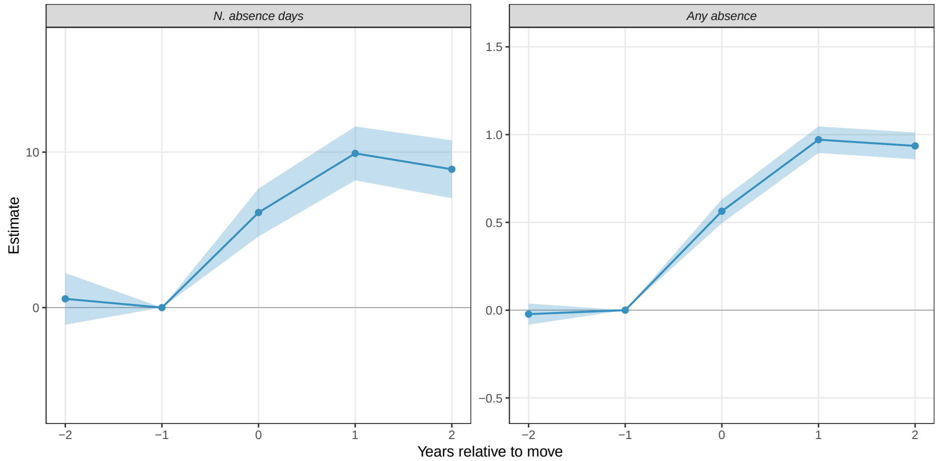

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, restricting the sample to movers whose entire event window falls before 2020. To keep a sufficient number of movers within this period, we require two years of observation before and after the move rather than three. The outcome is the number of absence days in the left panel and an indicator for any absence in the right panel. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Workers without children. Figure C.3 re-estimates the event study restricting the sample to workers with no children at home throughout the event window. The estimates closely match those from the full sample for both outcomes, confirming that the absence response reflects workers' own absence rather than care provided to others.

Figure C.3 — Event study for absence outcomes, workers without children.

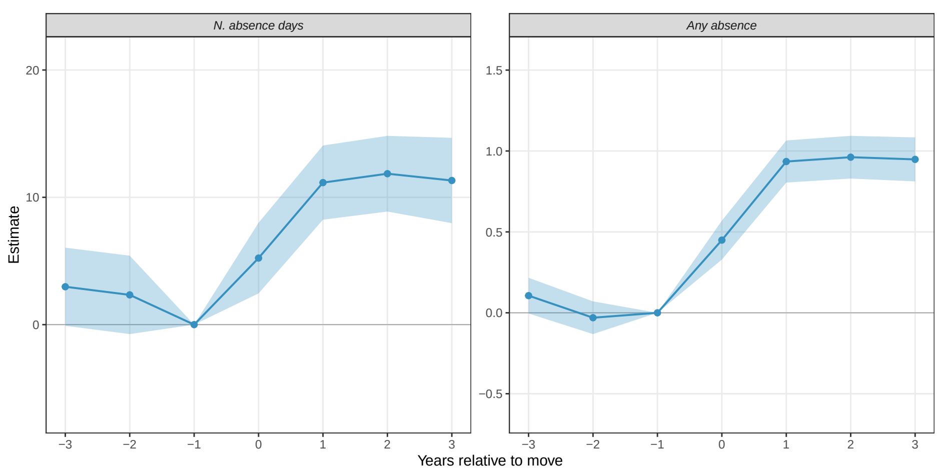

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, restricting the sample to workers with no children at home throughout the event window. The outcome is the number of absence days in the left panel and an indicator for any absence in the right panel. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Placebo test. Figure C.4 re-estimates the event study after replacing each mover's true salience differential with a random draw from the empirical distribution of δ_i, holding the move date fixed. Under this assignment, a worker's salience differential bears no relation to the firms they actually moved between, so any genuine workplace response should disappear. The resulting estimates are flat and close to zero throughout for both absence outcomes, confirming that the response we estimate reflects the actual salience gap between origin and destination rather than an artifact of the event-study specification or the composition of movers.

FIGURE C.4 — Placebo test: randomly assigned salience differentials.

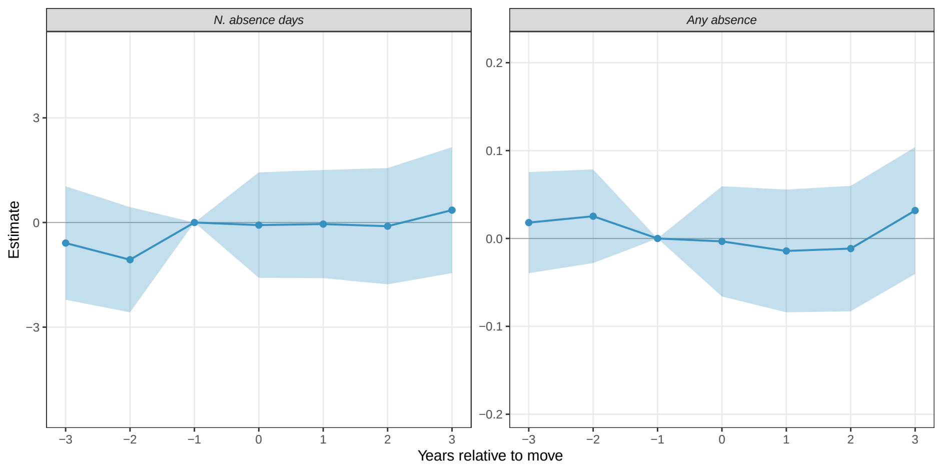

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2) after replacing each mover's salience differential δ_i with a random draw from its empirical distribution, holding move dates fixed. The left panel reports the number of absence days; the right panel, the indicator for any absence. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Pre-move salience differentials. We construct an alternative salience differential $\tilde{\delta}_i$ as in (1), but using only reports filed between 2013 and 2015. Because the mover sample requires three years of observation at the origin before the move and three at the destination afterwards, all moves occur between 2016 and 2019, so this window precedes every move by construction and $\tilde{\delta}_i$ is pre-move for all movers. Figures C.5–C.8 re-estimate (2) with $\tilde{\delta}_i$ in place of δ_i and find the same pattern for every outcome, though the absence response is attenuated. This is expected: $\tilde{\delta}_i$ is constructed from at most three reports per firm rather than the full window, and therefore measures salience with more noise.

Figure C.5 — Event study for sickness absence, pre-move salience differentials.

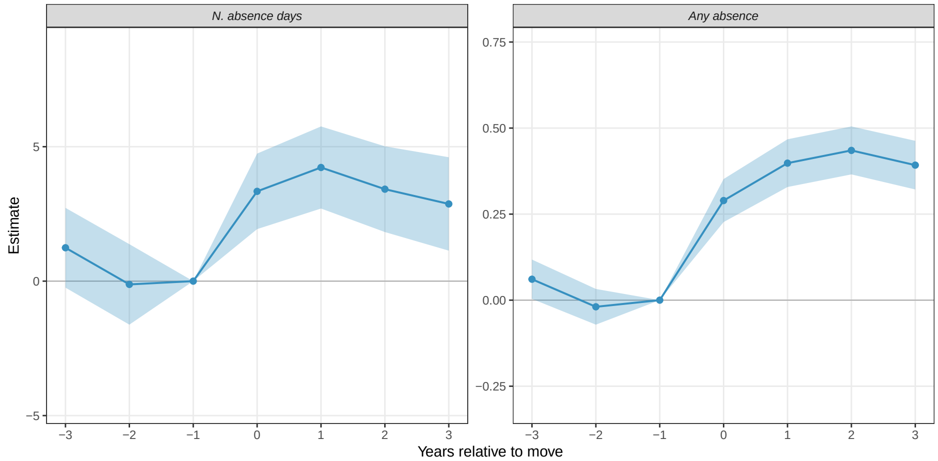

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, replacing δ_i with the pre-move differential $\tilde{\delta}_i$ constructed from reports filed in 2013–2015 only. The outcome is the number of absence days in the left panel and an indicator for any absence in the right panel. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Figure C.6 — Event study for prescription drug use, pre-move salience differentials.

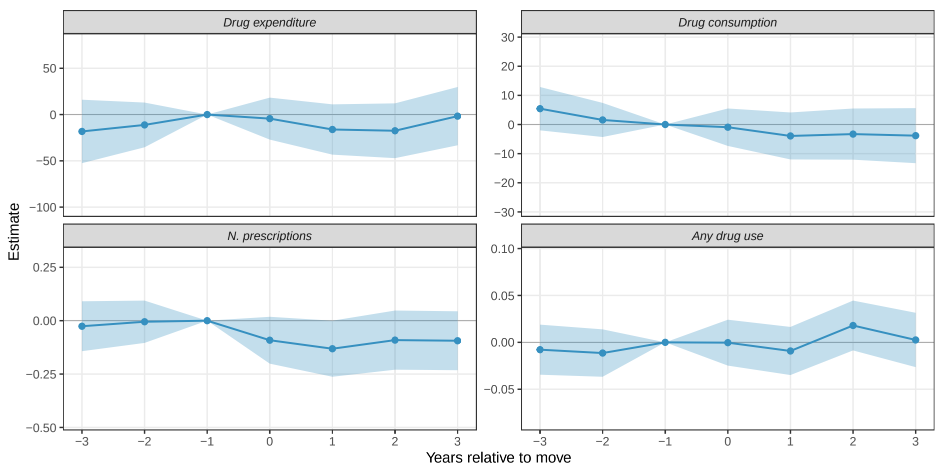

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, replacing δ_i with the pre-move differential $\tilde{\delta}_i$ constructed from reports filed in 2013–2015 only, for four measures of prescription drug use: annual drug expenditure, annual drug consumption, the number of prescriptions, and an indicator for any drug use. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Figure C.7 — Event study for mental-health hospitalizations and specialist visits, pre-move salience differentials.

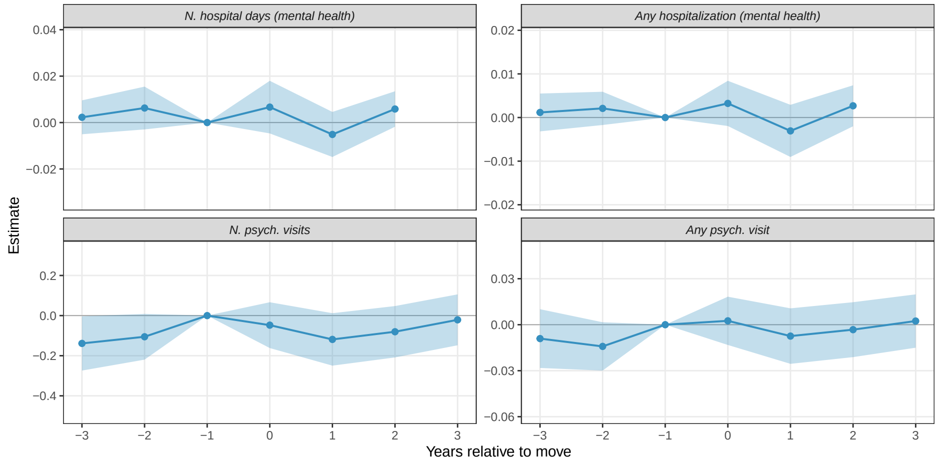

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, replacing δ_i with the pre-move differential $\tilde{\delta}_i$ constructed from reports filed in 2013–2015 only, for mental-health hospitalizations (the number of hospital days and an indicator for any hospitalization) and psychiatric and psychological visits (the number of visits and an indicator for any visit). The hospitalization register is available only through 2018, so the hospitalization estimates are reported up to $r = 2$; the visit outcomes run through $r = 3$. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Figure C.8 — Event study for respiratory and musculoskeletal hospitalizations, pre-move salience differentials.

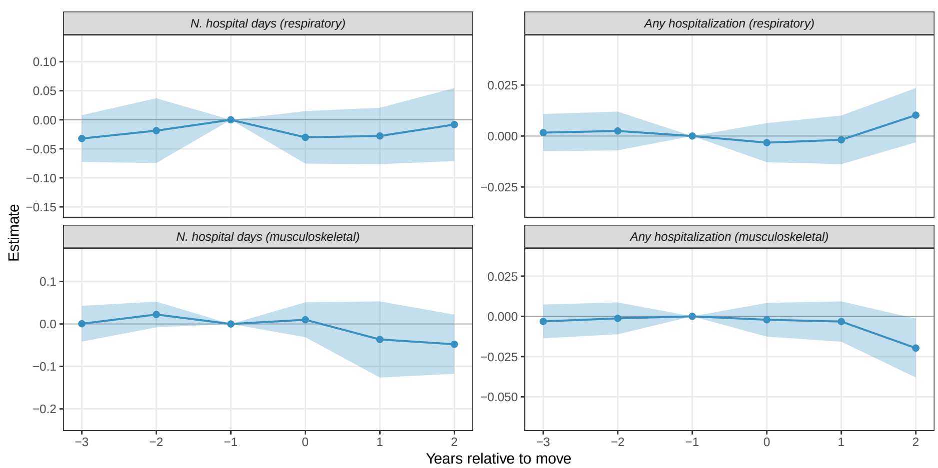

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, replacing δ_i with the pre-move differential $\tilde{\delta}_i$ constructed from reports filed in 2013–2015 only, for hospitalizations attributable to respiratory and musculoskeletal conditions (the number of hospital days and an indicator for any hospitalization). The hospitalization register is available only through 2018, so estimates are reported up to $r = 2$. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Alternative clustering. We re-estimate (2) clustering standard errors at the origin–destination firm pair, at the origin firm, at the destination firm, and two-way on origin and destination. The choice affects only the standard errors, leaving point estimates unchanged. Figure C.9 reports the results. Confidence intervals widen, but every post-move estimate remains significant, with the exception of the move year on the extensive margin.

Figure C.9 — Event study for sickness absence, alternative clustering.

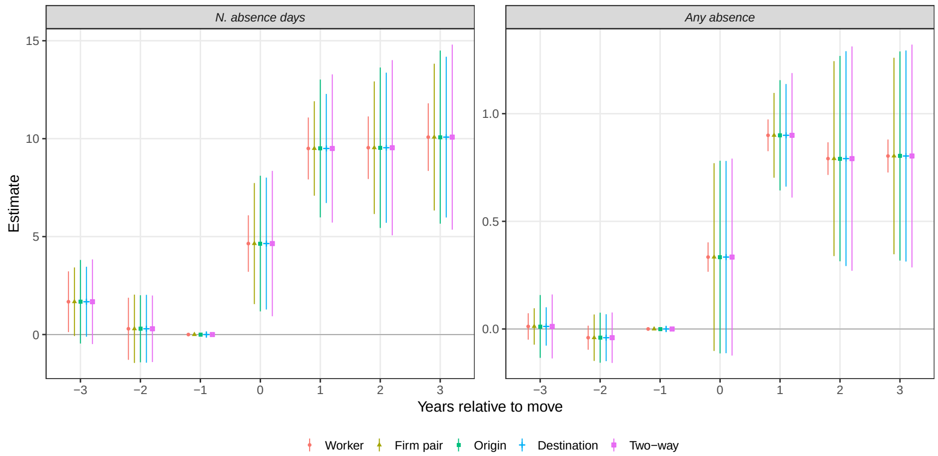

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, under five clustering schemes: the worker (our baseline), the origin–destination firm pair, the origin firm, the destination firm, and two-way on origin and destination firm. The outcome is the number of absence days in the left panel and an indicator for any absence in the right panel. Point estimates are identical across schemes; only the confidence intervals differ. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Vertical bars denote 95% confidence intervals.

Pre-move absence. Down-movers begin from higher baseline absence than up-movers (Table A.3). We therefore control for it directly, augmenting (2) with each worker's average number of absence days over the pre-move years, interacted with event time. Figure C.10 reports the results. The extensive margin is essentially unchanged, while the response in absence days attenuates but follows the same pattern.

FIGURE C.10 — Event study for sickness absence, controlling for pre-move absence.

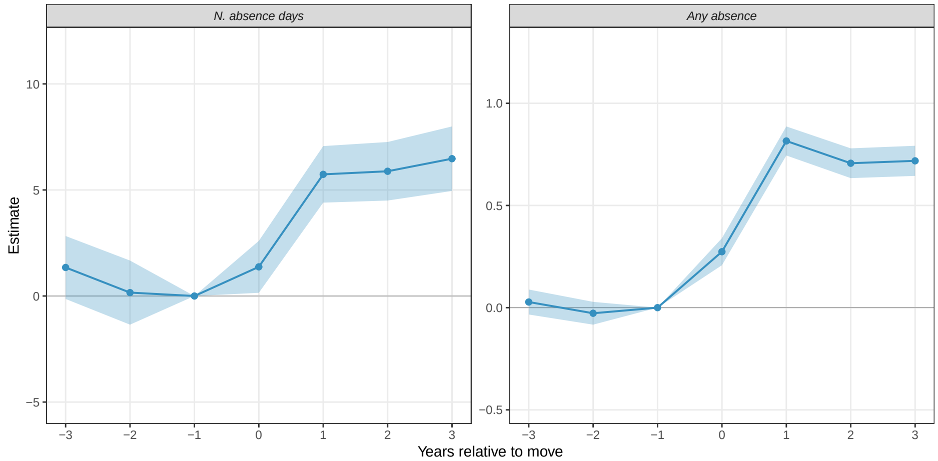

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, adding each worker's average number of absence days over the pre-move years, interacted with event time, as an additional control. The outcome is the number of absence days in the left panel and an indicator for any absence in the right panel. The specification includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator. Standard errors are clustered at the worker level; shaded bands denote 95% confidence intervals.

Conditioning on firm-characteristic differentials. A worker who moves to a higher-salience firm may simultaneously move to a larger, better-paying, or compositionally different one, in which case the salience differential would partly stand in for these other changes. Figure A.13 in Appendix A shows that this is not the case: for each firm characteristic, we construct the origin–destination differential as in (1), replacing the salience score with the characteristic, and find all correlations with the salience differential to be negligible. Consistent with this, Figures C.11–C.14 re-estimate (2) augmented with these differentials, each interacted with event time, and find estimates essentially indistinguishable from the baseline for every outcome.

FIGURE C.11 — Event study for absence outcomes, with additional controls.

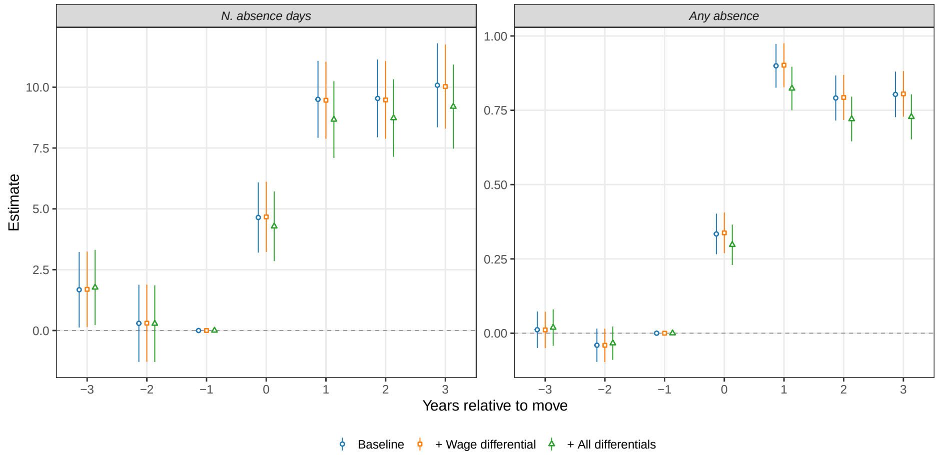

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, across three specifications. The baseline includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator, as in Figure 6. The second specification adds the origin–destination differential in mean hourly wages, interacted with event time; the third adds the analogous differentials in all the firm characteristics shown in Figure A.13. The outcome is the number of absence days in the left panel and an indicator for any absence in the right panel. Standard errors are clustered at the worker level; vertical bars denote 95% confidence intervals.

FIGURE C.12 — Event study for prescription drug outcomes, with additional controls.

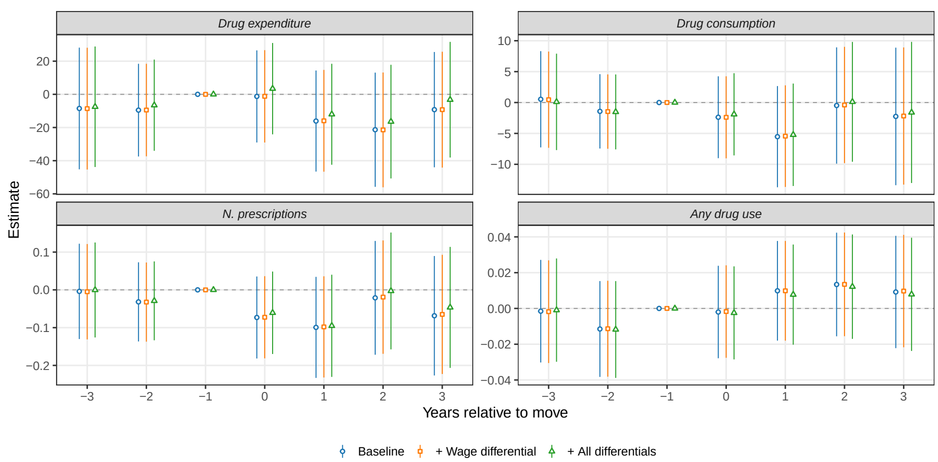

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, across three specifications, for four measures of prescription drug use: annual drug expenditure, annual drug consumption, the number of prescriptions, and an indicator for any drug use. The baseline includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator, as in Figure 7. The second specification adds the origin–destination differential in mean hourly wages, interacted with event time; the third adds the analogous differentials in all the firm characteristics shown in Figure A.13. Standard errors are clustered at the worker level; vertical bars denote 95% confidence intervals.

Figure C.13 — Event study for mental-health hospitalizations and specialist visits, with additional controls.

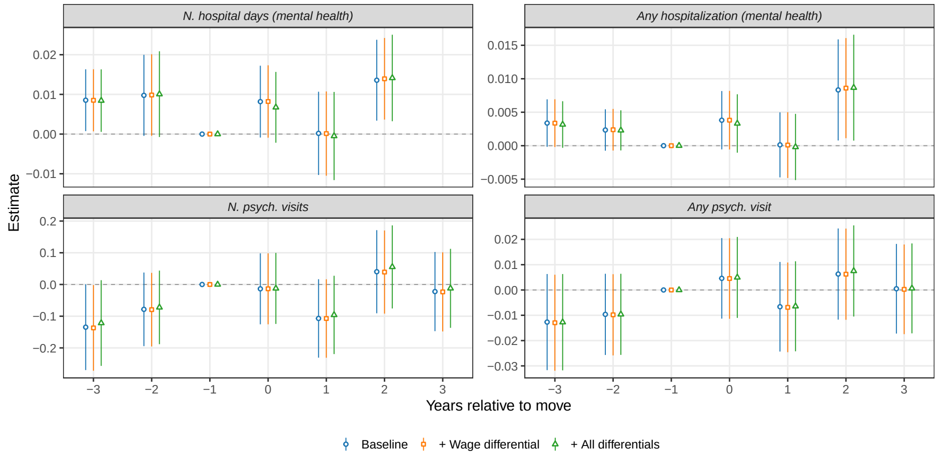

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, across three specifications, for mental-health hospitalizations (the number of hospital days and an indicator for any hospitalization) and psychiatric and psychological visits (the number of visits and an indicator for any visit). The hospitalization register is available only through 2018, so the hospitalization estimates are reported up to $r = 2$; the visit outcomes run through $r = 3$. The baseline includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator, as in Figure 8. The second specification adds the origin–destination differential in mean hourly wages, interacted with event time; the third adds the analogous differentials in all the firm characteristics shown in Figure A.13. Standard errors are clustered at the worker level; vertical bars denote 95% confidence intervals.

Figure C.14 — Event study for respiratory and musculoskeletal hospitalizations, with additional controls.

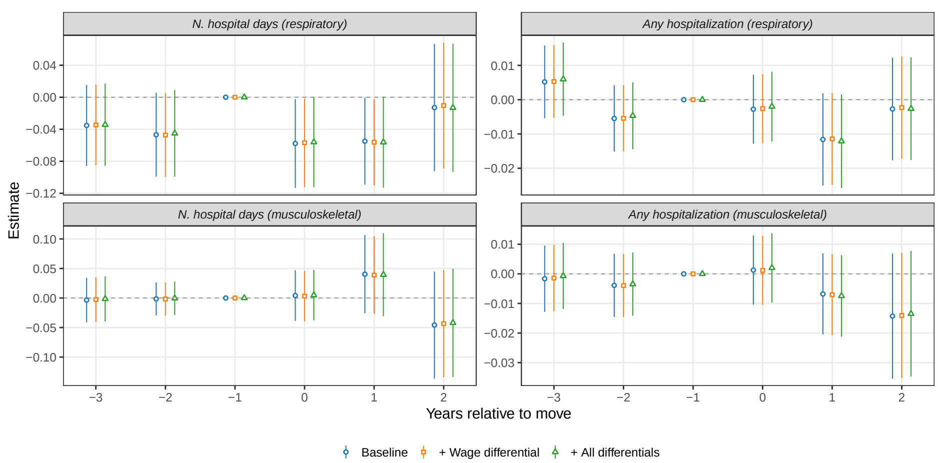

Notes. The figure plots the estimates $\hat{\theta}_r$ from (2), with $r = -1$ normalized to zero, across three specifications, for hospitalizations attributable to respiratory and musculoskeletal conditions (the number of hospital days and an indicator for any hospitalization). The hospitalization register is available only through 2018, so estimates are reported up to $r = 2$. The baseline includes worker, year, age-by-year, and industry fixed effects and a blue-collar indicator, as in Figure 9. The second specification adds the origin–destination differential in mean hourly wages, interacted with event time; the third adds the analogous differentials in all the firm characteristics shown in Figure A.13. Standard errors are clustered at the worker level; vertical bars denote 95% confidence intervals.

D Stayer Analysis

A worker's exposure to well-being salience can change in only two ways: the worker changes firm, or the firm's salience changes around a worker who stays. The mover analysis of Section 5 exploits the first source of variation, and supports a causal interpretation under an identifying assumption—that the timing of moves is unrelated to contemporaneous changes in health—for which we provide direct evidence. This appendix turns to the second source.

We focus on workers who never change firms during the sample window, whom we call *stayers*, and test whether changes in a firm's salience are associated with changes in its workers' absence. The sample comprises 1,238,369 such workers, observed over 4,910,784 worker-years. The exercise is descriptive—within-firm changes in salience may themselves respond to workforce conditions—but it speaks directly to whether the absence response of Section 5 is driven by the firms or by the workers: among stayers, any variation in the salience a worker faces originates entirely with the firm, so a firm-driven response should leave a trace in this sample.

We regress each outcome on indicators for the quartile of the employer's salience score, computed annually over the full population of firms, together with worker and year fixed effects. We use quartile indicators rather than the continuous score for two reasons: within-firm variation in the raw score partly reflects transitory fluctuations in reporting language, and the indicators respond only to changes large enough to move a firm across the annual distribution, filtering out this noise; in addition, they impose no functional form on the relationship between salience and absence. Because each stayer is attached to a single firm, the worker fixed effects absorb both worker heterogeneity and all time-invariant firm heterogeneity, so the coefficients are identified from firms moving across the salience distribution over time.

Figure D.1 reports the estimates for the absence outcomes. Consistent with a firm-driven response, the estimates for absence days rise monotonically across the salience quartiles, attaining statistical significance in the top quartile. The extensive margin displays a similar though weaker pattern, with positive estimates in the upper half of the distribution that remain statistically indistinguishable from zero.

Figures D.2–D.4 repeat the exercise for the medical outcomes. In contrast to absence, no gradient emerges: the estimates are small, display no ordering across the salience quartiles, and are statistically indistinguishable from zero throughout.

The stayers therefore reproduce, using a different source of variation, the central pattern of the mover analysis: absence moves with salience while the medical outcomes do not. Because the two samples share neither workers nor the source of variation in salience, the same pattern emerging in both is hard to attribute to anything about the workers, and points to the firms as the drivers of the absence response.

Figure D.1 — Stayer analysis for sickness absence.

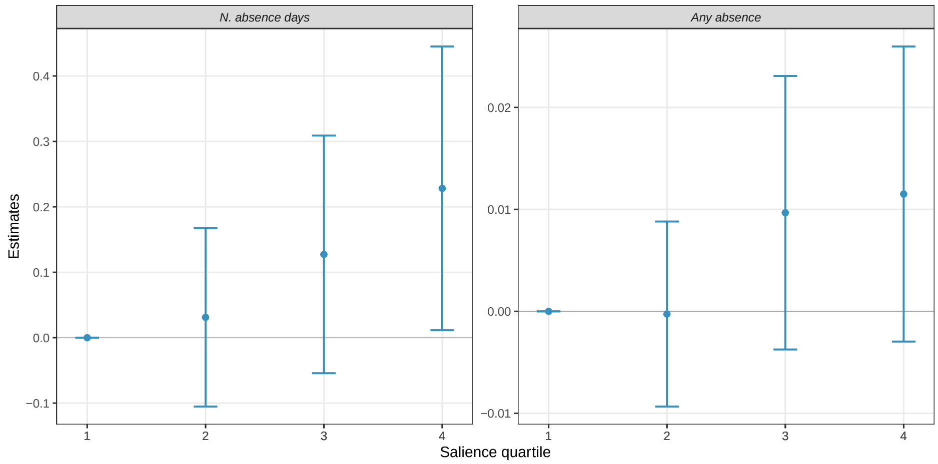

Notes. The figure plots the estimated coefficients on the salience-quartile indicators from regressions on the sample of stayers, with the bottom quartile as the reference category. The outcome is the number of absence days in the left panel and an indicator for any absence in the right panel. The specification includes worker and year fixed effects. Salience quartiles are computed annually on the full population of firm-years. Standard errors are clustered at the firm level; vertical bars denote 95% confidence intervals.

Figure D.2 — Stayer analysis for prescription drug use.

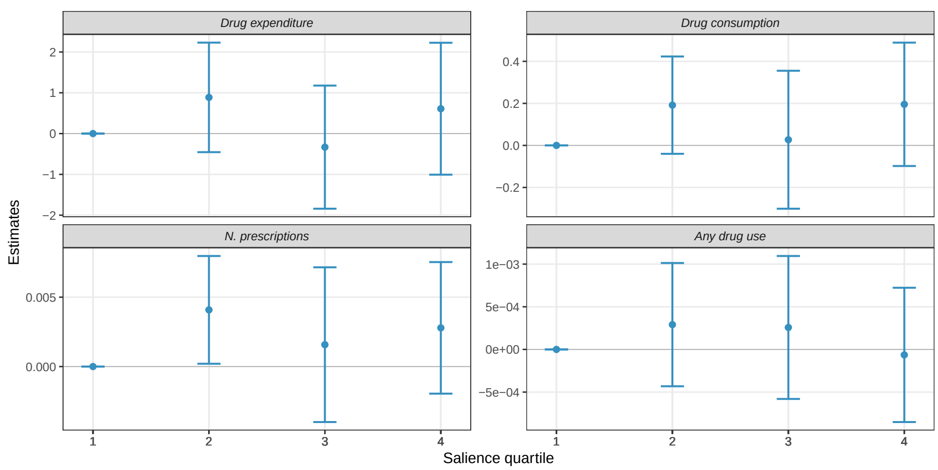

Notes. The figure plots the estimated coefficients on the salience-quartile indicators from regressions on the sample of stayers, with the bottom quartile as the reference category. The outcomes are annual drug expenditure, annual drug consumption, the number of prescriptions, and an indicator for any drug use. The specification includes worker and year fixed effects. Salience quartiles are computed annually on the full population of firm-years. Standard errors are clustered at the firm level; vertical bars denote 95% confidence intervals.

FIGURE D.3 — Stayer analysis for mental-health hospitalizations and specialist visits.

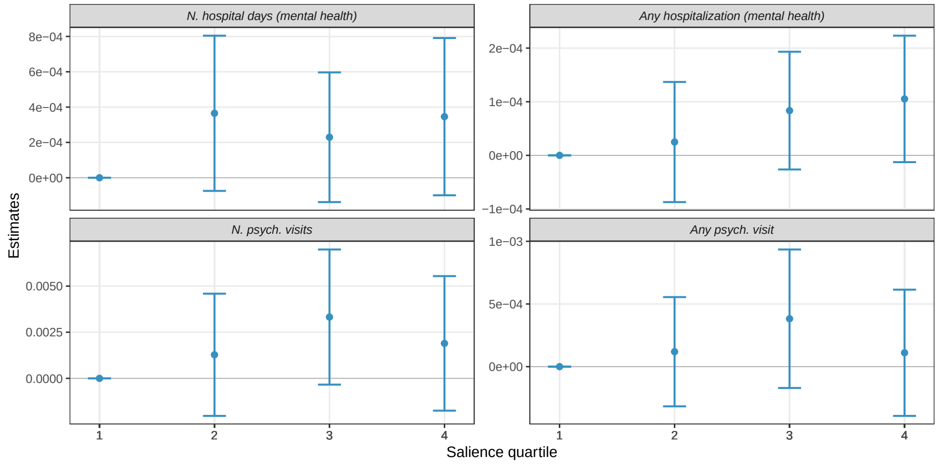

Notes. The figure plots the estimated coefficients on the salience-quartile indicators from regressions on the sample of stayers, with the bottom quartile as the reference category. The outcomes measure mental-health hospitalizations (the number of mental-health hospital days and an indicator for any hospitalization), and psychiatric and psychological visits (the number of visits and an indicator for any visit). The specification includes worker and year fixed effects. Salience quartiles are computed annually on the full population of firm-years. Standard errors are clustered at the firm level; vertical bars denote 95% confidence intervals.

FIGURE D.4 — Stayer analysis for respiratory and musculoskeletal hospitalizations.

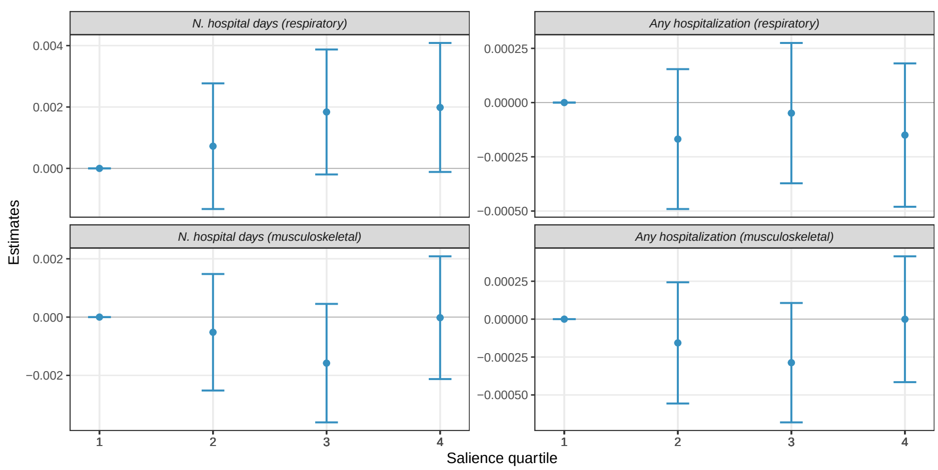

Notes. The figure plots the estimated coefficients on the salience-quartile indicators from regressions on the sample of stayers, with the bottom quartile as the reference category. The outcomes measure hospitalizations attributable to respiratory and musculoskeletal conditions (the number of hospital days and an indicator for any hospitalization). The specification includes worker and year fixed effects. Salience quartiles are computed annually on the full population of firm-years. Standard errors are clustered at the firm level; vertical bars denote 95% confidence intervals.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999): "High wage workers and high wage firms," *Econometrica*, 67, 251–333.

Ahammer, A., A. Packham, and J. Smith (2024): "Firms and worker health," *National Bureau of Economic Research.*

Amore, M. D., B. Larsen, M. Bennedsen, and Z. Zhao (2025): "The Economics of Working Environments: Evidence from Linked Survey and Register Data," *Available at SSRN 5562518.*

Ash, E. and S. Hansen (2023): "Text algorithms in economics," *Annual Review of Economics*, 15, 659–688.

Bach, L., R. Baghai, M. Bos, and R. Silva (2021): "How do acquisitions affect the mental health of employees?" Research Paper 21-21, Swedish House of Finance.

Badinski, I., A. Finkelstein, M. Gentzkow, and P. Hull (2023): "Geographic variation in healthcare utilization: The role of physicians," Tech. rep., National Bureau of Economic Research.

Barry, M.-L., B. J. Bruns, S. Kandemir, J. Klose, V. Smirnov, and P. Tillmann (2025): "The Emotions of Monetary Policy," *Available at SSRN 5099686.*

Belloni, M., L. Carrino, and E. Meschi (2022): "The impact of working conditions on mental health: Novel evidence from the UK," *Labour Economics*, 76, 102176.

Bonhomme, S., T. Lamadon, and E. Manresa (2026): "A Users' Guide to Uncovering Worker and Firm Effects: The ABC of AKM," *Journal of Economic Perspectives*, 40, 195–214.

Boyd, R. L., A. Ashokkumar, S. Seraj, and J. W. Pennebaker (2022): "The development and psychometric properties of LIWC-22," *Austin, TX: University of Texas at Austin*, 10, 1–47.

Burn, I., P. Button, L. M. Corella, and D. Neumark (2022): "Does ageist language in job ads predict age discrimination in hiring?" *Journal of Labor Economics*, 40, 613–667.

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

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

Clark, S. (2015): "Vector space models of lexical meaning," *The Handbook of Contemporary Semantic Theory*, 493–522.

De Vries, E., M. Schoonvelde, and G. Schumacher (2018): "No longer lost in translation: Evidence that Google Translate works for comparative bag-of-words text applications," *Political Analysis*, 26, 417–430.

Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova (2019): "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 conference of the North American chapter of the Association for Computational Linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.

Fan, A., S. Bhosale, H. Schwenk, Z. Ma, A. El-Kishky, S. Goyal, M. Baines, O. Celebi, G. Wenzek, V. Chaudhary, et al. (2021): "Beyond english-centric multilingual machine translation," *Journal of Machine Learning Research*, 22, 1–48.

Finkelstein, A., M. Gentzkow, and H. Williams (2016): "Sources of geographic variation in health care: Evidence from patient migration," *The quarterly journal of economics*, 131, 1681–1726.

Gennaro, G. and E. Ash (2022): "Emotion and reason in political language," *The Economic Journal*, 132, 1037–1059.

Gorodnichenko, Y., T. Pham, and O. Talavera (2023): "The voice of monetary policy," *American Economic Review*, 113, 548–584.

Hansen, S., M. McMahon, and A. Prat (2018): "Transparency and deliberation within the FOMC: A computational linguistics approach," *The Quarterly Journal of Economics*, 133, 801–870.

Hoff, N. v., G. Mellace, and S. Menon (2025): "Gender Differences in Healthcare Utilisation–Evidence from Unexpected Adverse Health Shocks," *arXiv preprint arXiv:2509.01310*.

Hollingworth, S. and T. Kairuz (2021): "Measuring medicine use: Applying ATC/DDD methodology to real-world data," *Pharmacy*, 9, 60.

Koschnick, J. (2025): "Teacher-directed scientific change: The case of the English Scientific Revolution," Tech. rep., EHES Working Paper.

Kosmidis, S., S. B. Hobolt, E. Molloy, and S. Whitefield (2019): "Party competition and emotive rhetoric," *Comparative Political Studies*, 52, 811–837.

Kárpáti, D. and L. Renneboog (2024): "Corporate Financial Frictions and Employee Mental Health," *Journal of Financial and Quantitative Analysis*, 59, 2256–2298.

Licht, H., R. Sczepanski, M. Laurer, and A. Bekmuratovna (2024): "No more cost in translation: Validating open-source machine translation for quantitative text analysis," Tech. rep., ECONtribute Discussion Paper.

Lucas, C., R. A. Nielsen, M. E. Roberts, B. M. Stewart, A. Storer, and D. Tingley (2015): "Computer-assisted text analysis for comparative politics," *Political Analysis*, 23, 254–277.

Madsen, I. E., F. Diderichsen, H. Burr, and R. Rugulies (2010): "Person-related work and incident use of antidepressants: relations and mediating factors from the Danish work environment cohort study," *Scandinavian Journal of Work, Environment & Health*, 435–444.

Marmot, M. G., H. Bosma, H. Hemingway, E. Brunner, and S. Stansfeld (1997): "Contribution of job control and other risk factors to social variations in coronary heart disease incidence," *The Lancet*, 350, 235–239.

Marmot, M. G., G. D. Smith, S. Stansfeld, C. Patel, F. North, J. Head, I. White, E. Brunner, and A. Feeney (2013): "Health inequalities among British civil servants: the Whitehall II study," in *Stress and the Brain*, Routledge, 61–67.

Mikolov, T., K. Chen, G. Corrado, and J. Dean (2013a): "Efficient estimation of word representations in Vector Space," *arXiv preprint arXiv:1301.3781*.

Mikolov, T., I. Sutskever, K. Chen, G. S. Corrado, and J. Dean (2013b): "Distributed representations of words and phrases and their compositionality," *Advances in Neural Information Processing Systems*, 26.

Nekoei, A., J. Sigurdsson, and D. Wehr (2024): "The economic burden of burnout," *CESifo Working Paper*.

Osnabrügge, M., S. B. Hobolt, and T. Rodon (2021): "Playing to the gallery: Emotive rhetoric in parliaments," *American Political Science Review*, 115, 885–899.

Pennington, J., R. Socher, and C. D. Manning (2014): "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.

Reimers, N. and I. Gurevych (2019): "Sentence-bert: Sentence embeddings using siamese bert-networks," *arXiv preprint arXiv:1908.10084*.

Rettl, D. A., A. Schandlbauer, and M. Trandafir (2022): "Employee health and firm performance," .

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

Rugulies, R., B. Aust, I. E. Madsen, H. Burr, J. Siegrist, and U. Bültmann (2013): "Adverse psychosocial working conditions and risk of severe depressive symptoms. Do effects differ by occupational grade?" *The European Journal of Public Health*, 23, 415–420.

Tausczik, Y. R. and J. W. Pennebaker (2010): "The psychological meaning of words: LIWC and computerized text analysis methods," *Journal of Language and Social Psychology*, 29, 24–54.

Wieclaw, J., E. Agerbo, P. B. Mortensen, and J. P. Bonde (2005): "Occupational risk of affective and stress-related disorders in the Danish workforce," *Scandinavian Journal of Work, Environment & Health*, 343–351.

World Health Organization (2024): "ATC/DDD Toolkit," https://www.who.int/toolkits/atc-ddd-toolkit, last accessed: July 7th, 2025.